

Ten Lectures on Mind and Language

By

Mark Turner

马克·特纳

心智与语言十讲

Copyright © Mark Turner 2009

These lectures were delivered by Mark Turner in May 2009, when he was the forum speaker for the Seventh China International Forum on Cognitive Linguistics. They were originally published as a book in the Eminent Linguists Lecture Series, 世界著名语言学家系列讲座。

Contents 目录

Summaries of the lectures in Chinese

Lecture One: Human Meaning

Lecture Two: Conceptual Integration

Lecture Three: Double-Scope Cognition

Lecture Four: Mental Packing and Unpacking

Lecture Five: Big Ideas

Lecture Six: Working in the Mental Network

Lecture Seven: Frame Blends

Lecture Eight: Blended Stories

Lecture Nine: The Nature of Language

Lecture Ten: Grammar

汉语导读

人类的意义

概念整合

双域认知

心理打包和解包

纵思千古

运用心理网络

框架整合

故事整合

语言的本质

语法

Summaries of the lectures in Chinese

汉语导读

第一讲

Human Meaning

人类的意义

作为整个系列讲座的第一讲，本讲的主要目的是介绍概念整合理论的背景。本讲内容大致可以分为两部分：第一部分是序曲，主要是对认知语言学研究现状的概述，第二部分正式进入本讲的核心话题——人类的意义。在第二部分中，Turner 教授对人类认知能力的特点进行了分析，并由此引出概念整合理论。

在本讲第一部分，Turner 教授首先肯定了世界以及中国认知语言学的发展现状与潜力，赞扬了认知语言学学会、学术期刊以及专业论坛等在学科发展中所起的积极作用，特别是“中国认知语言学国际论坛”(CIFCL)在学科发展中起到的巨大推动作用。之后他重点介绍了认知科学网(Cognitive Science Network,简称 CSN, <http://ssrn.com/csn/index.html>)。CSN 是社会学研究网 (Social Science Research Network,简称 SSRN, <http://www.ssrn.com/>) 的分支，是为研究者提供免费服务的非营利性网站。研究者可以从 CSN 上查询或下载各种论文，也可将自己的论文上传到 CSN。研究者上传到 CSN 上的论文不计为发表并保有对该论文的一切权利。CSN 接纳不同语言的论文，作者只需提供英文标题及摘要。CSN 为广大认知语言学研究提供了便捷的交流平台。

在第二部分，Turner 教授首先指出，人类拥有许多其他动物所不具备的独特能力，如艺术、音乐、宗教、语言、数学、高级工具使用、高级社会认知等。考古学证据显示，所有这些能力都出现于大约 5 万年前的旧石器时代晚期。从进化的尺度上看，5 万年是非常短的一段时间，抛却类现代人的漫长进化时间不算，单是解剖学上的现代人类就已存在了 15 万年，而认知上的现代人类却仅仅有 5 万年的历史。这一现实使人不得不思考以下问题：5 万年前的人类身上到底发生了什么变化？是什么使这些认知能力突然一起涌现？

关于人类特有的认知能力，神经学、生理学、认知科学等学科一直都在进行研究。但以往的研究往往采用分解法 (reductionism)，即将人类的认知能力割裂成不同的部分，并对不同的部分，如艺术、音乐、语言、数学能力等进行单独研究，以期在理解各个部分之后，便可得到关于人类总体认知能力的合理阐释。Turner 教授认为，分解法事实上存在谬误。他两次引用著名神经学家 V. S. Ramachandran 的论述来表明：对于某些复杂整体的研究，坚持分解法有多么荒诞可笑。而人类的高阶认知能力 (higher-order cognition)，正是这样一个不可割裂的复杂整体。

Turner 教授认为，人类的高阶认知能力都具有同一个来源，那便是概念整合 (conceptual integration)或称概念合成 (conceptual blending)。概念整合，尤其是其中的双域型整合 (double-scope blending) 可以形成全新的、创造性的层创结构 (emergent structure)，并使得人类可以超越

当时当地的尺度 (local scale), 形成抽象思维的能力。当人类进化出概念整合能力之后, 上述艺术、音乐、宗教、语言、数学、高级工具使用、高级社会认知等便在同一时间纷纷出现, 而且各个能力之间相互依赖, 相互协调, 共同发展。

Turner 教授以一个简单的例子初步介绍了概念整合的机制。一位伴郎在婚礼上突然想起自己远在他乡的女友以及三个星期之前同女友潜水的情景。这两个情景, 即婚礼和潜水, 是完全冲突的。两者的时间、地点全无相似, 参与者也大相径庭。然而进化却赋予了人类可以在同一时间并容并整合两个冲突场景的能力。那位伴郎并不会因此产生思维混乱, 相反, 他可以在两个不同情景的元素之间建立连接: 他可以将自己想象成为婚礼上的新郎, 将女友想象为新娘, 将现时新郎和新娘之间的婚礼想象成自己和女友之间的婚礼。当然, 他在脑中想象并运行这一情景的最终结果可能是女友在婚礼上说不会嫁给他, 又出现一个与现实中的婚礼完全相反的、崭新的情景。

最后, Turner 教授以一个谜语结束本讲: 一位和尚日出之际从山脚下往山顶上走, 日落之前到达山顶, 随后在山顶冥思。第二天他又于日出之际从山顶往山脚下走, 并于日落时分抵达山脚。这两次旅途中, 该和尚是否可能同一时间出现于同一地点? 该谜语的解答要依靠概念整合来得出。Turner 教授留下这一悬念, 留待第二讲以此为例, 深入介绍概念整合的机制。

第二讲

Conceptual Integration

概念整合

在第二讲，Turner 教授利用大量的实例向听众介绍了概念整合这一心理认知过程，以及概念整合网络如何在语言中运作。本讲的内容也是整个概念整合理论的基础。

首先，Turner 教授承接第一讲中“婚礼”的实例，深入讲解为何人类具备同时思考两件事情的能力。Turner 教授提出人类的记忆和概念整合可同时发生这样的假设，记忆可为概念整合提供良好的素材。在“婚礼”一例中，男子将自己和远在他乡的女友投射到了眼前的婚礼中，以自己 and 女友填充婚礼框架中新郎和新娘的角色，产生了一种整合性的模仿 (blended simulation)，并由此产生了他和女友举行婚礼的虚拟情景。

接着，Turner 教授以第一讲结尾时留下的“和尚之谜”为例，阐述了基本概念整合网络模式。解决“和尚之谜”需要利用概念整合的思维，我们可以假设山脚和山顶各有一个和尚同时向对方走去，那么他们相遇的那一地点则是该问题的答案。该整合网络具有四个心理空间 (mental spaces)，即两个输入空间 (input spaces)，一个类属空间 (generic space)，一个合成空间 (blended space or blend)。输入空间一包含和尚上山的时间、旅程、方向等结构，输入空间二包括和尚下山的时间、旅程、方向等结构。跨空间映射 (cross-space mappings) 在输入空间的对应成分 (counterparts)，如和尚、旅程等之间建立联系，并将之投射到类属空间，这样类属空间便包含了两个输入空间共有的抽象结构：行走的和尚、和尚的位置、山顶和山脚之间的旅程、旅程的时间以及没有特定方向的运动，等等。然后，输入空间的元素被有选择性地投射到了合成空间，经过组合、完善、扩展在合成空间内形成了层创结构：两个和尚分别从山脚和山顶出发并在途中相遇。

Turner 教授重点强调了整合中的三个问题并进行了分析。第一是选择性投射。在“和尚之谜”一例中，并不是输入空间中的所有结构都被投射到了合成空间中。例如，和尚上山和下山不是在同一天内，但这种时间上的冲突并没有被投射到合成空间。第二是框架的引入。当合成空间具备了一些元素后，我们会借助一些常见的框架来组织这些元素。如在“和尚之谜”中，上山的和尚、下山的和尚、山顶到山脚的旅程、上山的运动、下山的运动等等很容易使我们想起两人相遇的框架。于是该框架便被引入到合成空间中，成为合成空间中产生的不同于输入空间的层创结构。第三是使用适于整合空间的语言。例如，在“和尚之谜”中，我们可以说 “*The monk meets himself* (和尚遇到了他自己)”。由于概念整合的存在，我们可以很容易地理解这样一句在日常生活中不合逻辑的话。

概念整合具有不同的类型。本讲的重点是镜像型 (mirror network)。所谓镜像型，是指两个输入空间具有相同的组织结构。“和尚之谜”就是典型的镜像型概念整合网络，因为输入空间的组织框架都为“某和尚在某一天中沿着山路朝某一目的地行进”。本讲中的“盒子实验”、“田径单纪录图”以及“与康德辩论”等也都属于镜像型。另外，Turner 教授简

单讨论了双域型概念整合网络 (double-scope network)。他指出双域型是最为高级的一种整合方式，两个输入空间具有不同的组织框架，合成空间的组织框架同时包括来自两个输入空间的部分结构。例如《哈罗德和紫色暗笔》[Harold and the Purple Crayon]的故事，哈罗德 (Harold)用蜡笔画出的东西都能变为现实，于是他可以画苹果充饥，画月亮指引道路。输入空间一的画画框架和输入空间二中的现实世界框架完全不同，但是通过选择性投射，我们可以在合成空间中得出哈罗德用笔画出世界这一层创结构。关于概念整合的不同类型以及双域型概念整合，Turner 教授将在第三讲中进行详细介绍。

最后，Turner 教授简略介绍了概念整合的最终目标，即到达人类思维尺度 (achieve human scale)。压缩冗余成分 (compress what is diffuse)。获得全局顿悟 (obtain global insight), 强化重要关系 (strengthen vital relations)。产生故事 (come up with the story) 以及合众为一 (go from many to one)。另外，Turner 教授还谈到了概念整合的构建原则，如对应连接 (matching counterpart connections)、类属空间 (generic space)、整合 (blending)。选择性投射 (selective projection), 层创意义 (emergent meaning)。此外，Turner 教授还简单介绍了概念整合过程中的重要关系 (vital relations), 如变化 (change), 同一 (identity), 时间 (time), 空间 (space), 因果 (cause-effect), 部分整体 (part-whole), 表征 (representation) 等。同时，概念整合还要受到一些控制原则的制约，如压缩原则 (governing principles for compression), 拓扑原则 (topology principle) 等。以上这些目标、原则、重要关系等，Turner 教授会在下面的几讲中深入阐释。

第三讲 Double-Scope Cognition 双域认知

本讲内容主要分成两部分：第一部分，Turner 教授深入讲解了第二讲中提到的概念整合的不同类型以及各个类型间的区别与联系，第二部分，Turner 教授提出了关于概念整合以及认知研究的四个命题。

概念整合主要分为四种基本类型，即简单型 (simplex)。镜像型 (mirror)。单域型 (single-scope) 和双域型 (double-scope)。在简单型整合中，输入空间一包含特定框架及其角色，输入空间二包含无框架组织的元素，跨空间映射将角色与元素进行匹配。例如，在 *Paul is the father of Sally* (保罗是萨莉的父亲) 中，输入空间一包含的是在人类进化文明史中形成的“家庭”框架以及“父亲”、“女儿”的角色，输入空间二中包含的是 Paul, Sally 两个元素，跨空间映射在角色与相应的值之间进行匹配。

其余三种类型的两个输入空间都包含各自的框架，三者的差别在于框架兼容的方式不同。镜像型中，由于两个输入空间的框架相同，不存在冲突，因此两个框架可以直接融合。第二讲中的例子，如“和尚之谜”、“与康德辩论”等等都属于镜像型。单域型的两个输入空间在框架上有所冲突，合成空间的框架以其中的一个为主。例如，在 *Newton is the father of physics* (牛顿是物理学之父) 中，输入空间一是“家庭”空间，输入空间二是“物理学”空间，跨空间映射在牛顿与父亲、子女与物理学之间建立了对应连接，但合成空间的组织框架仅仅来自输入空间一，即“家庭”或者说是“父子”框架。双域型的两个输入空间也在框架上存在冲突，但与单域型不同的是，合成空间的框架由两个输入框架共同产生。例如，在 *The child is father of the man* (儿童是成年人的父亲三岁看老) 中，输入空间一仍然是“家庭”或“父子”框架，输入空间二是我们的背景知识框架——成年人都是由儿童长成的，都来源于儿童。合成空间框架从两个框架中分别提取了一部分。例如，合成空间的框架从输入空间一中提取了“父亲”、“孩子”的角色，从输入空间二中提取了时间顺序，即因为孩子成长为成年人，所以先有孩子，后有成年人。在四种类型的概念整合中，双域型无疑是最高级的一种，也是将人类与其他哺乳动物区别开来的重要认知特征。

关于整合网络的类型，Turner 教授特意强调了两点。第一，四种类型之间没有严格的界限。例如，*The Pope is the father of all Catholics* (教皇是所有天主教徒的父亲) 就介于简单型和单域型之间。第二，还有一些概念整合网络不属于以上任何一种，它们的两个输入空间框架既不相同也不冲突，可以很好地融合在一起。例如，*The last second violinist is always the spy* (最后一个第二小提琴手总是间谍) 的两个输入空间框架就不存在任何冲突。一个人既可以是“乐队”框架中的小提琴手，也可以是“谍报”框架中的间谍。

Turner 教授还提出了关于概念整合和认知研究的四个命题。第一，概念合成是基本认知模式，没有概念合成，就没有人类的高阶认知。第二，认知研究是一项艰巨的任务，因为认知科学家面对的是“后台认

知” (backstage cognition), 是进化刻意隐藏起来让我们注意不到的那部分。第三, 人类通过整合构建“故事”, 虽然“故事”同实际情况并不一致, 但是却是一种重要的认知方式。换句话说, 整合产生虚构 (fictivity), 而虚构在很多领域中都扮演着重要的角色。因此, 我们有必要重新审视真理、科学、认知等概念, 以及怎么做才是获得了真理。最后, 词语并不等同于好的理论。人类是命名主义者, 我们倾向于认为如果语言中存在一个词, 那么这个词的所指便理所当然地存在, 而且不同词的所指理所当然地不同。但 Turnei-教授指出, 词语仅仅是词语, 我们不应被词语所误导。尤其是在认知研究中, 我们应避免先人为地认为 categorization (范畴化), analogy (类比) 等都是基本层次的认知机制, 然后再努力寻找它们的不同。

此外, Turner 教授指出, 我们通常认为整合都是先由输入空间开始, 最后得到合成空间, 但事实是人们可以在整个合成网络的任何一个地方入手。例如, 根据 Chris Johnson 的研究, seeing (看见) 和 knowing (知道) 在人类早期是融合在一起的, 随后才慢慢分开。在这个例子中, 先有合成空间, 然后才逐渐分解成两个空间。

最后, Turner 教授分析了拉辛的戏剧《菲德拉》中的一幕。希腊英雄忒修斯 (Theseus) 前往克里特岛, 与公主阿里阿德涅 (Ariadne) 相爱, 在阿里阿德涅用线团协助下, 杀死了迷宫中的怪物弥诺陶洛斯 (Minotaur)。但由于酒神狄俄尼索斯 (Dionysus) 宣称阿里阿德涅是他的妻子, 忒修斯后来娶了阿里阿德涅的小妹妹菲德拉 (Phèdre)。忒修斯婚后四处游荡, 而菲德拉爱上了与她年纪相仿的继子——忒修斯的儿子希波吕托斯 (Hippolytus)。菲德拉用间接的方式向希波吕托斯告白说: “为什么不是你去了克里特岛? 你可以成为我姐姐救助的那个人。不, 我将成为救助你的那个人。一个线团是不够的, 我会和你一起进入迷宫, 无论是生是死都在一起。”通过对该故事的解析, Turner 教授说明了从输入空间到整合空间的整合过程实现了重要关系 (vital relations) 的压缩。例如, 输入空间中忒修斯与希波吕托斯、阿里阿德涅与菲德拉之间的类比关系 (analogy) 被压缩成了合成空间中的同一关系 (identity)。

第四讲

Mental Packing and Unpacking

心理打包和解包

本讲围绕心理打包和解包展开，主要分为四部分：第一部分，Turner 教授以“旅行”为喻说明心理打包和解包过程，第二部分，Turner 教授对比了关于人类语言能力的两种理解方式——提取与使用、打包与解包，指出后者具有更强的阐释力，第三部分，Turner 教授以大量实例对概念整合如何实现打包和解包作了更深入的分析；最后，Turner 教授以致使运动句型 (caused-motion) 和致使结果句型 (resultative) 为例，说明了打包与解包在语言、尤其是语法中的作用。

首先，Turner 教授以旅行为喻点明本讲的主要内容：打包和解包。人们旅行时不需要把所有东西都带在身上，只需一个简易的旅行箱，装上一些必需品，如牙刷、电源转换器、钱包等。旅行结束后，箱子可能会稍有变化，如添加了新的东西或改变了打包的方法等。Turner 教授认为人类的思维以及意义的建构正如旅行时的打包和解包。我们不是将所有的意义都放在脑中，而只是携带框架等认知必备品，然后在使用时将其解包。这些认知必备品在使用的同时得到改变、发展以及扩充等。

之后，Turner 教授对提取与使用、打包与解包进行了对比。提取与使用是关于人类语言知识和能力的另一种理解方式。在这种理解方式中，所有的意义结构都是存储在头脑中的。Turner 教授以日晷为例，说明了前者的谬误。日晷中阴影的产生是不同系统相互作用的结果，对阴影含义的理解过程不是从大脑中提取之前储存的阴影含义，而是根据已有的知识和认知能力，逐步解读出蕴涵在阴影中的含义，进而生成识别日-时间的能力。

接下来，Turner 教授阐释打包和解包的过程可以通过概念整合来实现。人类可提取输入空间事物的抽象结构，将之投射到整合空间中，并在大脑中存储整合空间的概念结构。例如，年、月、习惯等抽象概念，这些概念本来是不存在的，是我们将 365 天打包成一年，将大约 30 天打包成一个月，将每天重复进行的事情打包成习惯。一旦打包的过程完成，我们便不再需要整个网络，而是可以直接对由打包形成的概念进行操作。例如，在年、月、习惯等概念形成后，我们可以说 *In a leap year, we add a day to February* (每逢闰年，我们要在二月加一天) 或者 *I'm banishing my habit* (我在改掉自己的习惯)。此外，人们可以根据具体的情况，在解包和重新打包的过程中对原来打包在头脑中的概念加以改变。Turner 教授以美国加利福尼亚塔霍湖 (Lake Tahoe) 的野生熊为例进行了说明。游客凭以往在动物园等地的经验，形成“喂养行为对动物有益”这一观念。但对于野生熊，如果游客有意或无意之中喂食，那么熊就会形成在人类居住地附近寻找食物的习惯，由此造成人类对熊的恐慌，最终熊被人类射杀。因此，塔霍湖各处都有宣传画，通过 *A fed bear is a dead bear* (喂熊即杀熊) 这一醒目的标语来改变游客对喂养行为的观念。*A fed bear is a dead bear* (喂熊即杀熊) 是概念整合对因果链进行压缩的结果。也就是说，概念整合对“喂养行为”进行了重新打包。

最后，Turner 教授讲述了打包与解包在语言中的实际应用。例如，

致使运动从句是英语的一个基本类型：施动者发力，使受力物体按照一定方向运动。英语中存在这样的使役动词，如 *I throw the ball over the fence* (我把球扔过篱笆) 中的 *throw*。但使动句型的抽象结构形成后，我们可以将许多非使动动词嵌入这一结构中，如 *I read him to sleep* (我把他读睡了) 中的 *read*。在这个过程中，输入空间一是使役句型的结构，输入空间二是没有框架的元素，如 *I, him, read* 等。概念整合将分散的元素安放到框架的适当位置上，就形成了上述例句。

综上所述，本讲的核心是：人类通过不断的打包与解包再打包的过程认知世界，掌握知识，使用语言> 打包与解包是可以通过概念整合实现的。

第五讲 Big Ideas 纵思千古

在本讲中，Turner 教授首先回顾了第四讲“打包和解包”的主要内容。他强调：人们使用话语来表达意义，这其实是一个压缩后的概念，实际上话语并不具有意义，话语只是一种工具，一种提示（prompt），人们利用这些提示构建意义。我们构建的意义无限宽广，超越时间、空间的限制，也远远超越我们自身的局限。但因为我们可以建立概念整合网络，这些在时间和空间上跨度庞大的意义或概念（big ideas）可以被压缩在我们能够理解的人类尺度以内。第五讲即是围绕这个话题展开的。

其他物种的认知机能受时间和空间的局限，很少能够超越局部范围去理解事情。为何人类思想和感觉涉及长期的远程的因果模式呢？在此，Turner 教授借用了影视作品中频繁出现的时间机器来比喻人类的大脑。时间机器可以在整个宇宙和任何年代穿梭，人类的大脑就是这样一台时间机器。与此同时，我们拥有发达的概念整合能力，可以压缩超越自身局限的事物。例如，在一部呼吁环保的电影《难以忽视的真相》(*An Inconvenient Truth*) 的结尾部分，影片中播放的是从 40 亿英里之外看地球的情景，同时伴有画外音：

Everything that has ever happened in all of human history has happened on that dot. All the triumphs and tragedies, all the wars and all the famines, all the major advances. That is what is at stake—our ability to live on planet Earth, to have a future as a civilization. Future generations may well have occasion to ask themselves, "What were our parents thinking? Why didn't they wake up when they had the chance?" We have to hear that question from them now.

人类历史上发生的所有事情都发生在那个点上，包括所有的胜利与悲剧，所有的战争与饥荒，所有的重大进步。但是现在，我们能否继续生活在地球上，我们的文明是否还有未来，都是个未知数。我们的后代可能会问自己：“我们的父母当时在想什么？他们为什么不在还有机会时醒悟呢？”我们必须现在就聆听这个问题。

以上涉及到在现实世界中我们无法把握的距离和时间，即 40 亿英里和未来几代乃至几十代。但是概念整合却可以把距离和时间压缩，以至于在合成空间中，40 亿英里仅仅是眼前的距离，几代乃至几十代的时间间隔完全消失。

概念整合的这种压缩能力也是人类得以构建事物及统一体的基础。例如，评论家 Hugh Kenner 在谈论诗人 Marianne Moore 的诗歌 Poetry 时形容这首诗是“scarred by all those revisions (因屡次修改而留下伤痕)”。在此概念化过程中，所有在这个标题下发表的不同版本的文字被压缩成了一个元素，具有一个唯一的统一体，即 Poetry 这首诗。在有了统一体之后，我们可以说“这首诗经过屡次修改”。我们甚至可以将其进一步隐喻化成生命体，于是便出现了诸如“The poem is scarred by all

those revisions (这首诗因屡次修改而留下伤痕)”之类的表达。人类自身的统一体构建也是如此。概念整合将处于不同时刻的人类个体相同的那部分压缩成统一体，不同的部分压缩成变量。我们将过去的记忆与现时的情感系统整合，在整合空间内让现时的情感系统对过去的事件作出反应，这样便可以在整合空间构建“过去的自己”的概念并与“现时的自己”相联系。整合也同样适用于“未来的自己”。我们将年龄经验的框架与现时的情感推理整合，就得到一个延展了的自己。正是因为有一个统一的、连续的统一体，我们才会有“惩罚”、“报复”、“救赎”等能够联系一个统一体在不同时刻所做的两件事的概念。

最后，Turner 教授简单谈及了人类如何识别其他人思维的问题，并就概念整合理论与 Tomasello 的理论在此问题上的阐释进行了对比。Turner 教授认为人类可以用特有的双域整合能力来猜测其他人的想法，甚至虚拟动物等的思维，为无生命的事物安加意图等。Tomasello 的理论则侧重于社会认知，他认为社会认知是区别人类和其他动物的关键，也是人类识别其他人思维的基础。Turner 教授认为两者的理论并无矛盾之处，而概念整合理论则更深一步，因为在概念整合理论体系中，社会认知也是概念整合的结果。

第六讲
Working in the Mental Network
运用心理网络

本讲延续 Turner 教授前几讲的风格，以具体例子阐释理论。首先，他从双及物动词 (ditransitive) 入手。下面例句 (1) 中的 hand 是一个典型的双及物动词，语法结构为“名词+动词+名词+名词”，表达的是一个基本的人类场景 (human scene), 即将某物进行物理传递的过程。例句 (2), (3) 是同样的结构，表达的同样是转交某物的过程，但是意义上却有些许差异：例句 (2) 的转交过程历时较长，例句 (3) 中交递的并非实体物质，甚至句中的“她”自己原本就没有“头痛”这种东西。

(1) *I handed her the eraser.* (我把橡皮递给了她。)

(2) *She bequeathed him a farm.* (她把农场遗赠给了他。)

(3) *She gave him a headache.* (她让他感到头疼。)

Turner 教授认为，像例句 (2), (3) 这种描述非人类基本尺度情景的句子之所以也能够使用例句 (1) 的句式，完全是因为概念整合的结果。概念整合可以创造压缩，而已经压缩的结构可直接用于输入空间，创造新的整合。Turner 教授认为，例句 (1) 中的“交给某人某物”的场景被压缩后，生成了固定的框架和语法结构“名词+动词+名词+名词”。此框架和语法结构可以用来构造一个输入空间，再与另一个输入空间中的元素进行整合，就可以形成像 (2), (3) 等句子。这与第四讲中提到的致使运动结构的整合是同一个道理。

Turner 教授指出，他与 Fauconnier 教授近年来愈加意识到，压缩是概念整合最至关重要的一个目标。关于输入空间元素之间的重要关系如何被压缩进而转化成整合空间内的新关系，Turner 教授在此前的很多例子——如 *fed bear is a dead bear* (喂熊即杀熊)——中都有详细的解释。本讲中 Turner 教授着重讲解了之前未涉及的一种新压缩方式，即将相异 (disanalogy, 或译为“非类比”) 压缩成特征 (feature)。例如，对于 *Put the green tea in front of the missing chair* (把绿茶放到缺失的那把椅子前) 一句，Turner 教授的解释如下：想象一下，一个圆桌旁目前摆放了五把椅子，但是之前有六把。这样我们在头脑中会形成两个心理空间，我们可以注意到两个心理空间的相异，并利用概念整合将这种相异压缩成合成空间的“缺失”。在合成空间中，“缺失”便成为了如同“木制”、“金属”等一样的特征。事物具有某种特征是一个基本的人类场景，在我们的把握尺度之内。我们用“形容词+名词”来表达事物的特征。因此，便如同 *wooden chair, metal chair* 一样，我们可以说 *missing chair*。将相异压缩成特征在生活中十分常见，由此形成的词组也数不胜数，例如 *money problem, security problem, insulin coma, arousal problem, insulin death, food emergency, honesty crisis, rice famine*, 等等。

本讲的最后，Turner 教授结合具体例子指出了概念隐喻理论与概念整合理论相比的不足之处，即概念隐喻理论过于简单化。Turner 教授认为，绝大多数的隐喻都不能以从源域到目标域的简单映射来解释。他以历来被认为是最典型的隐喻 *time is space* 说明了这一点。下面两个句子中，句 (4) 完全正确，而句 (5) 却不合乎语言习惯，甚至让人无法理

解，可见时间域具有自己额外的结构。

(4) *Minutes are quick but hours are slow.* (分钟快'小时慢。)

(5) **Inches go by faster than feet.* (英寸比英尺快。)

Turner 教授认为，我们对于时间的理解具有两个输入空间。输入空间一是空间结构，输入空间二是我们的个人感觉。因为有输入空间二的存在，才会存在 *The eight-hour work day is longer on Monday than it is on Friday* (星期一的八小时比星期五的八小时长) 等表达。也就说，时间域的结构不仅仅来源于空间域，还来源于个人感觉域。因此，概念隐喻理论的源域一目标域解释在绝大多数甚至可能是全部的情况中都不能适用。相比之下，概念整合理论涉及多个空间多次映射，更具阐释力。

第七讲 Frame Blends 框架整合

Turner 教授主要围绕框架整合概念展开第七讲，他列举了大量丰富多彩的实例，包括诗歌、笑话和小说等，说明了框架整合在日常生活中的普遍性、重要性和基础性。

此前几讲中，Turner 教授一直在谈论概念整合如何使人类的各种高阶认知能力成为可能。从本讲开始，Turner 教授将概念整合与人类其他的两个能力相结合，即框架化的能力和构造故事的能力。本讲以框架整合为主，第八讲以故事整合为主。框架是我们用于指导行为的一个组织结构，例如图书馆的框架包括书、图书管理员、房间以及在图书馆的操作程序等。Turner 教授首先以他多次提及的“婚礼”之例简要说明了什么是框架整合。一位伴郎在婚礼之上突然想起自己远在他乡的女友以及三个星期之前同女友潜水的情景。婚礼和潜水是两个完全不同的框架，但是伴郎可以在两个不同框架的元素之间建立连接：他可以将自己想象成为眼前婚礼上的新郎，将女友想象为新娘，将现时新郎和新娘之间的婚礼想象成自己和女友之间的婚礼。这就是所说的框架整合。

某些框架整合在全世界都是非常普遍的，但是另一些却是极具文化色彩的，因此框架整合是基于文化的，可以在一种文化中非常具有影响力而在另一种文化中根本不存在。例如，Turner 教授分析的美国笑话 *Why God won't get tenure*。（上帝为什么不能得到续聘？）就带有极其浓重的文化色彩。这个笑话根据美国大学的聘任制度列出了 15 条理由为什么上帝不能得到续聘，如他只发表过一项著作，而且是希伯来语写作，书中没有任何参考文献，等等。

随后 Turner 教授探讨了框架整合如何创造新的语言或者新的语言使用方式。例如，Catullus 的诗歌 *Elegy 101* 中的主人公可以称他兄弟的骨灰为“兄弟”。“兄弟”一词的这种用法，显然不是因为某本词典列出该词有“骨灰”的意思，而是因为我们通过框架整合将活着的兄弟与他死后的骨灰对等起来。再比如，*jail bait* 是一个较近出现的合成词，用来指代未成年但很具吸引力的少女。其中 *jail* 来自犯罪框架，*bait* 来自钓鱼框架，我们将两个框架整合之后，*jail bait* 就可以用来表示“引诱成年男子犯罪的少女”。

Turner 教授强调，尽管他列举了很多诗歌、笑话的实例来说明框架整合，但是框架整合也频繁出现于我们对于严肃重大事件的理性思考中，帮助我们作出判断。例如，曾经有一个关于二战的著名论断：*If Churchill had been prime minister in 1938 instead of Neville Chamberlain, Hitler would have been deposed and World War II averted.*（如果 1938 年时英国首相是丘吉尔而非张伯伦，那么希特勒就会被除掉，二战也不会发生。）此后，撒切尔夫人在应对前南斯拉夫问题时，就借用过这个论断，并称 *We need to be just like that in Yugoslavia*（我们在南斯拉夫问题上也该如此），意在说明不能对南斯拉夫采取绥靖政策，而该迎头痛击。

此外，Turner 教授总结了关于多义词的四个原则。第一，用于输入空间的表达也适用于合成空间，并用来指代对等物。例如上文称骨灰为

“兄弟”的例子。第二，在输入中的表达组合可用来在整合中选取结构，尽管这些组合对于输入来说是不合适的。例如，在对负数的框架和平方根的框架进行整合得到虚数的框架之后，我们可以在整合中使用“负一的平方根”这样的表达，尽管该表达在原来的两个输入空间，即负数空间和平方根空间，都是不适用的。第三，我们总可以为合成空间中的层创结构找到词语表达，尽管这些词语不能运用到输入本身。例如，第二讲“和尚之谜”中提到的表达 *The monk meets himself* (和尚遇到了他自己)，该表达完全适于层创结构，但是却不能用于输入空间。最后一条原则是对以上三条的总结，即整合不可避免地扩展了词语的使用，但是我们几乎感觉不到这些扩展。

最后，Turner 教授谈到了虚构现象，特别是 Talmy 虚构运动理论中的框架整合现象。他认为概念整合理论可以解释虚构运动。例如，在句子 *The mountain range goes all the way from Mexico to Canada* (这条山脉从墨西哥一直延伸到加拿大) 中，输入空间一的框架为某物体从源头到目的地沿着一个路径移动，输入空间二的框架为山脉，经过框架整合后，我们就可以得出上句所表达的虚构运动。

第八讲

Blended Stories

故事整合

本讲的主题是故事整合。Turner 教授通过大量的实例展示了人们如何将不同故事的人物整合在一起，形成新的意义。

我们曾在第三讲中探讨过的拉辛的戏剧《菲德拉》(Phèdre) 就是一个整合故事的例子。其中有两个独立的故事和两组人物。在原来的故事中，忒修斯 (Theseus) 去了弥诺斯宫，在恋人阿里阿德涅 (Ariadne) 的帮助下杀死了弥诺陶洛斯 (Minotaur)。而在菲德拉 (Phèdre) 对故事的整合中，希波吕托斯 (Hippolytus) 和菲德拉分别充当了忒修斯和阿里阿德涅的角色，成为了一对恋人，尽管在现实的故事中，他们是继子与继母的关系。菲德拉利用故事的整合达到了她的目的，间接地向希波吕托斯表白。

故事整合还可以发生在两个以上的独立故事之间。例如，威廉-巴特勒-叶芝 (William Butler Yeats) 的诗 <《在学童中间》 ("Among School Children")。该诗以一位花甲老人的口吻，描述他以政府官员的身份重返学校的所思所想。开篇讲述老人访问学校，面对一群孩子，这是第一个故事。接下来老人回忆起他曾深爱的一个女子，这是第二个故事。然后老人回忆起他所深爱的女子讲述的她儿时的事情，这是第三个故事。老人将所有这些整合在一起，想象曾经的恋人孩童时期的样子，并将其与眼前的女孩重合，于是诗中就有了 And thereupon my heart is driven wild / She stands before me as a living child (想到此，我的心简直就像发了疯癫 / 她仿佛一个活生生的孩子站在我面前) 这样的诗句。

故事整合可以多重整合混杂在一起，非常复杂。例如《夏洛特姨妈和国家美术馆的画像》(Aunt Charlotte and the NGA Portraits)。故事的主角之一是一只名叫奥尔佳 (Olga) 的“塞尔克” (selkie)。塞尔克平时是海豹的样子，生活在海里，但是当它们脱下海豹皮衣服就会变成人形。如果男人抱走了女塞尔克的衣服，她们就会嫁给他，直到拿回衣服为止。奥尔佳的衣服被一个男人抱走，但她不愿嫁给那个男人。男人将衣服藏在一幅画中，画上有一条大运河，奥尔佳没有衣服掉到水中就会淹死，因此她自己无法取回衣服。后来一个孤独的孩子夏洛特和奥尔佳成为了朋友，她在国家美术馆其他画像中的孩子们的帮助下，进入了藏着奥尔佳衣服的那幅画，帮助奥尔佳取回了衣服。最后奥尔佳重返大海，夏洛特自己也长大了，上了大学并嫁给了一个有钱而且很爱她的丈夫。整个故事是已经年迈的夏洛特向她的小侄孙女玛格丽特 (Marguerite) 讲述的，玛格丽特与她当年年龄相仿并同样感到孤独无趣。

在这个故事中，奥尔佳的世界里有一个图画般的世界，不同图片中的人可以相互走动。故事将图画世界和奥尔佳的世界进行了整合，同时将夏洛特和玛格丽特进行了整合，而最大的整合则是十来岁儿童在阅读故事的时候，与书中主人公的一个整合。这些多重整合都是人在不知不觉的情况下完成的。

Turner 教授最后重点分析的故事选自莎士比亚的历史剧《亨利六世》(King Henry the Sixth)。贵族塔尔博特 (Talbot) 与法军血战，分别七年

的儿子约翰违背父亲的意愿来到战场，最终战死。文中有一句 **Had death been French, death has died today.** (如果死神是一个法国人，那今天就是他的死期)。这里首先涉及到死亡的拟人化过程，该过程为先从具体的各种原因导致的死亡，如老死、病死和事故致死等等中抽象出来的一般的、抽象的死亡，然后再通过整合形成人格化的死神。这是一个常规化的固化整合。之后，再将人类与死神之间的争斗和战士与对手之间的争斗进行整合，通过对应元素连接，通过选择性投射等就形成了上述例句中的结构。

本讲最后，Turner 教授简要提及了其他几个故事以及油画中的整合。所有这些实例都向我们展现人类概念整合的能力和构造故事的能力之间的整合，以及这种整合产生的影响。

第九讲

The Nature of Language

语言的本质

Turner 教授在第九讲主要解释的问题是语言的本质及起源。

首先 Turner 教授归纳了在语言本质和起源问题研究中经常出现的两个误区：一是认为某种显著、突然的后果一定是由于某种显著、突然的原因造成的，而非渐进的原因。因此，语言的起源一定是因为人类在身体和大脑构造上突然发生了某种巨大的变化。第二是对功能与器官之间关系的误解。例如，当我们看到负鼠 (opossum) 用尾巴挂在树上时可能会认为，尾巴的进化是为了能让负鼠挂在树上。但事实上，早在负鼠能挂在树上之前，它们的长尾巴就已经在那里了。

然后，Turner 教授指出了解释语言起源的理论应该具备什么特点。他认为，一个好的解释语言的理论应该：(1) 认识到语言的奇特性：没有证据表明语言能力的进化存在中间状态，世界上不曾存在初级的、不完备的语言系统；(2) 不承认语言能力是由某种突然的巨变引起的，即避免上述的第一个误区；(3) 为语言的发生是持续的进化所带来的结果，且进化之路上的每一个变化具有自适应性；(4) 解释什么样的心智活动在这条道路上以什么顺序发展；(5) 解释什么持续的变化带来了怎样的独特性以及如何带来这种独特性；(6) 具有多方证据证明人类的确在此假设的道路上进行心智活动；(7) 解释语言能力不存在中间状态，而引起语言产生的认知能力却是连续的，存在中间状态；(8) 具有支持这些中间步骤的现代人类解剖和行为证据；(9) 解释其他相关的人类特有的能力也是在同样一条持续进化的道路上产生的。

Turner 教授认为，用概念整合理论解释语言的起源符合上述很多特点。他认为是双域整合能力导致了语言的产生。双域概念整合是四种概念整合类型（简单型、镜像型、单域型、双域型）中最高级的一种，是人类区别于其他物种的特征，也是语言、艺术、科学、高级工具使用、社会认知、数学、音乐等人类高阶认知能力的基础。双域整合可以把分散的、不可操控的意义压缩成有效的、可理解的、深刻的、使人类可以立刻识别的基本场景。整合能力的发展是渐进的，每一步都具有自适应性。有充足的证据可以证明整合能力的发展是有中间状态的，例如有些动物只具备简单型整合能力。而人类目前既有双域整合能力，也有简单整合能力。因此，双域整合不是一个剧变，而是整合能力连续发展的结果。

虽然概念整合是连续的，但它带来的结果却有剧变性。就好像我们将一块冰放在零下 20 摄氏度的房间里，然后一度一度地逐渐提高温度，只有当温度达到零度时，冰才会开始融化。语言是整体出现的，它是一种新的行为，一旦整合能力发展到关键的双域整合阶段便会出现。语言要么完整出现，要么完全不出现。文化不会停止在“简单”语言阶段。一个语法体系一定会是一整套的体系，可以组合以表达任何情景。语言起源背后的机制包含中间阶段，但是这些阶段在语言里是找不到的，因为一旦达到类键阶段，完整的语法就会作为整合的整体产物快速出现。

“快速”不是指瞬间出现，说它“快速”是因为它可以以文化尺度

计量，而不是以进化尺度计量。语言从深层意义来讲，必须是有潜力的，必须能够表达我们遇到的新情况。唯一的可能就是人类大脑能够将这些新情况和我们已经知道的结合起来，之后产生一个可以理解的整合以及相应的语法形式，以便使这些语法形式能够表达这些新情况。我们不需要发明新语法，我们需要构想一个整合，让已经存在的语法发挥作用。我们也不需要很多新词，而是构建一个整合，让已有词汇可以表达更多含义。只有这样，一个人才能用相对固定的词汇和基本语法形式表达无限丰富、永无止境的世界。关于这一点，Turner 教授已经在前面的第六讲和第七讲中作过说明。

最后 Turner 教授以一个文字体系背后的概念整合为例结束了本讲。一名女子在阅读丈夫从前线寄来的家书时，从一个角度上看，她只是在看纸上的符号。其他动物，如狗、鸽子经过训练之后也完全可以办到。但是我们知道她做的事情与狗和鸽子所做的事情是完全不同的。这里面涉及两个输入空间，输入空间一是一个女人在独自看一些符号，输入空间二中有她的丈夫以及他丈夫拥有的讲话、交谈的能力。经过整合之后，合成空间中形成层创结构：她在与她的丈夫交谈。当然，这是一种特殊的交谈，她丈夫不能回答她的问题。但是她丈夫可能会想象出她的问题，并在信中写“你一定想知道我什么时候回来”之类的话。总之，我们之所以能将一连串的符号看成词，是由于我们可以将这些符号与人类讲话时的声音整合，也就是说，概念整合使语言系统和文字系统成为可能。

第十讲
Grammar
语法

作为讲座的最后一讲，本讲的核心是语法问题。首先 Turner 教授回顾了前面各讲中一直强调的一个论点，即语言本身并没有意义，但它可以使我们进行整合，从而产生意义。换言之，语言是提示整合的系统。此外，语言也是复杂动态系统 (complex dynamic system) 和复杂适应系统 (complex adaptive system)。上述语言观是本讲内容的理论基础。

本讲中涉及的语法结构数量较多，Turner 教授将其分成了八大类，分别是：

- (1) Yof 结构，例如：*Paul is the father of Sally.* (保罗是萨莉的父亲。)
- (2) 需要依靠构建虚拟来理解的词汇，如 *safe*
- (3) 致使运动结构，例如：*I throw the ball over the fence.* (我把球扔过了篱笆。)
- (4) 致使结果结构，例如：*He boiled the pan dry.* (他把锅给煮干了。)
- (5) 双及物致使结构，例如：*I handed her the eraser.* (我把橡皮递给了她。)
- (6) 名词组合，例如：*jail bait*
- (7) 形+名组合，例如：*missing chair*
- (8) 一个词内的词汇组合，例如：*Chunnel*

以上八类中，第 (2) 至 (7) 类 Turner 教授在前面各讲中已经有所涉及或进行过详细讲解，本讲的重点是第一种结构。Y of 是指以下例子中的结构：

- (9) *Paul is the father of Sally.* (保罗是萨莉的父亲。)
- (10) *Necessity is the mother of invention.* (需要是发明之母。)

准确来说，该句法结构可概括为“*X be Y of Z*”。无论是简单的例句 (9) 还是隐喻性的例句 (10)，该结构在我们头脑中唤起的都是同样的映射图式和组合方式：即 X 和 Z 确定了输入空间一中的元素，Y 确定了输入空间二中的元素，be 动词确定了 X 与 Y 的对应连接。我们所要做的是确定输入空间二中的另一个元素 W 并与 Z 进行对应连接，从而完成概念整合的过程。Yof 结构具有递归性，即一个 Yof 表达之后再接另外一个 Y of 表达，例如 *Ann is the boss of the daughter of Max* (安是马克斯的女儿的老板)，重复这些表达就是要求我们重复映射操作，上一次概念整合的结果可以充当下一次概念整合的输入。

第 (8) 类中的结构在之前各讲中没有提及。所谓一个词内的词汇组合，是指如 *Chunnel* 由 *Channel* 和 *tunnel* 两个词，取前一个的头和后一个词的尾而合成。这种形态上的合成能够提示框架上的合成。值得注意的是，词汇的形态合成往往依赖于语言中已有词汇的形态特点。例如，因为法语中英吉利海峡和隧道分别是 *La Manche* 和 *Tunnel*，所以同样是英吉利海峡的海底隧道，在法语中就没有用类似 *Chunnel* 的一个词来表示，而是用一个名词词组 *Tunnel sous La Manche*。

此外，Turner 教授特意提到了 Mandelblit 关于希伯来语中 *binyan* 的研究。*binyan* (复数 *binyanim*) 是一种结构体，希伯来语中每个动词都

是将三到四个辅音根与 binyan 结合形成的。不同的 binyan 具有不同的意思。例如, hif'il 可以令动词带上使动义, 如下例所示:

(11) Hamefaked heric et haxayalim.

the-commander **run-hif'il**_{past} direct-object-marker the-soldiers.

The commander made the soldiers run. (司令官让士兵跑步。)

如此, 英语中的句法整合在希伯来语中可以用词汇整合来实现。

最后, Turner 教授谈到了普遍语法的问题。Hauser, Chomsky 和 Fitch 于 2002 年写的文章 "The faculty of language: What is it, who has it, and how did it evolve?" 中提到, 我们需要区分广义的语言官能和狭义的语言官能。广义的语言官能包括运动系统、认知一意愿系统、递归的计算机制, 而狭义的语言官能仅包括递归, 而且是语言官能中唯一的人类特有的组成部分。他们进一步提出, 甚至狭义语言官能都可能是由于语言之外的原因而进化的。

Turner 教授认为, 递归是语言中的普遍现象, 从简单的 *Ann is the boss of the daughter of Max* (安是马克斯的女儿的老板) 到复杂的 *[that was panned by the reviewer [who was kissed by the actress [who was escorted by the director [who was insulted by the reviewer]]]]* ([被评论家侮辱的[导演护送的[女演员亲吻的[那个评论家所批评的电影]]])。而语法获得递归性的根本原因在于我们的认知具有递归性, 而这种递归性又是通过概念整合的叠加, 即上一个整合的结果可以作为下一个整合的输入来完成的。但 Turner 教授强调概念整合绝不是普遍语法 D 概念整合是操作, 人类都具有这种能力。概念整合仅仅是一个非语言特有的普遍语法 (non-uniquely-linguistic universal grammar) 的一部分。人类的语言还涉及到许多其他东西, 如注意力、记忆力等等

Lecture One Human Meaning

Nihao.

Let me say how wonderful it is to be here in a nation with such a long and distinguished history and tradition of scholarship in language, such wonderful traditions of philosophy and linguistics and philology and research. I am especially grateful to Professor Li Fuyin for organizing this, so everybody, let's clap for Li Fuyin. And I am equally grateful to all these wonderful volunteers. To prove to you that I do not speak Mandarin, I am going to try to read their names, and I want them to stand up so that we can acknowledge them: Hu Yanan, Chen Mo, Li Lingmin, Ma Sai, Wang Fan, Wu Shan, Xiong Liqin, Yang Jie, Yin Shuying, Yuan Wenjuan, Zheng Lingyan. Thank you all for the wonderful help that you have given us and for making this conference possible.

But I also want to thank all of the participants here today, because it's quite amazing that you here. This is Dragon Boat Festival weekend! And I am told that this is the very oldest of the holidays observed in China. You could be home, all of you, eating the *zongzi*, and drinking yellow wine, right? But no, you have come to study Cognitive Linguistics here at Beihang University. So we must be particularly respectful of the sacrifice you are making to come here, and I will have to try to give good lectures in honor of QU Yuan, the great poet who was a great student of language and a fabulous practitioner of language and also a very dedicated man. I'll do my best in the lectures. Otherwise, perhaps the dragon will come eat me.

It is very important for us to establish an international community in Cognitive Linguistics, to build on the work that has gone before. China has become a great leader in Cognitive Linguistics in its national efforts and through the efforts of the participants in this audience and in other capitals and provinces in China and other cities. We are hoping to develop more programs for students throughout the world and to create a broader basis for research and collaboration. I hope that maybe someday—since Beihang University and Professor Li Fuyin have become such great leaders in Cognitive Linguistics—we will see a doctoral program located here in Beijing, which would be one of the great things to celebrate in Cognitive Linguistics throughout the world.

At Case Western Reserve University, where I hold my academic appointment, we have a program. We offer an MA in Cognitive Linguistics. This is the first year that we have offered it. We have our first class of students. They come from all around the world and they work with us. We have visitors who are postgraduate students, and also, we have visiting professors who work with us and who come to work in our workshop. They come from Brazil and Portugal, from Denmark and Germany, from Spain, from

all over. And we have two Chinese scholars who are working with us in a workshop now at Case. In fact, we have quite a large community of students and researchers already and many more applications for admission. It's possible to come to study Cognitive Linguistics at Case Western Reserve University at the graduate level, or, as you call it, the postgraduate level, for a Master's degree. If you inspect the Case website under Cognitive Linguistics, you will find the program and what it does and what the curriculum is and how people study. We of course use in this program the wonderful *Oxford Handbook of Cognitive Linguistics*, co-edited by my distinguished colleague here in the audience, Dirk Geeraerts. I am very grateful to him, too, for having done such a superb job in giving a broad picture of the role of Cognitive Linguistics.

The question that we are here to try to answer at the beginning is: What is Cognitive Linguistics? What is language? What is the nature of this enterprise? And I would say that, in a way, I think the progress so far in cognitive linguistics is almost sufficient to let us drop the word *cognitive* from *Cognitive Linguistics*, because we see the great influence of Cognitive Linguistics throughout the study of language, and its ability to combine a number of different fields.

In Cognitive Linguistics, broadly, the view is taken that language is a branch of cognition; language is a branch of mind, not a separate, partitioned modular activity that can be studied on its own, which of course it cannot. Leonard Talmy, for example, who was a speaker in this series, says at the beginning of his book *Toward a Cognitive Semantics* that he doesn't quite see why he is using the word *cognitive*, since of course semantics is always cognitive, so he should just have called it *Toward a Theory of Semantics*. Practically, we can't just drop the word "cognitive" yet, but we can continue to push even further the basic principle that language is a branch of mind, not an autonomous system.

This view of language is not surprising when you think of it. Evolution is a tinker. The great evolutionary biologist Francois Jacob described evolution as a tinker. A tinker finds things that are available and builds on them, modifies them just a little, tries to fit them together in order to new things. Of course, operating that way may create emergent structures; it may lead to emergent phenomena. But evolution does not start from scratch. Or at least someone who maintains that evolution has started from scratch for a particular human activity, like language, has a great deal of explaining to do, because that is not the principal way that evolutionary biology works. There are many different evolutionary mechanisms and most of them support this view of evolution as a tinker.

We are here to talk about cognitive science, Cognitive Linguistics, and the study of the distinctiveness of human beings. One way to look at the distinctiveness of human beings is to observe all the things that they can do that other species can't do, and then try to figure out how each one came about and to take each one—art, music, language, fashions of dress, advanced social cognition, advanced tool use, and so on—and then to break them into parts and to try figure out how each of those parts works, independent of the others, and hope that a human being is a linear sum of all these little separate parts; that a human being is a combination of all these little parts.

I argue that that is in many ways a questionable way to approach the study of human cognition. The quotation that I will read to you to give you a sense of my perspective on this comes from the great neuroscientist V. S. Ramachandran, in *Phantoms in the Brain: Probing the Mysteries of the Human Mind*. He wrote:

For over half a century, modern neuroscience has been on a reductionist path, breaking things down into ever-smaller parts with the hope that understanding all the little pieces will eventually explain the whole. Unfortunately, many people think that because reductionism is so often useful in solving problems, it is therefore also sufficient for solving them, and generations of neuroscientists have been raised on this dogma. This misapplication of reductionism leads to the perverse and tenacious belief that somehow reductionism itself will tell us how the brain works, when what is really needed are attempts to bridge different levels of discourse. The Cambridge physiologist Horace Barlow recently pointed out at a scientific meeting that we have spent five decades studying the cerebral cortex in excruciating detail, but we still don't have the foggiest idea of how it works or what it does. He shocked the audience by suggesting that we are all like asexual Martians visiting earth who spend fifty years examining the detailed cellular mechanisms and biochemistry of the testicles without knowing anything at all about sex.

This is the sorrow of the reductionist path in the study of higher-order human cognition. Human higher-order cognition is a great complex, but not, I think, a linear sum of little things.

We have heard the objection even today in this room: Is it not ambitious to try to explain all of these higher-order human behaviors at once? The view I offer is that a human being, when it comes to higher-order cognition, is a seamless whole. A human being does not have one brain when it is doing mathematics, a different brain when it is speaking language, a different brain with different abilities when it is engaging in social cognition, a different brain when it is walking.

A human brain, when it comes to higher-order cognition, is a seamless system, and human higher-order cognitive abilities come up together; they come up together in the infant. The infant develops these higher-order abilities simultaneously. It may well be that it would be impossible to learn these individual behaviors if you had to learn them separately. But they all help each other. Social cognition helps language, walking helps social cognition, vision helps social cognition, gesture helps language. All of these things work together. It is not clear that it is easier to explain them separately than it is to explain them as a unified set of abilities that human beings have, that they developed evolutionarily, phylogenetically, and that they develop in a lifetime ontogenetically, and this is indeed what we are going to look at.

It's not so difficult these days to find these sorts of views expressed very widely in the study of language. Leonard Talmy's work, for example, has become quite popular and routinely referred to in successors to transformational-generative, principles-and-parameters linguistics. Norm Chomsky himself, in a *Science* article with W. Tecumseh

Fitch and Marc Hauser, in 2002, talking about what might be exclusive to language, says that there might be only one thing that is exclusive to language, only one linguistic operation that is not shared with many other cognitive abilities, and it's recursion—our old, classic recursion. But then he says in the article, maybe not even recursion, because recursion can come from navigation behavior in animals that need to search for and locate food.

Of course, that *Science* article is part of a great debate. You are all familiar with that debate. But the idea that language is a branch of higher-order human cognition—that language draws upon and cooperates with other human abilities—is by now no longer a minority position. It's very common across the study of linguistics generally.

In fact I wonder whether, in 50 years, people will not be wondering what the debate was about anyway; just where did these people disagree? The disagreement has never been about whether human beings are uniquely equipped for language or whether this is biological. Of course, in some sense, they are uniquely equipped for language and of course it is biological. What else could it be? If you have a frog and a human baby, the frog will never learn language no matter what you do. The human baby will learn language if it is born into anything like a human community, and this is naturally because of their biological endowments. There is no doubt about that. The question is rather whether or not our ability for language draws upon other cognitive capacities, and other mechanisms, or is an autonomous capacity, and I think the evidence you hear presented in these lectures and the research broadly in linguistics is moving and has moved every year, every five years, more and more in the direction of recognizing that human beings are complex adaptive systems that develop a great range of capabilities, and that it is of course natural and to be expected that higher-order human capacities, like language, draw upon capacities that are deep in the evolutionary line.

Programs in Cognitive Linguistics are advancing, and now we need a place where we can publish all of this research and your papers and all of the research that is taking place in the world. Publication should be quick, free, open-access. It should be arranged so that the authors—you—retain the copyright. We have some fine journals, of course. We have *Cognitive Linguistics*, which is a rigorously peer-reviewed journal. We have the new *International Cognitive Linguistics Journal* edited by Professor Li Fuyin. But I want to make you aware of another possibility that is brand-new. There is a worldwide research network called the Social Science Research Network and it is an online service by academics, for academics. It is always free to upload your research—whenever you are ready—to the Social Science Research Network. It is always free to search the research other people have uploaded. It is always free to download the research that they have uploaded. The Social Science Research Network has been a great success over the last 12 years in law, and business, and economics, and political science. And the Social Science Research Network has just created a Cognitive Science Network, one of whose branches is Cognitive Linguistics. Here is a slide presenting its website. If you just point your web browser to ssrn.com, you will see this service. You can click on cognitive science, as you see here in the slide.

There are already 250,000 papers posted on the Social Science Research Network.

They are free. There have been 28 million downloads of these papers. Here, in this slide, you see what the Cognitive Science Network looks like. The Cognitive Science Network is not for profit. I am its director and I take no salary. This is a purely research activity. The Cognitive Science Network never takes copyright. You can always upload papers, and search and download papers free. You retain the copyright. You retain the rights. If you want to take your paper down, that's just fine. Leave an abstract.

Across the entire world, we don't know of a single print journal that regards posting something on the Cognitive Science Network as publication, so you are free to publish your papers in journals later. You can remove your papers at any time or leave an abstract and a link. And this is important — papers can be posted in any language, certainly including Mandarin. All you have to provide is an English translation of the title and an English translation of the abstract, a little paragraph, that's all. So when you go to my webpage, for instance, you see various articles and next to the articles you see “CSN version,” and if you click on the CSN version, your web browser takes you to my author page on the Cognitive Science Network. Remember, this is all free. You could set up your SSRN CSN author page right now, in this room, if you had online access. You could do in 10 seconds. Everyone could do this now in 10 seconds. Here, in this slide, are some more of my papers on the Cognitive Science Network, all free. When you click on one, you see the abstract. You can also download the paper, if the author has posted it.

If you have old papers, I invite you to put them up. If you have new papers to which you hold the electronic rights, put them up. If you have published papers to which you do not hold the electronic rights, ask the publisher for the electronic rights back and then put them up. There have been 28 million downloads all around the world.

The upshot is that Cognitive Linguistics now has a supplementary place for the publication (or rather the provision) of all your research. CSN does not compete with journals in the field. That is a different kind of service. You can put something up on the Cognitive Science Network and then later submit it to a journal and if the journal accepts it and prints it, you can leave it on the Cognitive Science Network or take it down, no problem. You have all the rights.

The Cognitive Science Network comes with a wonderful set of search mechanisms. Suppose you are interested in, for example, clitics in French analytic causatives, or some other linguistic topic. Suppose you are interested in iconicity in Chinese ideograms. You just type into the search engine—which you see here in this slide—what you are interested in, and up come all the papers in that field.

It is really wonderful for me to direct the publication of linguistics articles on CSN, because when I started my studies, the name “Cognitive Linguistics” did not exist. When I started in cognitive science, there was no field called “cognitive science.” Professors of cognitive science like myself all have degrees in something other than cognitive science, because those degrees did not yet exist. There were no journals.

We used to print things up and mimeograph them and then hand them out. I remember I had two copies of all of Leonard Talmy's papers. They were in blue mimeograph form and I had one box on the east coast of the United States, where I

lived, and I had another box at my mother's house in San Diego, where I was born, and I would tell people on both coasts, "Oh you really need to know about fictive motion or you really need to know about rubber-sheet typology and language." And they would say "What?" And then I would say "Leonard Talmy," and they would say "What?" because of course these papers were published much in the way Charles Fillmore's papers were published, often in places that were very, very hard to get to. That's why I had them in my boxes. And I would respond, "Just a minute", and I would copy the paper and then hand the copy over. I was a distributor of articles of Cognitive Linguistics, but no more. Now we have journals. We have an international association. We have the Cognitive Science Network, so there is no need for me any longer to keep a box of papers at my mother's. I can just say "Put them on the Cognitive Science Network," and this is really wonderful.

An awful lot of work has been done by so many people to turn cognitive linguistics into a worldwide discipline, and it's all available to you. At CSN, we send out every week a little announcement of new work in Cognitive Linguistics. It arrives in your email. This slide shows the most recent announcement, for May 15th, 2009. You see here, in that announcement, description of some new research by Alan Cienki, by Karen Sullivan, by Anna Pleshakova, by Nathaniel Smith, and by Shweta Narayan.

You can click on these links you see here in this slide to view the abstract and, if you are interested, download the paper. If you go to Google and type in "Cognitive Science Network", you will see the Cognitive Science Network as the first listing, and the organizational page with the keywords as the second listing. There it is. You just click, so it's not hard. You do not have to remember anything. Actually you do have to remember two words. They are proper nouns. Everybody say "Mark Turner." One more time, "Mark Turner." Good! If you go to markturner.org, you see this page, and you click right here, where it says "Cognitive Science Network." That is all you need to know: *Mark Turner*. Then you can get to the Cognitive Science Network. My website will take you right there.

When we look at the presentations that will happen in the general conference this afternoon, we are looking at tomorrow's research today. You are tomorrow's research today, and that is what the Cognitive Science Network is designed to provide. It is up to you. Please provide.

The title of today's lecture is "Human Meaning." That is a normal topic for Cognitive Linguistics, because language is not independent of meaning, and moreover the meaning that language depends upon is human meaning. It is the kind of meaning that is available for a human being. 50,000 years ago, more or less, during the Upper Paleolithic age, archeological evidence began to amass of a radical transformation in the Homo sapiens line. Of course we are on a cline with animals. Of course we share many, many things with mammals and with primates. Of course we do, but we are also strikingly different in our abilities. We have art, music, religion, fashions, symbols indicating social status, sign systems, mathematical insight, scientific discovery, advanced tool use, advanced social cognition.

You are looking at me and I am looking at you. All we are seeing is some photons

striking the retina, and all we are hearing is some longitudinal waves striking the ear. But it seems as if we can see what each other is thinking. I can try to figure out whether you are bored, or in agreement, or wishing I would switch topic, or not understanding my English, just by looking at your eyes and your gestures and how you hold yourself.

This ability is radically distinctive in the animal world. Nonhuman animals have severe limitations on this kind of thinking, even when working at the full extent of their abilities in central domains like eating and dealing with dominants. Chimpanzees, with a lot of human help, seem to be able to get a little bit of this social cognition, just a little bit, to be able to understand that another chimpanzee is a goal-directed agent. But they seem to have no ability to understand that another chimpanzee has false beliefs. While it is important to recognize that we are on a cline with animals, the biggest scientific question is to explain why we are so different. We celebrate diversity all around the world—which is good—but the hard scientific question is why there should be much diversity at all.

Let me explain that. The kinds of abilities we are talking about come up only in the blink of an eye evolutionally, just the last fifty thousand years—maybe sixty, maybe seventy, maybe eighty. We are all captives of the archaeological record, and if archaeologists make a new discovery of a cave in Austria tomorrow, perhaps we will need to adjust our understanding of these developments by moving them back in time a little. But that's not the point. The point is that there has been evolutionarily a very recent, major change in human behavior. This slide in my slideshow presents only the last six million years since our most recent common ancestor with the chimpanzee line, pan troglodytes and pan paniscus (that's the regular chimpanzee and the bonobo). Here, in this slide, is a reconstruction, two reconstructions, of how our line might have developed. Here, in this slide, are homo erectus, homo ergaster, and so on. The point is—it's only at the top of this line, in this little, teeny, tiny slice, just a thin thread right at the top, where you see these behaviors that we all take for granted. This is a great riddle: why should all of these behaviors be here at all? For about two billion years, they were not here. Life got on just fine. They are very new, evolutionarily. Where did they come from?

If you look at tool use, one of the things you see in the hominid line is that there were certain tool sets used in this line. By a toolset, I just mean there is a rock, and it gets shattered so as to produce a cutting-edge and maybe a grinding part or something like that. There are toolsets that last in the evolution of our ancestors for hundreds of thousands of years, species-wide, without any change, no innovation, no creativity. Hundreds of thousands of years, not four thousand years, but hundreds of thousands of years with no change, no invention, no creativity. That toolset was as far as the species could get and hang on to. If you give this toolset to a three-year-old now, the three-year-old will do something new within 15 minutes. That is what we can do: we are immensely creative, immensely innovative, immensely able to create emergent structure to put together things that did not exist there before, and it is not cognitively costly for us to do this integration. It is what we do all the time, effortlessly. This is what we are built to do.

Here are some other examples of this history. I want to emphasize that the

phenomena we are studying are very bizarre. They didn't have to be here. They are very recent. They do not look recent to us because we don't live for 50,000 years, right?

We are familiar with the kind of diversity that comes from evolutionary creation of structure: when biological evolution produces different species, we are used to seeing different behaviors in the different species. But those kinds of differences in behavior take a very, very long time, an evolutionary time, to develop. For example, one species of great ape may eat one way and another species of great ape may eat another way and that's evolutionary, not cultural.

But cultural innovation is quite different. It operates at lightening speed. Through cultural innovation, new things can be invented in 15 seconds, 15 minutes, 15 years, 150 years, 15,000 years, which is nothing.

To give you a contrast, think of the difference between *Homo sapiens sapiens* and *Lemur catta catta*. *Lemur catta catta* is the ringtail lemur. *Homo sapiens sapiens* is all of us. We are the only subspecies of our species and the *Lemur catta catta* is the only subspecies of the species *Lemur catta*, so we look just like them, right? The answer is no. If you look at the ringtail lemurs in one valley, and you walk over the hill down to the other valley, they are just the same. Now sometimes you see some differences in animals that belong to the same species but in different locations because the differences depend upon the environment. So if there is a lake, now you can wash your food, but if there is no lake on the other side of the hill, you can't wash your food. This is a difference, but that is a difference that's pushed by the environment. It is environmentally elicited, and of course there are variations. The female ringtail lemurs do not act exactly the way that the male ringtail lemurs operate, of course. But the range of behaviors and structures available to ringtail lemurs is species-wide.

But for human beings, if you are in one valley, and you walk over the hill down to another valley, the human beings in the second valley might speak a different language, wear different clothes, eat different food, eat differently. It might be that the things that were very diplomatic in one valley are very offensive in the other valley. We get these great variations that are not brought about by evolutionary biological provision of capacities to the entire species. I, for instance, am a speaker of English. I am not a speaker of Mandarin. It is not true that I, at my age, can form every sentence of Mandarin including those that I have not heard, but all the Chinese can. We have a capacity to develop new and different structure rapidly, and the question is—what kinds of mechanisms give us that capacity?

The story for the lemurs is pretty much true of the chimpanzees as well. Richard Wrangham has worked on whether or not chimpanzees have culture and it is very impressive that some chimpanzees have learned to use a shoot to get the termites and eat them, others to wash food in water. But I ask you to remember that these inventive behaviors are attached to the most basic elements of being a chimpanzee, to eating, to reproduction, to feeding, fighting, fleeing, and reproduction. And these inventive cultural behaviors are very fragile. The offspring can pick them up, but they are forgotten very easily, and there are only a few of them, whereas for human beings this is just what we do—creatively and rapidly invent new meaning, new behaviors, including meaning and behavior not restricted to local, basic operation like feeding,

fleeing, fighting, and reproduction.

If an anthropologist from Mars were to come here, we would look to that anthropologist like just another great ape, and the reason we would look like another great ape is that *we are*. In fact, we share the common mammalian brain, which is quite complicated. The human brain is not a general mush. The common mammalian brain has many pathways and a lot of anatomical structure, and we have all of that because we are mammals. In fact, we also share the common primate brain, which is very complicated. So how can we be so different from other primates? What happened?

Here, in this slide, is another presentation, just to show the location of human beings in the great ape lineage. We structurally and anatomically look like the rest of them pretty much, and in fact anatomical human beings have been here for about 150,000 years. By “anatomically modern human beings,” I mean people who have the anatomical structure that we all have. But cognitively modern human beings—that is, human beings who have all of our cognitive abilities—seem to have been here for only about 50,000 years, maybe more. I do not have a time machine that I could use to settle the dating, but if I received a time machine for my birthday, I know what I would do with it. If there were only one place I could go in my time machine, I would go back in time 50,000 years and look at human beings to see what they could do.

So what happened? Why is it that, starting fifty, sixty, seventy thousand years ago, you suddenly have art, you suddenly have all kinds of inventiveness and creativity? Well, the answer is, I think, we have these behaviors because of the kinds of brains we have.

There have been many stories trying to explain where human beings could have come from, many stories of our origins. One of those stories comes from the Greeks—the story of Prometheus and fire. There are similar stories—of a god, or ancient astronauts, or anything more advanced than we are—in which the more advanced being brings us what we do not have. In these stories, human beings learned from more advanced beings. These stories are entertaining blends: we all know that we learn from our parents, who are more advanced agents, so we can blend that scene of learning from parents with the situation of our ancestors, and in the blend, there are analogous of our parents—that is, there are more advanced agents teach our ancestors learned.

There is also a legend in which a Chinese cave dragon blew on monkeys and turned them into human beings. There’s of course the anthropological (oh, thank you, now I have two cups of tea, seriously, these assistants are just amazing. I have to move to China, thank you, *xiexie*)—the anthropological story, in which we all come from Africa. In that story, a group of human beings grew restless and decided to leave, in two waves, called “Out of Africa One” and “Out of Africa Two.”

There is another story of our origin—I will call this story the “Forbidden Fruit Blending story.” This story, which is the one I tell here in these lectures, is the story of the evolutionary development of a higher-order capacity out of a mental operation that is shared all the way back in the mammalian line in rudimentary fashions. It is a story of a kind we find often in science—a story in which a small change in causes produces a big change in the effects. That is the story I am going to start talking about today and introduce you to more fully this afternoon. You can find an entry-level

presentation of that story in a book I wrote, *The Literary Mind: The Origins of Thought and Language*, and you can find applications of that story in the book *Cognitive Dimensions of Social Science*.

The most comprehensive presentation of the blending story is contained in a book Gilles Fauconnier and I wrote, *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. Gilles of course was a speaker in this series, before Chris Sinha. All of your questions by the way about this capacity called “blending” are answered at <http://blending.sanford.edu>. So you just go there and ask questions. It is all there, hundreds of papers. This slide shows Gilles. I only wish he were with us here, but he has been here before.

There are major questions that we are going to try to take on in the blending story, not because we are ambitious, but because we must take them on. I think that our cognitively modern human capacities came up together and that trying to explain them separately leads us to error. They came up together phylogenetically, it seems. They also come up together ontogenetically, in the child. The blending story takes these capacities to be a group: language, art, music mathematical insight, scientific discovery, religion, advanced social cognition, refined tool use, advanced music and dance,

Have you ever noticed that your dogs don't dance? A dog will hop up and down with a human being, if you urge or train it to, and you can condition animals to almost anything. But have you ever noticed that when music is playing, the dog does not beat the time with its paw? There is nothing biologically stopping it. It has a paw and it can move that paw; it could beat; it could beat three-four time; it could do all kinds of things, but it doesn't. It is hard to realize that it doesn't because we've seen so many cartoons and so many movies in which dogs behave in ways that dogs don't behave. We have seen so many representations of fictional animals and it's so easy to project what we know in our minds onto our understanding of other animals. It is very easy to overestimate what animals can do.

By the way, let me say, I am very fond of animals and we could talk for a very long time about how brilliant they are . . . corvids, all those blue jays, scrub jays, New Zealand rooks, and their episodic memory, and their ability to use a tool to get another tool . . . we can talk about itborder collies who round up the sheep . . . all these kinds of animal behaviors . . .

I am not overlooking the abilities of animals. On the contrary. But human beings have very many crucial and impressive behaviors that nonhuman animals do not have. What makes these behaviors possible? This is the big question. How do they work? How can new meanings arise?

Evolution might very well give you some specific ways to construe some specific situation. If you see your prey, you might think, “Wow, there is food,” or perhaps not think consciously at all, but chase it down. You get it; you respond. Or suppose you see a big scary animal that makes a noise and comes at you, and you run away. Yes, evolution can certainly provide these species-wide frames, tied to the local scale of the organism and its species behavior.

But human beings come up with new meanings all the time. They come up with new things all the time, new constructions all the time. How can they do that? How can

new meaning arise? And how does the mind go beyond the local scale?

I want to pause for just a second to emphasize this capacity to go beyond the local scale. I will come back to this capacity again in this series of talks. It is one of the major open questions in science. Other animals have impressive abilities to do things such as remember where they store the nuts, or figure out which of three or four other agents to be allied with while hunting. But all of these behaviors are tied to the local moment, to the local scale: a visual field that we can see, a small amount of time, a few agents. When you see an animal acting at local scale, it can be very impressive.

But a human being thinks without effort and constantly beyond the local scale. Human beings remember when they were five years old. They construct a sense of personal identity that runs through their whole life. They think about their great-grandchildren who are not even here yet. They establish trust funds for their great-grandchildren who are not even here yet. They think about global-warming. They think about international relations in the areas of the world they have never visited. They work at a scale of intentionality and time and space that goes very far beyond the local scale for the species. Yes, it is true, dolphins have signature whistles and seem to have some kind of genetic ability to recognize each other. Yes, it is true that the chimpanzee in the Swedish zoo, the genius chimpanzee, stores up rocks to throw at the visitors tomorrow. This has been cited as evidence that perhaps animals are not tied so completely to the local moment. (Endel Tulving is the one who talked about mental time travel, and how good human beings are at it compared to other species.) But notice that even that remarkable genius chimpanzee who stores those rocks is doing something that goes only the tiniest bit beyond the local scale and is for the most part tightly tied to the local scale. His action is to pick up a rock and throw it: it's a rock, it's his arm, it's a bodily action. The visitors are there every day, so the situation is routine, one he is conditioned to. He can see the visitors, right in front of him, on the other side of the fence. The chimpanzee has been in a human environment that imposes great regularity. Getting rocks in position to throw not now but tomorrow, where tomorrow is just like today, is as far as the chimpanzee gets. He does not, for example, ever practice throwing rocks at a tree so that he can throw them with greater accuracy at the human beings. By contrast, practicing throwing rocks at a tree or another target is something human children do all the time. They practice many things.

How could human beings have developed this ability to go beyond human scale? Where does that come from? Human beings transcend local scale. I am not talking here about how evolution has built us—or any animal—to behave at local scale in ways that are consequential for events beyond the local scale. So for example, we don't need to know about nutrition to feel hunger, because evolution has built in us something at local scale—namely hunger—and that something at local human scale cause behaviors that have long range consequences. But we don't need to be aware of that causation. We do not have to be thinking about our great-great-grandchildren in order to feel attracted to a member of the opposite sex, because evolution has built desire into us at local scale, and desire will have long-range consequences. That's not the kind of long-range consequentiality I am talking about. All animals have that, and of course

evolution can do it. It does it for fish. It does it for seagulls. It does it for worms. That kind of long-range consequentiality is a different question.

The question for us is, how can we *think* about something that's far beyond our scale? This is a bigger problem than we are likely to imagine. Charles Scott Sherrington, the great English neurophysiologist, pointed out that, for thought, biologically, all you have is the central nervous system in your body. To be sure, we have each other and we have artifacts and environments, but, biologically, all we have is the central nerve system in our bodies.

Thought is run by your biology, but your biology is always just now. It is always present. Your brain, Sherrington wrote, is "an enchanted loom where millions of flashing shuttles weave a dissolving pattern, always a meaningful pattern though never an abiding one, a shifting harmony of sub-patterns."

And with this *present* biological capability, you are able to have *long-range* thoughts, but how? How can you think of the past? So that is one of the major questions I will take up in this series of talks.

Another major question is, where does new meaning come from? There have been some previous attempts to answer this question. Evolutionary psychology takes the standard evolutionary view that evolutionary decent can build in you, build in our species, a few frames of meaning, and that's true. This evolutionary psychological view gives us a good, short ride, but it cannot answer the question of where new meanings come from, because mostly we come up with new meanings on a time scale that is incredibly faster than evolution can possibly work. There's no evidence that basic mental operations of human cognition have changed in the last 50,000 years or so. But look all the cultures that have been developed in that time, right?

Another attempt to explain how new meaning can arise comes from learning theory: you can extract schemata from the experience that you have and then apply those schemata to other situations. That's a pretty good story, but it has a crucial problem. The problem is that a great deal of what exists in the world came up because we invented it and put it in the world, so it is not the case that we could have extracted schemata for those things from our experience and then used them. The invention preceded the existence in the world. For example, human beings did not originally get the idea of a restaurant by extracting a schema from all the restaurants they encountered. More seriously, consider laws and institutions. We invented them; we created them. The story that you extract schemata from your experience does not give us an explanation for how all these things got here before anybody experienced them. How did we invent them mentally before we had experience of them?

Another attempt to explain how new meaning arises comes from complexity theory of the sort that is studied at the Santa Fe Institute. In this view, human beings are complex adaptive self-organizing systems. I think there's a great deal of truth to that view, and I'll be talking a little bit about how we invent complexities through conceptual blending.

Going back to the reductionist story, one idea is that human beings have some independent traits. I illustrate this reductionist view in this slide: You get religion in

a box, and science in a box, and art in a box, and refined tool use in a box, and fashion, and language, and music and so on, with the idea that there is a different box for each of these traits, and somehow evolution gave them to you, one at a time, independent of each other.

I think, on the contrary, that what happened is that human beings evolved a higher level of a basic mental operation. The basic mental operation is conceptual integration, “blending”—which we will talk about throughout these lectures. The higher level of blending that evolved for cognitively modern human beings is “double-scope” blending, which I will define and discuss this afternoon. Once human beings had the advanced form of this basic mental operation, culture became possible. Once they had double-scope blending, other higher-order mental capacities became possible. Those capacities had to be developed. Languages had to be developed. Art had to be developed. All these behaviors had to be developed. Tools had to be developed.

I tried to make a slide to illustrate the difference between the view of human behavior as a collection of different boxes and the view I will be proposing. The first sketch I made, shown in this slide, is misleading, because it could be misread as suggesting that the behaviors came up in a linear sequence. So I made a different set of slides, showing, here, the clan of abilities that arise out of conceptual integration, “blending.” In this slide, you see the capacity for conceptual integration increasing along a gradient over evolutionary time. At the top of that gradient comes “double-scope blending.” Cognitively modern human beings are the only species that can do double-scope blending. As you see in the slide, once double-scope blending is possible, then a host of related behaviors develop out of it. All of them become possible simultaneously.

That is the overarching story I offer as an answer to these major questions, namely: How can we think beyond local scale? Where do new meanings come from? Why are we so distinctive? What is language?

This group of abilities that human beings possess—language, advanced social cognition, mathematical invention, and so on—should not be thought of as independent. They should be thought of as a group, as a pack, as a coordinated and cooperating set of abilities that all are made possible by the same basic mental operation—double-scope blending.

There are Cub Scouts all over the world. They are an organization of children. They follow the “Cub Scout Law of the Pack.” The Cub Scout Law of the Pack has three parts. The first part is, “*the Cub Scout follows Akela*”. You don’t need to know who Akela is. In the blend I am making, I want you to think of Akela as the mental ability for double-scope blending. The second part is, “*the Cub Scout helps the pack go*.” The third part is, “*the pack helps the Cub Scout grow*.” Think of our higher-order human abilities as cub scouts: they belong to a pack. They follow double-scope blending. And any one of them helps the entire pack, and the entire pack helps any one of them. Any one of them helps the others develop and work; and they all help the individual ability develop and work. This pack of higher-order human abilities is a long alliance.

I want to make it clear just this once in this series of talks that the “double-scope

blending” story of human abilities is not a triumphal story. Mostly in discussing human abilities, we sound positive and proud. Look at human beings: they invented complex numbers; they launched rockets into space; they circumnavigated the globe; they can paint.

But it’s also the case that we are a species of great sorrow and great suffering and great psychopathology. There is no evidence that members of other species spend twenty years being paralyzed with guilt for something they didn’t do in one moment when they were a teenager or that they construct fantasies in which terrible things happen and then feel bad. Human beings think of counterfactual scenarios. They think of the hypothetical. They think of what isn’t and what might have been and things that might not be, and they think of these things reaching far beyond local scale, and this often brings great sorrow and great suffering.

Rilke noticed this. He wrote in the *Duino Elegies*, “The shrewd animals notice that we’re not very much at home in the world we’ve expounded.” I don’t know how shrewd the animals are, or how much they notice, but human beings often have the feeling of not being at home in our world. Cognitive Science and Cognitive Linguistics—let’s just make it very clear at the outset—is not a story of triumph. It is a story of the human condition and it comes with many wonderful things and many things that are hard.

Now let me tell you what conceptual blending is, and then I will take questions.

Imagine that there is a man who is at a wedding. He is a groomsman. He is helping at the wedding. And he is in San Diego, where I was raised. The wedding is taking place on the cliff overlooking the Pacific Ocean, where I like to swim. Let’s say that our understanding of this little scene is one “mental space.”

In the slide, this yellow circle is supposed to represent this mental space, with the groomsman and the wedding. Of course there is no such circle in our brain. It’s just a representation. So don’t ask where the yellow circle resides in the head. It is a symbol for a theoretical construct. A mental space is a fairly compact mental array, fairly simple, that you can manipulate mentally in a moment of thinking and acting.

While the groomsman is at this wedding, he is thinking about the fact that three weeks ago, he was diving with his girlfriend off Cabo San Lucas, looking for treasure. In the slide, this blue circle represents the mental space of his diving off Cabo San Lucas with his girlfriend. Notice that the girlfriend is in the blue circle but not the yellow circle: she is in Cabo San Lucas diving with him, but she is not at the wedding with him. And he is thinking of both the wedding and the diving.

The fact that the groomsman can think during the same brief interval of time about both the wedding and the diving immediately presents us with a major scientific problem—why should human beings be able to think simultaneously of two stories that do not correspond to each other, that are in conflict with each other? The groomsman is thinking of the wedding here and now and also of the diving and the girlfriend and Cabo San Lucas. Shouldn’t evolution shut down simultaneous thinking about two clashing things? Doesn’t thinking about two conflicting things present the possibility of confusion, of loss of focus?

But notice that he does not get confused. He does not swim down the aisle even

though in the blue mental space—the diving space—he is swimming. He talks normally at the wedding even though in the diving space he’s got something in his mouth. He does not mistake the bride for a shark. If you are able to call these very different stories to mind simultaneously, what stops you from getting confused?

This is such a big problem in cognitive science that Arthur Glenberg in 1997 wrote a brilliant paper called “What memory is for” in which he said that mostly memory in other species seem to be subordinated to what they are doing in the moment. So if you want to get out of the forest, you can remember how you got in, but you keep that memory tied to and subordinated to the present moment. You remember what you need to operate in the here-and-now. But human beings constantly call up multiple stories that are not needed for operating in the present moment. Terry Deacon has been working on this problem, along some of the lines presented in blending theory, for about eight years.

But the groomsman can take another mental step, as we see in this slide. He can connect the two mental spaces—the wedding and the diving—by vital relations. I represent those relations by lines between the yellow circle and the blue circle. The groomsman can use a frame relationship to connect the bride in one space and the girlfriend in the other, and to connect the groom in one space and himself in the other, for example. And then he can engage in selective projection, as we see in this slide: that is, he can take parts of each mental space—the wedding and the diving—and project them to a new imaginative mental space, a blended mental space, which I illustrate in this slide with a green circle. And in this new blended mental space, he can be marrying his girlfriend right here. The blend has the wedding, and its roles for bride and groom, but it has the groomsman as the groom and his girlfriend as the bride.

The groomsman is not deluded. He knows that this is a fantasy, but he can think about it. He can think, “You know, this is a pretty good idea. I have never thought about it before. I’ll make a plan to achieve that.” That’s creativity, that’s emergent structure. He is coming up with something he has never seen before, namely, marrying his girlfriend. And he can like it and try to act upon it.

Alternatively, as he is running the simulation in which he is marrying his girlfriend, he can come to the moment when the officiant asks her, “Do you take this man to be your lawfully wedded husband?” and realize that at that moment she might say, “I would never marry you”, and he can feel that it’s true and that in some ways he hasn’t been thinking properly about his relationship with his girlfriend.

In this blend (the green circle), there is selective projection: only some of what is in the input mental spaces (the yellow and blue circles) comes into the blend. And in this blend, there is emergent structure, namely, here he is marrying his girlfriend, right here on the scene. Notice that in the yellow circle and the blue circle, he is *not* marrying his girlfriend. The marriage to the girlfriend is not simply copied from the input mental spaces, because it is not in those spaces. It emerges in the blend, the green circle.

This blending, this invention of new meaning, this creation of new structure, is what we do all time. This is what we are capable of all the time. It seems remarkable to us in this example, because I am explicitly and painstakingly presenting to you a visible

example, but in fact this kind of blending happens all the time. 99.99% of blends are never noticed by the people making them.

Blending is the subject of *The Way We Think*. I want to give you a riddle that we present in that book, and then I will stop. The riddle, or rather its solution, is an example of a conceptual blend. It is the riddle of the Buddhist monk, which is an old riddle. Carl Dunker used this riddle in psychology, and it is quoted by Arthur Koestler. It is a riddle for you, a brain teaser:

A Buddhist monk begins at dawn one day walking up a mountain, reaches the top at sunset, meditates at the top overnight until, at dawn, he begins to walk back to the foot of the mountain, which he reaches at sunset. Make no assumptions about his starting or stopping or about his pace during the trips. Riddle: is there a place on the path, which the monk occupies at the same hour of the day on the two separate journeys?

Let me show you a little movie of the Buddhist monk. Here is a Buddhist monk, and he walks. He starts at dawn, and he walks up to the top. He reaches the top at sunset, sits down, meditates overnight until dawn comes again. He gets up and he walks back down the mountain path, reaching the bottom at sunset. So let's see this again: starts at dawn, goes up to the mountain path, however he goes, reaches the top at sunset, sits down, meditates overnight until dawn comes around and he gets up and goes back down the path and reaches the bottom at the sunset.

Now the question is, you may not know what point it is, but is there a point on the path that he occupies at the same hour of the day on the two separate days? I won't answer that riddle; you will answer that riddle. I will put you in the position to see the answer of that riddle at 3:30 pm today, and we will explore how you did it as an example of conceptual blending, selective projection, emergent meaning and basic mental operation that makes language possible.

Thank you. XIEXIE! We will now take some questions, if you like.

Lecture Two Conceptual Integration

You'll be facing me for many, many days, each day in the bitterness of late afternoon. So we need to find some way to wake up and pursue research. I'll try to begin each session with some music. The tune you just heard was from Brazil.

I am very happy, for two reasons. Reason number 1: somebody handed me this wonderful packet of abstracts of presentations that are going to be delivered in the afternoons after my talks. They have wonderful titles, like "Iconicity in Chinese Sign Language", "Semantic Field Again: Synonymous Homogenous Construction and English Vocabulary Acquisition", "Cognitive Development, Metaphorical Mapping, and Expansion of Lexico-semantic Categories", "The Role of Frame-shifting in Humour from Social Perspectives", "Embodied Realism in Cognitive Linguistics". This is terrific. It's tomorrow's research today. So of course I'd like you to go post all of these on the Cognitive Science Network, just sign up in SSRN.com, become a user, post your research, and there it is. The whole world will be able to see it. In fact, while we were having lunch, the next installment of this little email message that we send out every week to announce new cognitive linguistics research just appeared. Here it is on the slide. You see that for this week, we are announcing "How Compression Give Rise to Metaphor and Metonymy" by Gilles Fauconnier, "When Do Understanders Mentally Simulate Locations?" by Nian Liu and Benjamin Bergen, and other sorts of research. We announce some new research every week.

An example of how you might use the Cognitive Science Network can be given right in these talks. We heard Professor Geeraerts talk today about the role of polysemy in cognitive sociolinguistics. Gilles Fauconnier and I have an article on polysemy and conceptual blending: it analyzes the way in which the basic mental operation of conceptual integration makes polysemy possible. It presents four different ways in which varieties and kinds of polysemy can come up in conceptual integration networks. That paper is available on the CSN site. So you could go to the site, search for polysemy, and there would be the paper for you to download, along with many other articles. No waiting, no fuss, no muss.

The second reason that I am very, very happy is that I went to "Sculpting in Time" café next to Deep Tennis after lunch, and I asked for a latte, a *caffè latte*. And here in this slide you see a picture of what they served to me. This is a beautiful latte. They call this "latte art," which is a culinary blend and a mental blend. In "latte art," you pour the espresso into the steamed milk, but you drizzle the espresso in such a way as to paint on the surface of the milk. I just drank this. But before I drank it, I took a picture of it to show you, because this is the kind of inventiveness that comes from the creative mental operations human beings possess.

To take the surface of the milk, and turn it into a canvas for art, is very inventive.

Cognitive Linguistics, if you recall, is a branch of cognitive science in which language is principally seen as a branch of human cognition. And as these lectures go on, we will pick up the pace and move faster and bring more and more technical features and aspects of language into analysis. But it's important here at the beginning to

remember that these analyses are for the most part not specific to language. These mental operations are, for the most part, not restricted to language. They are basic mental operations that make new meaning and higher-order cognitive performances possible.

And so we need to focus on the way in which these mental operations work generally so that we can see how they work in language. There are many ways in which conceptual integration networks make language possible, and make the production of language possible, and make the change of language possible. We'll be going through some of those.

Let us pick up where we left off talking about conceptual integration. Remember you saw that you can find some of this material in the book, *The Way We Think*.

How many of you were here this morning? I think most of you were. I recognize faces. You recall we talked about the man who is at a wedding party, engaged in the wedding party. But while he is at the wedding party, he is remembering the fact that three weeks ago he was diving with his girlfriend at Cabo San Lucas. And this is a major scientific question in cognitive science—why evolution has not stopped us from developing the ability to think of two things simultaneously when they're incompatible and where one of them is not serving present action and in fact goes against present action. Isn't there a danger that we would get confused? Isn't there a danger that we would make mistakes in action? Why do we have this ability? I mentioned Glenberg's analysis of this problem. I mentioned Terry Deacon's project to try to work on how evolution might have lifted, or eased, the inhibition against thinking about what is not relevant. What might have happened evolutionarily to lift the disinhibition of rival concepts, rival conceptual networks, conflicting ideas that do not serve the present moment.

The brain is all about activation, as we saw in that quotation from Sherrington about the brain as an enchanted loom. We are often shown a picture of the brain. Here is one, and I'll show you others tomorrow or the next day. It is easy in looking at the brain to think "Ah, there is a brain, so that is a thought".

But no. That brain in the picture is dead. That brain is not good for thought at all. What matters is not just the object, but what's active in the brain, and it is important always to focus on the fact that activation varies. We tend to think of our concepts and our abilities as things that are always there and that are always ready. But not so. What your brain can do depends upon what is active. And it's a real question: Why you should have a kind of brain that can activate things multiply and simultaneously? I will talk about that as we go along.

A hypothesis presented in original blending theory (Fauconnier and Turner 2002) is that perhaps human memory and conceptual integration evolved simultaneously so that memory would supply better material to conceptual integration, so that we would be more inventive. Of course, that puts it teleologically, so we would need to cash this out if we wanted to be accurate: the change in human thought so we could do this simultaneous activation managed to survive and build, through various evolutionary means, in the line of human descent.

Anyway, there is the man at the wedding party. He's thinking about diving with his girlfriend at Cabo San Lucas, but she is not here at the wedding party. He makes a connection between the bride and his girlfriend, and then he blends them. He projects selectively to the blend. Notice the girlfriend is not here at the wedding, but she is projected into the blend. There was no bride in Cabo San Lucas, but the role for *bride* is projected into the blend.

And he is now mentally running a simulation. It is a blended simulation. It has emergent structure, selective projection, it has running the blend, and it has emergent structure. Now, in the blend, he is marrying his girlfriend and contemplating the prospect. Notice that this event is not in either of the input spaces. In neither of the input spaces is he marrying his girlfriend. But now, in the blend, there is emergent structure in the blend. In the blend, he is contemplating marrying his girlfriend in fact right here.

Conceptual integration is an operation. It is not a diagram. There are no circles and arrows in the brain. And certainly conceptual integration does not look like this fetish four-space diagram that I'm putting up here. This is just something I can point to, for the purposes of teaching and discussion.

Conceptual integration always has a network of mental spaces that need to be connected in order to be creative. There are always input spaces. Connections of vital relation are always constructed between them. Blending cannot get started without some kind of connection between the inputs, some kind of conceptual link. We call the great majority of these connections "vital relations" for reasons we will discuss. More connections can develop as we construct the conceptual blending network. Indeed that is one of the virtues of conceptual integration. In blending, we see projection to the blend—typically selective projection, not just wholesale copying or cut-and-paste. In blending, we typically see emergence of structure in the blend that is not contained in either of the input spaces, such the wedding to the girlfriend.

The purpose of a conceptual integration network is not to create a blend that eliminates the network. On the contrary, often the purpose of the blend is to make it possible for us to command the network, to grasp and hold onto a network that otherwise would be too large and strange and alien for us to understand with our human brains.

The blend gives us a human-scale *platform*, a place to stand, something we can *grasp* that is *congenial* for human cognition and from which we can organize a conceptual network that might otherwise be beyond our mental abilities, as it seems to be quite beyond the mental abilities of every other species.

Now, I gave you the riddle of the Buddhist monk and I asked you to solve the riddle.

A Buddhist monk begins at dawn one day walking up a mountain, reaches the top at sunset, meditates at the top overnight until, at dawn, he begins to walk back to the foot of the mountain, which he reaches at sunset. Make no assumptions about his starting or stopping or about his pace during the trips. Riddle: is there a place on the path that the monk occupies at the same hour of the day on the two separate journeys?

So, here, in this movie on the slide, we have a monk. I will narrate his actions as they happen: at dawn, he rises, he walks up to the top, which he reaches at sunset, sits there all evening, meditates overnight, rises at dawn, and walks back down, and reaches the bottom at sunset. What's the answer to the riddle? Who has an answer? Who has an answer? Is it true? Is there a place that the monk occupies on the same hour of the day on the two successive days?

You should not feel embarrassed. I must tell you that I've given this riddle to many physicists and mathematicians. And the first thing that most of the mathematicians say is, "Ah, that's the fixed point theorem". It isn't. But this can be a hard problem. I see a hand in the back, an answer in the back. Yes?

(From the audience): "Assuming that the distance between ..."

No, stop. No assuming is allowed. He can move anyway he wants.

(From the audience): "Ok, that guy goes up the hill . . . there's a point he would meet himself when he comes down."

Exactly, now notice what the person in the audience just did. He said that he—meaning the monk— would meet himself coming back down.

Notice: *it is impossible for a person to meet himself*. This doesn't happen. Next notice that the word *meet* is ungrammatical for this scene, that is, ungrammatical for the movie we watched of the monk going up the hill and coming back down. There is no meeting in that movie. The reflective *himself* has no referent for the scene, and so the linguistic construction *the place is where he would meet himself* is ungrammatical for this scene. You cannot use this language grammatically of the scene in the movie.

But the scene in the movie is not the scene to which the person in the audience is applying the phrase "meet himself when he comes down."

What has just happened in the audience is that someone has said to himself, "Ah, I know what I'll do. I'll take the journey on the first day and the journey on the second day, and I will blend them. I will superimpose them as if you had a video camera and you videotaped the monk on day one, and then you videotaped the monk on day two. A video camera just like this one. And you have two projectors, and you project the two movies of the two days on the same screen, and it's done—now we have a scene with two monks which are the same monk, and they meet, so the monk is meeting himself." There are two monks on the screen and they have to cross somewhere. They have to cross somewhere. This scene is a blend. It happens to be one that does not come quickly to mind for most people.

Notice that you understand this blend before I have shown you a movie of it on the screen. Once the idea of blending the two journeys comes to you, you understand the blend.

There is selective projection to the blend, that is, we do not project to the blend the date from the input spaces, because the date in one and the date in the other are two incompatible dates. We do project the time of day—noon, three o'clock, . . . And very interestingly, as we'll talk about later, we project the topology of space-time collocation. You have to preserve that under the topology principle, as we'll discuss later, in order for this to work.

There is emergent structure, namely, there are now two monks, not just one, and

they're the same monk and there is a meeting. There was no meeting in the original scene, but now that you have blended them, there is emergent structure. There is a meeting, because you know that when two people are approaching each other on the path and they go to opposite terminals of the path, they have to meet. That's a common human frame you have of interaction.

You can use the blend to understand this network of the monk's movement. You didn't quite see, many of you, that there would be a place on this path here that is identical to a place on that path there and that the time that the monk is at this place on day one is the same time as the time the monk is at this place on day two. You are working in a big network to understand the journey over successive days. But once we had the blend as a compressed anchor for the network, a compressed scene that is at human scale and that you are very good at grasping, then you understood its structure and how it connects to the rest of the network immediately.

You've never seen a person meeting himself. Now why can you imagine this thing that you've never seen before? It is not because you have a genetic module for understanding monks' meeting themselves. It is because you have the mental operation of blending. And this is not even a double-scope blend. This is a mirror network—which I'll discuss in a minute. Once you see this, you recognize the solution to the riddle.

The selective projection in this case is really selective. What would happen to you if you were walking along a path and suddenly you saw yourself coming toward you? What would you do? Notice that in the blend, the monks don't even notice their meeting. If you saw yourself approaching yourself, you'd probably scream "Ahhhh!" and run in the other direction, or more blandly ask yourself, "What is going on?" But no, the awareness of the other monk as such, which would let him recognize that there are two people, and a meeting, and that the two people are the same person, and in fact he is two people meeting—that basic ability for awareness does not get projected to the blend. The monks in the blend are not aware of the blend.

Moreover, you do not derive from this blend an understanding of where people can be located at various times. You do not say "Ah, I see now, I was stupid before. Now I understand how the world works. It is possible for people to meet themselves." In general, the monk's standard ability to make inferences about the world on the basis of experience is not projected to the blend. The monk in the blend, for example, does not say to himself (although you could write a science-fiction novel this way): "Wait a minute, is this yesterday? Yesterday I was coming up the hill. And now I am coming up the hill and I thought today was today". The monk in the blend is not getting confused about time. No. What is happening in this case is that the blend is being governed by our purpose in making the conceptual integration network. The purpose of this conceptual integration is not fantasy or entertainment, but rather to do what we call "solving for the inputs". You want to understand something about the relationship between the inputs. That is the purpose of the blend. The purpose is not to create a fantasy world in which people can suddenly meet themselves. All of this is complex, but recognizing the blend and its meaning, once it is suggested to you, is pretty easy for everybody. Most people, when they think about it, are amazed at how easy such

blending is, and they then say to themselves, “How could this blend happen?”

Here, in this movie on the screen, for example, is the Buddhist monk version two. In this movie, he moves at a different pace, very slowly at the beginning and then he rises up to the hill, and sits down, meditates overnight, rises at dawn and goes back down with a different speed. But that’s no problem, because he gets back down to the bottom. And if you blend them again, now with their different speeds, although in the previous version of the journey, he met himself above this cusp, this time he meets below this cusp. And you can even make kind of extreme examples, such as in this third move, where he races up and then goes very slowly at the top, because he’s tired, and sits down and meditates overnight, gets up at dawn, goes down very slowly because he wants to see the view, the beautiful view of the mountain top, and he has to rush back down in order to get back to the temple in time for dinner. And if you blend those, now he meets way up here, you see.

But the point is that he has to meet somewhere, because once you recognize that when you put these together, you have two people approaching each other on the path, they have to meet.

What probably hasn’t occurred to you is that you projected this structure from the inputs to the blend in such a way as to keep well-ordering of time and location. That is, you had a choice, and you made the choice: when the monk is on a certain spot on the path at a certain time of day, you brought that connection into the blend and you made it so that if he’s at a spot at time one, and time two is later, then time one and time two come into the blend with the same relationship in the blend as that they had in the input.

In fact, this specific projection is not necessary for the blend to give the right answer, mathematically. The requirement is more general: as long as you preserve the coupling of time and location from the inputs, and map the two identical spots on the path in the inputs to the identical spot on the path in the blend, then the monk has to meet himself somewhere. Following this topological constraint, you could project time-space couples from the inputs to the blend so that temporal moments are not well-ordered according to chronology in the blend. Then, in effect, the monks would hop instantaneously back and forth in the blend, and go up and down the path as you see in my gesture—so long as they start at dawn and end at dusk, and as long as you preserve those time-space collocations and map identical locations to identical locations, the blend will work mathematically: the monk will still meet himself somewhere. But it’s not so easy to see that the hopping monk blend works. What makes it easy to see is the meeting point is the congenial framing of two people approaching each other on a path.

Let’s review just a little. What is going on here—and hold in mind this all-important fact—is that I can talk about the blend because I already have language that suits the blend. I have language that suits the human-scale scene. And when I talk about the blend, you understand what it means for the network. I don’t have to have language for the entire network: I can talk about the blend, and you will connect the blend up to the network. If I say “the place is where he meets himself,” that is language that is grammatical for the blend.

But notice that in the riddle I asked you, I did not ask you where the monk meets

himself. I asked you a different riddle. The riddle I asked you to solve has to do with the relationship between the two inputs. We can talk about the inputs by talking about the blend. Language is now available to communicate about the entire conceptual network by using constructions we already have that fit the blend. Hold on to that point. It's one of the many important ways in which conceptual integration makes language possible. Because of blending and blends, you do not need new language to express new things. You can express new concepts and new structure in networks with old language. You already have everything you need. I did not have to invent grammatical structure. I did not have to invent lexical items. I could say, "the place is where the monk meets himself". And that is understood by you as conveying structure about the great network.

Here are input spaces. In the two input spaces, there is a path, and there is a monk. There is the ascent, and there is the descent. The connection between them is quite obvious, because this is a mirror network.

A mirror network is one in which the input spaces share the same conceptual frame, where by "frame", I mean the organizing structure, such as temporal structure, participant structure, modal structure—necessity, probability, things like that—, image schemata, and so on. In both of these, there is a monk who is traveling on a mountain path from a departure point to an arrival point in a space of a day, meaning from dawn to dusk. That conceptual structure organizes both of these inputs, which makes it very easy to make vital relations between them. The path is identical to the path. The time frame, not the date, but the time frame is identical to the time frame. The mountain is identical to the mountain. And so we can connect them. There is a generic space, or there can be a generic space, in which what is shared between the inputs is captured.

Now, do not think that the generic space is always already there, independent of the connection, or think that the reason we can connect two input spaces is that there is already a pre-existing generic abstract space that connects them all up. On the contrary, people often remanufacture what goes into the generic space as they are working on the conceptual integration network. They try different vital relations between the inputs to see if they can get a good network. So these things are dynamic and they change as we work and you see many examples of that. There is projection to the blended space. Notice something interesting in the Buddhist Monk network. The paths in the inputs are fused into one path in the blend. We projected two paths down into the blend as one path. So there's fusion here in this blending.

But the monks are brought in separately. They are not fused. The projection is selective. There is emergent structure in the blend, namely, the monk meets himself, something that does not happen in the inputs.

A typical network will be like this, but remember, please, conceptual integration is a mental operation. And we will go through a number of its features. It is not a diagram. I can show you conceptual integration networks that have 45 mental spaces where there are blends and the blends become inputs to other blends. Conceptual integration, to create conceptual integration networks, is a basic mental operation which we use again and again and again in constructing meaning. It has certain features: vital

relations between inputs, selective projection to the blend, emergent structure, creativity in the blend. This is the basis of human creativity that makes culture and invention of culture possible. And culture supports our use of conceptual integration. We all work together to support it.

Now I want to give you another set of examples to try to make what we have talked about today survive in memory, because these are very basic ideas that will become much more sophisticated as we go on.

I would like very much to run a set of experiments, which I call “blending box experiments.” A simple blending box experiment goes like this. I will give a slideshow on the screen to illustrate what I am talking about as we go along.

Suppose we show somebody something like this on a computer screen—a blue ball and a place for it to land. And we show them that if you press this blue button, then what happens to the blue ball is that it rolls diagonally across the screen from upper left to lower right, and lands here. Got that? Watch the blue button get pressed here. You see that the button turns black in the slide show when I press the blue button here on my computer. Press the blue button; the blue ball rolls. Now, you also see that there is a red ball, in the upper right. And when you press the red button, the red ball rolls diagonally, from upper right to lower left, where it lands. Watch the red button get pressed: see, that’s what the red ball does.

Everybody get this? Press the red button, that’s what the red ball does. Let’s go back. Press the blue button, that’s what the blue button does. And if you press the red button, that’s what the red button does. Now, you tell people they can press the buttons themselves, and if the balls land on the green oval here at the bottom of the screen, in the middle, you will give them a hundred yuan. A million. You tell them whatever. You give them an incentive. You show them this little computer with buttons. If the balls land on the green area, you’ll get lots of money.

You could also present this little scene to dogs. They wouldn’t care about the money and maybe not the computer screen. But suppose you built a box, a real box, with real pedals to push, so that it would in fact be easy for a dog lying down and facing the box to push one pedal with one paw and the other pedal with the other. Suppose we had in the box not blue and red balls but hard spherical yummy dog treats that smell good to the dog. Suppose it is all behind plastic glass, except for ventilation holes, and a nice, big hole here where the green oval is, so that if the spherical treats land on the green oval, they roll right out where the dog can eat them.

You could design a similar presentation for birds. New Zealand rooks could stand on one pedal and peck a button. You could design equivalent presentations in lots of ways, for lots of animals.

I would like to know at what age young children can make the balls land on the green oval. Notice that this is just a mirror network. It is like the Buddhist monk, in the sense that you press the blue button, the blue ball rolls, it goes on the trajectory, it stops, but does not land on the green; you press the red button, the red ball rolls, it goes on the trajectory, it stops, but does not land on the green. Can you make the balls land on the green oval?

I’d like to know if dogs can do it. This is pretty challenging. Now notice I have

shown you two scenes: press blue button, blue ball rolls, press red button, red ball rolls.

How are you going to get the balls onto the green oval? What's the answer? Press the two buttons at the same time—now that would be cut and paste. Notice that they don't land on the green oval. They miss each other. We don't want to build these experiments so that if you press the two buttons simultaneously, you get lucky, because it is too standard and automatic for the animal (including us) to do the same thing with the right and left. The human body is bilaterally symmetric. It is very easy to do things at the same time on both sides of the body. For example, raise both hands; now lower them; now put them out to your sides. You see, it's nothing.

So we make these experiments such that if you just flap your two hands down on the buttons simultaneously, what you get does not make the balls land on the red oval. Watch: press both buttons simultaneously, and that's what you get. It's no good. They do not land on the green. What you get is a full cut-and-paste combination of the two scenes. Nothing is left out. If you take everything from one and everything from the other and you put them together, no solution. So what's the solution?

(Audience calls out the answer and demonstrates with gestures.)

Exactly, first press the red button then the blue one. The red ball is higher up; it has a longer way to go. So you need to stagger the pressing of the buttons.

You understand how to solve this before you have seen a demonstration, because you can selectively blend the two scenes, in your imagination. In imagination, you hit the two buttons, but not at the same time. You press the red button so that the red ball rolls down, and then you press the blue button so that the blue button starts to roll. I have not shown you such a demonstration, and I did not talk about that blend, but you made the blend, and it gave you the solutions. You thought of this blend. You can come up with solutions that you have never seen. And you can do this not because there's a generic module for pressing buttons in a blending box. Nobody has ever seen this before.

Here's a demonstration in the slide show of what you are all thinking: see, the red ball rolls, then the blue ball rolls, they collide in the middle, and fall down, landing on the green oval. You have performed mental magic in seeing this before being shown this. Your mental performance seems to you like nothing, obvious. The reason that it seems like nothing is that human beings are blenders. Blending is what you do. You can put things together selectively, imaging a simulation, and not bringing in certain kinds of things. Notice that the trajectories of the balls are not brought in, the timing is adjusted, projected in a certain way. And you get emergent structure, namely, the balls bump into each other and land on the screen and you get a million yuan. Or if you are a dog, you get the yummy treat. Or if you are a bird, you get whatever you like. This is actual extremely inventive, extremely creative, but it's what we can do all the time. And I'll show you more. Notice again that there is selective projection and emergent structure.

Remember that this morning I said that human beings are not like other species, they can conceive of things that run way beyond human scale, beyond local time, beyond local space, beyond just the few agents that you just see right here in the situation in which you are interacting. I will give you an example.

Again, so far we are just working with mirror networks. We've not got to a double-scope network in these little examples. There was a man, Hicham el-Guerrouj, and in 1996 he set the world record in the mile. He ran the mile the fastest: Hicham el-Guerrouj. And the *New York Times* wanted to convey what his accomplishment consisted of. So they presented this graphic, which you see in the slide show. By the way, it is of course already blending that you can look at these little black lines and regard them as a representation of a race. Nobody is running. I can point at these little static black figures and say, "these people," and you are not bothered, because in the blend, they are people. You are not deluded, but in the blend, they are people. We will get to that.

Just for now, these people are the fastest milers from each of the preceding six decades, except Roger Bannister and Herb Elliott, who are both from the 1950s. Roger Bannister has to go in because he is the most famous miler, the first person, I think, who broke the four-minute mile. He's very, very famous. So this little figure is Hicham el-Guerrouj in 1996, this one is Steve Cram from 1985, this one is Sebastian Coe from 1979. What the *New York Times* has done is place these little symbols of the runners on the track where they would have been if they had run against Hicham el-Guerrouj.

Of course this takes some calculation. I'll work it out. They took the time of the runner for the mile and assumed that the runner was running at a constant speed and that's how they figured out where to place the runner on the track. Of course, milers don't run at a constant speed, so this is quite inaccurate in that way. Now you understand because of this blend the accomplishment of Hicham el-Guerrouj and how fast he is. It is a mythic race. Notice you are not deluded. You don't think that Roger Bannister ran against Hicham el-Guerrouj. They are 40 years apart, 45 years apart.

Once we have this blend, we can say, because we already have language for the blend, "*Oh, Hicham el-Guerrouj was really wonderful. He defeated Roger Bannister by 120 yards.*" That language is grammatical for the blend. It is not true of the inputs, because Hicham el-Guerrouj and Roger Bannister never competed against each other. Where is this 120 yards? It's only in the blend. Moreover, Hicham el-Guerrouj never defeated Roger Bannister. He never defeated any of these people in the inputs, but in the blend he defeats them.

Let's look a little bit at this selective projection. This is a "mirror network". Mirror networks come up often. It is a mirror network because each of these input spaces, and there are six of them, have the same organizing frame: a one-mile foot race. For the blend, one person is taken from each of these six inputs. Notice that, in each of those inputs, the person selected for projection to the blend is not just the winner. That person is the world record holder when they cross the line! Each of these people is a world record in the mile. But now, Hicham el-Guerrouj gets projected down into the role for winner in the frame for the blend, and, luckily, there are other roles for projecting these other winners from the other five input spaces. They are projected into the blend as losers! Up here they are the winners and world record holders, but now in the blend, they are losers. That is emergent structure. All of this is emergent structure. And we have language for referring to the blend that lets you understand the

cross-time complicated network, the relationships across all of these people.

With this language and with this human scale blended frame, we can say things like “*Hicham el-Guerrouj beat Roger Bannister*”, and “*Hicham el-Guerrouj defeated Roger Bannister by 120 yards*”. This language comes right up for the blend. It’s completely false for the inputs and in the inputs you wouldn’t understand it. But now you have language, not just a conceptual network that can be understood by grasping it through the blend. It’s a mythic race. You also have language that lets you understand the network without having to invent new constructions, new grammatical structure. This blend presents a kind of fictive interaction.

Fictivity is very important in human conception. It is in a way a shame that we use this word “fictive”, because we think of “fictive” as the opposite of “true”. The word “fictive” should almost be retired. Fictive conceptions are really important for people in understanding the truth. They may be fictive in part of the network, but they help you grasp the truth that runs across the network. These people on this track never actually interacted with each other, but you’re built to understand human interaction. That’s what you’re built for, to understand human interaction with a few agents. Construct the blend that has fictive interaction and you can understand it, and that let you understand the network. There are many different kinds of fictive interaction. Here, in this slide, are several references to work by Ana Margarida Abrantes and Esther Pascual, who work on fictive interaction in legal discourse, and trials, in theatricality and everyday life, and so on. We very frequently understand what’s going on in the outer-space network by a kind of fictive interaction or by fictivity in general. The Buddhist monk example, interestingly, is fictive interaction. In the blend there is an interaction between two people. There was no interaction in the input spaces. That interaction gives us this human scale frame.

Here is another fictive interaction. A modern philosopher is teaching a class, and he says:

I claim that reason is a self-developing capacity. Kant disagrees with me on this point. He says it's innate, but I answer that that's begging the question, to which he counters, in Critique of Pure Reason, that only innate ideas have power. But I say to that, what about neuronal group selection? And he gives no answer.

Note that it’s pushing it a little at the end there, because of course Kant did not know about neuronal group selection. The important thing is that there is no Kant in the classroom, when the modern philosophy is talking. Kant is dead. He is not there. One of the problems I have in talking about conceptual integration is that, since I need people to see what is going on, I give spectacular, pyrotechnic examples, like the Buddhist Monk and the Mythic Race, and these examples are misleading. Because when you look at those, it’s easy to think “Oh I see, conceptual integration is a goofy thing that happens in Saturday morning cartoons and in wild stories.” No, these are only the examples where you can see the blending. Notice that this way of talking about disagreeing with Kant, this is just the normal way to talk about it. As far as I can tell, when philosophers talk about the way they study and what they think in relation to

previous thinkers who have worked on the problem, the basic language that they use is this. This does not look exotic or bizarre to us at all. We say “*Kant disagrees with me*”. Now of course, Kant was dead before the philosophy was born. Kant cannot disagree with this man. And you know that, but in the blend, there is Kant and the philosophy and they are having a debate. In fact, Norman Melchert wrote a textbook called *The Great Conversation*, which presents the history of philosophy through such blends. This is a standard way of understanding a sequence of analyses by different philosophers. We are not deluded, but there is a fictive interaction in the blend and this means suddenly we have language to refer to the blend and let us understand the whole network. We can say things like “*disagree with me*”. That’s grammatical for the blend. “*He says it’s innate, but I answer . . .*”. Where does “*answer*” come from? The modern philosopher cannot, in the input spaces, answer Kant, but in the blend, he can. You know what that means; you can project it back to the network to which he counters. Kant cannot counter something that the modern philosophy says in the input spaces. But in the blend, he does, and you have this language: “*But I say to that, . . . , and he gives no answer*”. You have available all of the language for describing two people having a debate or a disagreement with each other. Indeed, all of the conception comes right up in the blend and allow us to understand this entire network. The riddle of Buddhist monk looks like this. There is a modern philosopher, there is Kant. The works of Kant are in the space with the modern philosopher but Kant himself is of course not. The modern philosopher knows that Kant existed. But Kant does not have knowledge that the modern philosopher existed. There’s a debate frame that one can recruit, and now you have a conceptual integration network.

I challenge you: go look at the newspapers, go look at books, you will find this kind of fictive interaction blending all over the place as soon as you leave this hall. So, for example, in France, there is an electronic summary, *The News*, that’s published in the newspapers. Newspapers have opinion columns. These opinion columns are written by journalists overnight and published in the paper in the morning and typically none of them knows what the other ones are going to say. They wrote them without having a debate. But the summary in *The News* of commentary on a particular point of government or politics uses just this conceptual integration network of debate. So one columnist will assert and, in the blend, another will contradict. One person will say to the other “*No, you are wrong*”. This is how the summary goes. This is in French. In their actual columns, they are not addressing each other. But if you look at the language of the summary in *The News*, you see—and I will read a translation in English here—that they are addressing each other. Columnist A is summarized as asserting, “*We should never negotiate with terrorists*” and columnist B is summarized as responding, “*On the contrary, you are wrong, that’s exactly what we should do*”. The summary uses all the discourse markers of debate. *On the contrary* comes up, available for the blend, so that we can understand this network.

This use of the debate blending network is extremely common. You find it even in places where no speakers are involved. Here’s an example: *The bean burrito is California’s answer to France’s Croque Monsieur*. You have this fictive interaction between different people eating different foods. You have your Croque Monsieur,

which is a sandwich, but I have my bean burrito. Another example is *Stag's Leap Chardonnay is California's answer to Corton-Charlemagne*. Stag's Leap is a winery.

Fictive interaction, fictive debate, fictive actuality even, are standard ways of understanding a mirror network involving agents who disagree or have disanalogies between them.

Again, what is happening in these cases is this: you have input spaces and vital relations between the input spaces. The generic space is possible. Dynamically, the generic space captures whatever vital relation connections are in place between the inputs at any given time of processing. You have selective projection to the blend, emergent structure, and cohesive structure in the blend. The emergent structure is the mark of creativity. The inputs and their relations become intelligible because they are anchored in the blend. Making a blend is a way to get “something more out of nothing but”; it is a way to start with things at one level and of one kind and to generate mentally things at another level and of another kind.

Let's look at double-scope networks. In summary, so far, we have looked at the basic mental operation of conceptual integration, which, in its rudimentary forms, we seem to share with many mammals. We have looked at the way in which blending connects mental spaces, projects selectively to a new conception, a blend, where there is emergent structure. And we've talked about many different purposes here. Some of the things we've seen are that blending allows us—or, rather, some of the things we've suggested and that we will be exploring are that blending allow us to understand mental networks that otherwise would be beyond our grasp, that otherwise would be beyond the scale we're built for. Blending allows us to understand what we are not built to understand, except that we have this blending ability. And blending also allows us, as we will explore, to take a relatively small system of grammar—grammatical structure is very complicated, but it's nothing compared to conceptual structure — to take a relatively small set of form-meaning pairs and to use them ad hoc potentially to talk about anything.

Now we are going to look at double-scope networks. Remember I said a mirror network connects inputs that share a frame. A double-scope network arises when the inputs have conflicting organizing frames, organizing frames that conflict on central organizing structure: causality, participant structure, modality, things like that. You might think that if you have two things that do not go together, that should be never put together, that you should never confuse, then that conflict would stop you from blending them. Indeed, this seems to be what stops other species completely cold. But on the contrary, putting things together that do not go together is what human beings really specialize in, mentally. The conflict, far from even being noticeable, seems to prompt us to use our conceptual integration capacity.

The most popular book in the world for three-year-olds is *Harold and the Purple Crayon*. Harold has a crayon. Here, in the slideshow, you see Harold, and here is the purple crayon. Here is the interesting thing—in his world, whatever Harold draws is real. So he wants to go for a walk. He needs light to go for a walk, so he draws a moon, and here is light, and the moon stays with him. And he needs a road and so, as it says — remember that this is a book for three-year olds — “he made a long straight path so

he wouldn't get lost. And he set off on his walk taking his big purple crayon with him". So what he draws is real.

What's going on here? Let's take the moon. There are Harold and the moon in the slideshow. In the mental space for the actual moon, you know that the moon gives light when you walk. Because it is so high, it sort of seems to stay with you, or you don't get away from the moon. You know the moon is not created by human beings. It's certainly not created by drawing.

But over here, in the mental space for drawing, you know about drawing. (It's already a blend, by the way, that this drawing can be a representation of the moon.) You know about drawing that you can make a sketch—the sketch does come into existence because a human being made. But the drawing of the moon doesn't give light. Now you project those two things—the moon gives light, and you can draw a moon—to the blend. Let's stay here for a minute. You project those two things into the blend, and now you have the moon that can be drawn, but it gives light and stays with you.

Such a thing, you have never seen. You have never seen a drawn moon that gives light. It's brand new. Nobody taught it to you. You haven't learned it. You can't produce it. But you have no trouble putting it together, mentally, and it's not because you have a genetic module for understanding that drawn moons give light—because they don't. The blend is false of our world. You should not have this if your conception is driven exclusively by reflection of the real world.

Here, in the blend, you have a world with a completely different physics. You have no trouble, and a three-year old has no trouble understanding it, either. So Harold goes on his big walk. He has a lot of fun. But like all three-year-olds who've had an adventure, he gets a little nervous, he wants to get home. But he cannot get home. He needs some food, so he draws the tree, and he eats the apple. But he does not want to leave all this uneaten fruit. So he draws a dragon to defend the apples; then he gets afraid of the dragon, and so he is backing up but his hand is shaking. And so that's water—the line he draws inadvertently with his shaking hand is water—and he falls in it, and he has to come up to draw a boat. It's very inventive for this little three-year-old. Children have no trouble understanding this. But finally he wants to get home and go to bed.

Harold's going to bed prompts for another blend. It's important. I'll call it, whimsically, a parent's blend. All parents want their children to go to bed. And you want the child to whom you are reading to blend himself or herself with Harold. Harold wants to go to bed. You want the child to go to bed. This is the blend in which Harold's wanting to go to bed is your child's wanting to go to bed. It is a kind of persuasion. All these books end up with kids going to bed, because you want the child to whom you are reading to go to bed. You do not only want them to understand that Harold wants to go to bed. You want them to make a blend.

Harold wants to go home. He cannot get home. He doesn't know how to get home. Who here is as smart as the three-year-old? How does he get home? (Answer from audience.) Yes, pretty much: he just draws his home. He remembers — but you know, it is a nicer story than that — he remembers that through his bedroom window he can see the moon. He's got the moon and he's got the purple crayon, so he draws his

bedroom window around the moon. And ipso facto, presto chango, he is home.

Now notice that this is action is translocation in space by drawing. This is a physics you've never seen. It is an emergent physics. It is a creative physics. It is an innovated physics. No three-year-old I know has any trouble understanding this, and moreover they don't have any trouble continuing the story of Harold, to say what Harold might do with his purple crayon.

This is a double-scope integration network. It's a dramatic one. Again, you might be thinking, "Oh I see, this is just a mental operation for goofy kids' books". On the contrary, the examples I am presenting here are the ones where I can make it clear that there is a blend going on. All these same features are going on in your head all the time. Almost all blending goes unnoticed. Blending is not special and it is not cognitively costly.

What has happened here is that there is a mental operation that has created new structure. What we need to ask is: what are the operations, mechanisms, principles, and constraints of this kind of creativity? Why should we have them? What's the neurobiological substrate?

There is an elaborate scientific effort to answering these questions. Some of the science is embryonic. Some of it is much further long than the other parts. It is an open project. And of course we hope many of the great insights will come from you. Throughout this series of talks, we'll be talking about certain features of this mental operation.

There are some overarching goals to conceptual integration. One is, Achieve Human Scale. When you're making a blend, try to get something that human beings are built to understand. Compress what is diffuse. Obtain global insight. Strengthen vital relations, which are connections between mental spaces. Come up with the story. Go from Many to One. "Go from Many to One" is an abbreviation for taking a vast range of things, too many to be contemplated, and compressing them into something that you can manage in such way as to give you insight into the Many. I will be going over many examples of that. This is a preview of coming attractions.

There are constitutive principles: matching counterpart connections, generic spaces, blending, emergent meaning. We can compose things. We can complete things. We can elaborate them in the blend.

There is a range of vital relations that we'll be going over. These relations have names like change, identity, time, space, cause-effect, part-whole, representation, role connectors, analogy. So, for example, just to give you a taste: between the moon and the drawing of the moon, there is a representation link and another analogy link, because the drawing of the moon has topological analogy to the shape that you see in the moon. This is extremely common. This blending network, for Harold and the purple crayon, takes a representation link that also has analogy link and compresses them into an identity in the blend.

There are governing principles for compression. Many to One involves compression. In compression, you take something that is too diffuse and compress it into something that you can understand without losing the network. One kind of compression involves borrowing. I won't go through all of these here, this because

we'll be looking at these again and again and again, many different kinds of compression, many kinds of results from compression. There are networks of compressions and conceptual blending that we find again and again and again.

Here, in this slide, is one kind of compression network. Here is another. You don't need to grasp this yet. We'll be looking at examples and I will call them up. The governing principles are optimality constraints, in the sense that they can compete each other. One of the governing principles is the Topology Principle. For example, the Buddhist monk compression will not work if you did not give top priority to the Topology Principle. You have to keep the space-time relationships in the inputs fixed as you project them down into the blend. If you don't do that, you don't solve the riddle. There are other governing principles, such as the Web Principle, and the Unpacking Principle. Each of these has a definition that we can go through and that we will go through in the coming days.

There are different kinds of networks. Conceptual integration is an operation that has principles. It works over all conceptual spaces. We have not found a conceptual domain in which it does not work or does not work according to the same principles. You can have a conceptual integration network that has lots of and lots of space. Some can become inactive. You might recover the inactive ones later on, and so on. When I show you in the slideshow just these four spaces, or I tell you there are certain kinds of networks, I am giving you a cognitive reference point on which to focus. I am not giving you the entire network, which is usually quite large. Ten is a cognitive reference point for all the real numbers, but there is an uncountable infinity of numbers on the real line. This four-space diagram is a cognitive reference point for blending networks, but there is an indefinite number of networks, of different forms. This sketch of a few spaces is giving only the highlight on which to focus. Some standard kinds of networks are simplex networks, mirror networks, single-scope networks, and double-scope networks. But these are only cognitive reference points. It's not the case that all networks belong to just these categories. On the contrary. Blending is a process, and it works over diverse conceptual networks, and I am only picking out a few standard sorts of networks in order to discuss the principles of blending.

Tomorrow what I will begin with is a literary example of a double-scope network from Racine. We will actually look in some detail at the language that is used in that passage and how the conceptual integration network that is put together makes it possible for the network to be commanded and understood by using expressions we already have, to express new meaning that comes up in the blend. Preview of coming attractions.

Later on in this set of lectures, we will be talking about the way in which conceptual integration makes language possible by making it possible the blend form and meaning, by making it possible to take different constructions and to blend them into a construct in a way that is selective, what has emergent structure. The operation of blending makes it possible to solve the central problem in the development of the language, which is how to use a relatively limited number of constructions with equipotentiality so you can talk about anything that you want to talk about.

The beginning parts of these lectures are going to be looking at the mental operations and the second half of these lectures are going to be looking at the role of these mental operations in language. This is what is characteristic of Cognitive Linguistics territory — language is a branch of cognition, and if we want to understand language, we need to look at the way in which human higher-order cognition makes it possible for language to come into existence and to operate. This is not to say that any one of these abilities—art, music, dance, social cognition—might not develop some structure that is more specific to that particular domain.

I should say that on the subject of the evolution of human higher-order capacities, I am in a minority. Most evolutionary stories about how human beings became what they are begin with the idea that one of these abilities was all-important and you developed that one ability and that ability it made it possible for you to do the others. So a common candidate is tool use. In this view, human beings use tools and make tools, and at some point they developed this ability for tool use, and this gave them the neurobiological computational power to do other things, and so suddenly you can have syntax. Another common candidate is the manipulation of the hand using an opposable thumb. Once you have the computational ability for that, then you have the right kind of neurobiological patterns to be able to do other things. Another common candidate is social cognition: we learned how to coordinate with others and we developed Machiavellian intelligence, and once we had social cognition, then the other things came along. Another common candidate for the all-important first ability is of course language. In this view, language was the big break-through. We were basically just like all those other apes, but we got language and then all these other things became possible: social cognition and so on.

The view I take — and I am in a minority here; this is an assertion; this is controversial— is that no one of these abilities came first, no one of these abilities ran up to perfection and maturity and then made the others possible. On the contrary, it is not anyone of these abilities but something behind them that makes us human and gives us all the abilities. This underlying all-important ability is the ability to make integrating networks of a slightly higher kind than any other species could make. Once we had that advanced mental operation—double-scope blending—, then all the other things became possible partly because they can come up together, they can feed off each other, and bootstrap each other.

Remember the Cub Scout pack —*The Cub Scout Helps The Pack Go*, and *The Pack Helps The Cub Scout Grow*. In these lectures, we are starting to study language by starting with the mental operation of blending, because that, that I think, is where the story starts. We will continue in later lectures to advance into the structure, function, and development of language phylogenetically and ontogenetically. Thank you.

Lecture Three

Double-scope Cognition

Today we're going to talk about double-scope cognition. I thought I would begin with a little history about the development of the theory. It began in earnest in 1992 and early 1993, when I was at the University of California, San Diego. I had been working on various problems in conceptual projection and so had Gilles Fauconnier. I gave a presentation at Ronald Langacker's house, and Gilles was in the audience when I talked about certain data I had found that could not be analyzed on any existing theory of conceptual projection, including my own. This data caused me difficulties and I indicated that I thought we needed some new theory and sketched some tiny beginnings. Gilles Fauconnier immediately stood up and said that he had similar examples and presented them. So we decided the next day that we would write a little article. But as soon as we started to write this little article, it exploded and now this collaboration has been going on for 17 years.

In that time, hundreds of other people, maybe by now a thousand, have contributed to the theory, sometimes quite substantially, not just in producing applications of the theory but also in offering theoretical developments. If you want to see the history of these things, go to <http://blending.stanford.edu>, and there you will see hundreds of articles and citations by people all around the world working in the theory of conceptual integration in a variety of fields.

You see in this slide the standard representation of blending, with four circles and some lines. This is a typical network, but the theory has nothing essential to do with four spaces or five or seven. In fact, what the diagram looks like does not matter at all. It's just a mnemonic device. Conceptual integration is a basic mental operation. It happens all the time. It happens according to a set of constitutive and governing principles and it happens over all conceptual networks. I put this diagram up simply so I have something to point at while I am talking.

In blending, there are always at least a couple of input spaces, sometimes many more. There are always some kinds of connections developed between the input spaces. There is always selective projection to the blend and emergent structure there.

You saw the example of Buddhist monk. Here it is again: the Buddhist monk gets up at dawn, walks up the mountain path, reaches the top at sunset and sits down, meditates over night, gets up at dawn, walks down the mountain path, reaches the bottom at sunset. You are supposed to prove, without assuming anything about how the monk moves, that there is a spot on the path that he occupies at the same hour of the day on the two separate days.

One way to show that there is such a spot is to blend the ascent and the descent selectively. That is, you do not bring into the blend the calendrical day. You do not bring in the usual psychology of someone who would be terrified to see himself, approaching himself on the path. You do not bring in normal physics. You bring in the monks but you bring them in separately. You bring in the path but you fuse the paths. And there is

emergent structure, namely, the monks meet. They must meet somewhere and where they meet is the spot on the path they occupy at the same hour of the day on the two successive days.

There is emergent structure in this blend, namely, the meeting, which solves the problem. The important thing in this case is not so much to have the blend as it is to use the blend to help you understand relationships in the network that you didn't previously understand. The purpose of the blend is to run the network. It's crucial in this case to use the blend to understand the relationship between the inputs.

In fact, in any conceptual integration network, what counts is always the entire network, not just the blend. Once you have this blend, you have language for referring to the entire network. You can say "*The place is where the monk meets himself*". You have this word "meet", and you have this word "himself". And they are perfect for the blend. They are ungrammatical for the input spaces.

In the original story of the Buddhist Monk's journey, there was no meeting and there was no "himself". It's grammatical to use this language of the blend and so you can talk about the blend in a way that lets people understand the network.

In a blending network, there are always input spaces. There are always vital relation connections between the input spaces. There is a generic space, often not recognized but frequently very useful, that contains what ends up being the shared structure.

In the network, it's not the case that the generic space is always there to begin with. There is projection to the blended space and emergent meaning in the blended space.

I said that there were types of networks. Today, what I want to do for about the first 15 minutes is spend a little time avoiding mistakes, avoiding misimpressions. When I show you these blends, they look goofy, strange, bizarre. The Buddhist Monk is meeting himself. There are riddles. This strangeness can be misleading. Blends are almost never strange and they are almost never noticed. They are what we do all the time.

Similarly, when I say that there are different kinds of networks, this can mislead. It's common for people who study blending to look at the network and say, "Oh, now is this a simplex? Is this a mirror? Is this a single-scope? Or is this a double-scope network?" They ask this as if the network must be one or the other. No, blending works over conceptual networks, all kinds of networks. But there are certain basic patterns that arise again and again. They are strong representatives of conceptual integration networks. So think of it, as I said before, like numbers: there are some reference point numbers, like zero and one hundred, negative one hundred, too. These are cognitive reference points, but there are many other numbers. There are numbers all in between zero and one hundred.

In the same sort of way, there are plenty of blending networks that are not exactly simplex, not exactly mirror, maybe a little of both. There are networks that are not exactly single-scope. They are not exactly double-scope. But it's good to have these as reference points.

For example, a simplex network is typically what, in the history of psychology, people have called framing. You might say “Wait a minute, if we already have a term for this and people have already recognized it, then why place it inside the theory of conceptual integration?” And that’s a major point. The point is these things are not separate operations. They are aspects of conceptual integration. Framing is one case of conceptual integration. The various cases of conceptual integration should not be taken to be separate phenomena.

A simplex network is a network with a frame, as one input, that is built to apply to a certain kind of conceptual domain, and the other input has suitable elements from that conceptual domain. For example, we have a frame of kinship and it can vary from culture to culture. But in our frame of kinship — this is the one that I have — there is always an ego point. That’s the I. That’s where you measure things from. And you have things like father and mother. There are no specific people in this kinship frame. It is a frame that is designed to apply to specific people. You can take particular people, like, say, Paul and Sally—notice in the slide that there is an input with just two people, Paul and Sally. In this input, the input on the right, there are Paul and Sally. And on the input on your left, there is a kinship frame that has ego and father and mother and so on. And you can integrate these two inputs, so that in the blend, Paul is the father and Sally is the daughter. This is a simplex blend and it seems to be one that many mammals can do.

For example, if you have a dog that has learnt from its — I should say first that dogs can be a misleading example, because they are domesticated; they have been bred to serve us or maybe we’ve been bred to serve them; that seems to be what goes on in domesticating; there’s combined pressure on both species. Well, to resume, if the dog has learned to play fetch with only its master, but another human being shows up, then the dog seems to be able to understand —we cannot interview the dog, but from its behavior it seems that the dog can understand—the dog seems to be able to understand that another human being can be in the role of someone who throws the stick. Indeed, the dog will get the stick sometimes and try to put it in your hand. That’s a simplex framing. It’s putting somebody new into a frame the dog has for playing fetch. And it doesn’t seem to be very difficult for the dog. It’s actually magic. It’s really, really, really impressive. But it looks very, very simple, very, very simple to us. There is emergent structure in this blend. It just doesn’t look to us like a very big deal. The emergent structure in the framing of Paul and Sally as father and daughter is that now there is a particular role in the blend, *father of Sally*, not just *father* but *father of Sally*.

And in fact you can put different people into that role. So, you can say, “*Oh, Paul is Sally’s new father.*” Maybe Sally’s mother was married to someone else ten years ago and now she is married to this man. And he has certain responsibilities in this role. He might be responsible for paying her parking tickets or he might be able to excuse her from school. This is a new role in the blend, *father of Sally*, which is not in either of the inputs. It’s not in the kinship frame input, because there is no Sally there. And it is not in the input with Paul and Sally because there are no kinship relations

there.

Now, you may be thinking, “What? This is so basic!” Indeed it is! It is very basic. Simplex blending looks basic to us. It doesn’t look creative and amazing, like the blend underlying Harold and the purple crayon, but Simplex blending is still run according to the constitutive and governing principles of blending.

It is easy to run up the complexity of blending very quickly, so that you go from simplex to double-scope. Please notice that there is a cline: it is not that something is either simplex or mirror or single-scope or double-scope, but that there is a gradient along which they lie. Networks do not fall into just four pigeonholes.

We can say, “*Zeus was the father of Sarpedon. From Mount Olympus, he watched his mortal son die.*” Now Zeus is immortal, but Sarpedon is mortal. In this case, we have an interesting blend in which of course the father is going to outlive the son and we have a son who has immortal father.

Zeus is the father of Athena. She was born out of his head, fully clad in armor. Notice that in this case, Athena has no mother. She does not have a childhood. She comes out of his head. It is an interesting blend, because his head is sort of a container, in a way like a womb in a woman that is a container. Athena came in from the birth frame, but not the woman and not the womb, and not the mating. The usual kinship frame has got father and mother. But in this blend, only the father comes down, not the mother.

Joseph was the father of Jesus. Now notice that in the story, Joseph did not impregnate Mary, but one can say, “*Joseph was the father of Jesus*” without any difficulty in many communities. What it means is that he is married to the mother. He has certain kinds of authorities and responsibilities.

The Pope is the father of all Catholics.

The Pope is the father of the Catholic Church.

George Washington is the father of our country.

Newton is the father of Physics.

Fear is the father of cruelty. Notice that in this case, there are no specific people at all. People are implied because it is people who fear and are cruel. But *fear* is in the role of *father* and *cruelty* in the role of the *child*.

You can do a great many things with the kinship frame. I wrote a book, *Death is the Mother of Beauty*, exploring what you can do with the kinship frame. What you can do by integrating a kinship frame is all over the lot. You can say, “*Paul is the father of Sally*” or “*The child is father of the man*”. That expression may seem very confusing, but one way to interpret it is that the man is influenced by whatever his childhood was. And the child then is formative for the man that the child who grows into. So, even the same person at different spots can be framed as father and son.

There are other kinds of networks, such as frame-compatible networks. Sometimes what you are doing is taking a frame that applies to a domain that it’s meant to apply to—that’s a simplex blend. Or you can blend two frames that clash but put one of them frames in control of the other—that’s single-scope. Or you can blend two frames — this is double-scope blending— two frames that clash and you are blending

them by taking parts of each frame for a new frame in the blend. Sometimes, what you are doing is taking one frame and another frame and they nest very nicely—that's frame-compatible blending. So, suppose you have a frame for *sports event* and you have in the *sports event* frame a *fan*. A *fan* is somebody who likes this sport and pays attention to it. And now you develop a new frame for a new kind of sport. Maybe, it's bicycling. It could be anything, frog jumping competitions. And here is the frog jumping competition frame, and over here is the fan frame, and you combine these frames, and now you have a fan of frog jumping competition and there is no clash between them. They combine. They nest.

You can see a frame-compatible blend in spy novels, when the convention is that when the culture exchanges happen, sending the symphony orchestra from one country to another, you always included a spy, so you could say something like "*The last second violinist is always the spy*". And there is no inherent clash between the frames. The person could be both the last second violinist in that music frame and the spy over in the espionage frame. In the blend, the role for last second violinist and the role for spy are identical.

We saw that one of the great things that blending can do is create a blend that is at human scale—a blend that is built for the way we operate and the way we think. Using that blend, we can understand a network that is complex and not suited to the basic patterns of human understanding. We can understand the network partly because it has the blend.

Once you have a blend that takes in part of the network and compresses it to human scale, using certain kinds of selective projections, then the creative blend has a new status: it is at human scale. It wasn't before, but now it is. And you can use that blend again. Often, what you get is blends becoming inputs to other networks, and blends of those networks becoming blends in another network. Through this recursive process, you can build human understanding all the way up.

Think of something like numbers. We use complex numbers, hamiltonians, quaternions, imaginary numbers, and so on, but we do not start off the child by saying "Oh, we are just going to teach you complex numbers right at the beginning—no point going through that old history of learning counting numbers first!" No, the child first learns counting numbers. It's a blend of walking and numbers of things in the basket. And then you have to introduce the idea of zero. Zero came up very late in the history of mathematics. And it's an achievement for a child. But once you have zero, in the blend, it becomes a number and seems as if it is at human scale. And you can walk the child up just as history walked up through proportions to fractions. That's not so easy. Many, many children have difficulties with fractions. Many human beings have difficulties with fractions and there is a reason for it. Rational numbers violate a lot of the preferences in blending. Rational numbers go against some of the governing principles (not, of course, the constitutive principles⁰). Anytime you find culture spending a year trying to bang something into the heads of children, you can tell yourself this is not congenial to basic human mental conception. And once you get rational numbers, it seems to us as if, "Who could have any problem?". But in fact, if

you look at *The Way We Think*, we explain in the book why people have problems learning rational numbers. In fact, rational numbers are the product of a very inventive blend with emergent structure and so on all the way up. Once you get the blend, it seems like something that one should understand. It seems at human scale. It's second nature.

Because of double-scope conceptual integration, we, unlike any other animal, have the ability to acquire not just conditioning, but also a kind of second nature. We are able to take things that are not at human scale and make them feel as if they are at human scale, like reading. Reading has been around at most only 8,000 years. It is not an evolutionary product. It took human beings at least 40,000 years to develop writing and reading. It was not easy. Nobody thinks there is a module for writing and reading. But once you got it, you look at writing as if it is completely straightforward. When you look at the letters on the page or when you look at the marks on the page, you can no longer see it just as marks on the page. It's words. You can't look at the writing now without seeing the words. It's human scale. There it is. It seems completely natural.

We get figurative interaction in these blends. You remember the standard way of talking about the development of ideas. We say things like, "*Kant disagrees with me on this point*". There is no such disagreement possible in the input, because Kant is not here. This person is not here. In fact, Kant is not even aware of the modern philosopher. In the blend, we have a fictive interaction and that lets us understand the network. Here is a slide showing the blending network for the Debate with Kant. This is not at all an unusual thing for human beings to do. On the contrary, they do it all the time. It's the basic way. You saw fictive interaction in the six runners from different decades all put together on one track so that in the blend you have a human scale situation. It's really quite elaborate. It gives you language for referring to the network. You can say things like *Hicham el-Guerrouj beat Roger Bannister; defeated Roger Bannister by 120 yards*. Again, these things are completely ungrammatical for the input spaces, but they are grammatical for the blend and we use them to understand the network.

You saw double-scope networks, in which you have frames that clash. In *Harold and the Purple Crayon*, you saw a frame for drawing and a frame for physical reality and a blend created that takes parts of each of those frames to make a new blended frame with emergent structure. A single-scope network is where one of the frames controls and subordinates the other. Material from other frame comes in, subordinated to the frame of the other. That's single-scope. I don't think there are really any pure single-scope networks, because human beings, as soon as they start, as soon as they have a single-scope network, they immediately start running the blend and produce double-scope structure in there. A double-scope network is one in which the frame structure for the blend takes elements from the frames of each of the inputs. For example, in the blend in *Harold and the Purple Crayon*, you now have the ability to move by drawing. This is a kind of physics you've never seen before. If you want to be

home, you draw your house and you're home. You want light, you draw the moon and you have light. That is emergent structure.

Now, this happens all over the place in art. Here, for example, in this slide, is Magritte's *Tentative de l'impossible*. In our blend for understanding this painting, what the painter paints is real. That's Magritte looking at his wife. As you look at the painting, you have to sort of decide what the selective projection is to the blend. Is she aware that she is coming into existence, for example? Or is she just inanimate until she is completed? There are similar kinds of decisions you have to make in construing the network in Harold. For example, if you make some art by mistake, does it account? You know we make erasures when we are drawing. There are no erasures in Harold's world. If you draw it, it is real. So without intending to draw something, he can draw it and then he has to deal with it. What's happening in these kinds of networks is that there is a representation link. The element and representations are blended into a unique element in the blend. When Harold draws the moon, the drawing is the moon; in the blend, the drawing is the moon. There is also often an analogy link between the object and its representation, not always by any means, but often. I will talk about all these various different ways of representation. And those representation links and analogy links can get compressed into a unique element in the blend.

I'm going to spend a little time on this next slide. Blending is basic, but because of the way I have introduced it so far in these lectures, where my emphasis is to get you to see that blending happens, and to see its principles, it may seem as if blending is something special and strange. In fact, my claim is — and you will be seeing this in the rest of these lectures — that there is nothing that you do that's higher-order cognition that's possible without conceptual integration. Even look at each other and understanding each other, my making a gesture, coordinated with my language, syntax, meaning, polysemy—all of these are possible for us only because we are able to do higher-order conceptual integration. You are always using blending all the time. Animals have abilities to do some kinds of blending and they also have capacities like basic narrative understanding. They can understand that here are some agents they have to interact with, and that those animals might be goal-directed. Chimpanzees seem to be able to understand that other chimpanzees are goal-directed; they seem to understand basic force dynamics. They have that. That's about as far as they go, it seems. They don't have the ability to understand false belief. There is no indication that they have any sense of personal identity that stretches over 50 years. They do not have the kinds of things we think of as a story. They do not have the kinds of abilities we have for language.

Basically, we don't see anything about our cognition, because we are not built to. Our mental operations are built for certain real environments and they are built for us to succeed. In cognitive science, what we are trying to do is to take those mental abilities and turn them to analyzing ourselves, but those mental abilities are not meant to analyze thinking. They are not meant to analyze language. We are tricking them. We

are trying to take things that are useful for something else and manipulate them to see a little bit into thinking. So, the cognitive scientist is always in the hard position, because consciousness is a thin, weak, tiny, little ability, good for certain kinds of things, like staying on task, focusing on something that you need to learn in a way that becomes more secure, like learning to read. But whatever you were doing in learning to read is now behind you; it has been installed in backstage cognition. You do not want to try to read in consciousness. If you did, it would take you forever. Almost all the important things in thinking are happening in backstage cognition.

Backstage cognition is very complex, but it doesn't seem that way to us. We look at colors; they look straightforward. The yellow looks as if it comes straight in my head, goes in my eye, and registers in my brain. Now, nothing like that is happening. There are no colors in the world. We create colors in our brains because we want to have stability in the world. The apple that is red at noon needs to look red at dawn and at dusk, so we can locate it and eat it, despite the fact that the light is radically different at dawn, noon, and dusk. We discount the illuminant. This is not controversial. Suppose I hold up a painting that has little squares of color, red and green and so on, and we are in a dark room, and I shine three projectors of light on the painting, and the three projectors project the wavebands of light to which the three kinds of cones in the eye respond. (There may be some women who have a fourth, so perhaps there is a reason for marital disagreements about how the coloring should be set on the television.) Suppose I set those three projectors to the settings that correspond to normal white light, and when I do, you see that this rectangle in the painting looks red and this rectangle in the painting looks green. If I take a spectrometer, and I measure the light that is reflected from the red for each of the wavebands, and I measure the light that is reflected from the green for each of the wavebands, then I know exactly what light has been reflected from the red and green. Now suppose that I change the projection settings on the three projectors so that the light that is now coming from the red is exactly what used to be coming from the green. I mean exactly. What color do you think the rectangle that used to look red will now look? What do you answer? Do you understand this? I have red, I have green. I measure the light that is reflected. Now I change the settings on the projectors, so that what looked like a red rectangle before is now reflecting exactly the light that used to be reflected from the green. What color does the rectangle that used to look red now look?

Almost everybody thinks it looks green or looks like a blend or something. No, it looks red as red can be. It still looks red, because how you see a color is not determined by the light that's reflected from the surface. It's vastly more complicated than that. But it will never seem that way to us. And that's the story for almost everything we can think. We are built for small stories. We are built to break the world into events, like moving the flower, and objects, like the flower. Physics does not break the world up in that way. In physics, every object is an event and every event is also an object. We are built to think that we have a personal identity, despite the fact that we change radically over time. Almost everything that we think is — if you look at it just for itself — false. It does not stand up. It is not what you want to have in a scientific theory, but the way you understand things, even scientific theories, is by using the

patterns you have for thinking and you never ever get beyond those patterns. This raises a question of truth and science.

For example, when you are looking at an equation for velocity, you make graphics in two dimensions. Here is time and here is distance, time versus distance. It's great. It's wonderful. It's right there in front of you. And someone may say, "Well, wait a minute, the rock that is moving is three-dimensional, and you have time and distance, that's two dimensional. So this is false. So you cannot use it." Well, of course it's a simplification, of course it's wrong in a sense. But the point is—don't focus on whether or not the blend is true or a particular space is true. The question is instead whether the *network* is useful. You could say "Oh, well the two dimensional graph is a projection of three dimensions on the two". But, notice that's false. In the world, there is never a projection of three dimensions on the two. There are shadows, reflections, mirrors, and so on, but not actually a collapsing of three dimensions into two. Or rather, we have no experience of that. That collapsing is something we can manufacture by blending. And it's extremely useful and it leads us to the truth.

Harold does the same thing. The world that he is moving in is three-dimensional. And the paper that he is drawing on is two-dimensional. And so we might say "Well, there's some kind of inconsistency here. Harold's brain doesn't work right. He is using two dimensions for three dimensions". But that's not the way to look at it. Using the two dimensions is something we do all the time to make a compressed, human-scale, manageable blend of three dimensions. The important thing is the *network*, whether or not there is truth in the *network* or usefulness in the *network*. We compress things. We make stories. We make stories out of almost everything, where by story, I mean narratives of a small number of agents that interact. For example, we do this all the time in science and this is how we are able to have science. We are not able to operate at sub-atomic levels, or relativist intervals of time and space. We don't have any experience of that, but we can conceive of such levels and distances by using basic human capacities for story, compression, and blending, and we're not deluded.

We do this not only for exotic purposes. For instance, we say something like *the acid etches the metal*. We are putting agency in the acid. We know when we think about it that, in fact, the agency is as much as in the metal as it is in the acid, because it is an interaction between the acid and the metal. Thinking of it as interaction is useful to us. Thinking of it as an interaction that is driven by one agent, the acid, is very useful to us and we hold on to this. We are not deluded; the network has true; the compressed blend is useful.

We make such compressions in mathematics. In mathematics, one of the standard things you learn is that a function carries one value into another. That's false. If you have something like a function, a linear function, one way of thinking of that is just as a set of pairs. For every value of X, there is value of Y. But now, there isn't actually an agent; there is no thing that carries. But we like to think of it that way. The function takes this thing from this space to that space, you know. The function "adds one" or "takes two to three" or "turns two into three." Of course, the function does not

turn two into three. Two is two. It stays two and it stays right there and two is never turned to three. But, we understand change. We understand turning the water into ice. There is a beginning state; there is an end state; it is very useful to us to think this way.

Suppose you have a function that is a rational function, meaning it's got a polynomial function in the denominator and a polynomial function in the numerator. There can be a value of the variable X , such that the denominator approaches zero. So, suppose I have some function that's, for example, $3/X-1$. We say, "as X approaches one, the value of the function approaches infinity". Or sometimes, you say it approaches positive infinity if you approach from the right, and it approaches negative infinity if you approach from the left. But of course, there is no approaching. The function does not change. These values are always there. What you are doing is taking the human ability with its limited attention to focus on something and saying "Ah, as I, the human being, focus on different numbers that are closer on the line to one, the value of this function gets bigger". We are taking something that is static and understanding it mathematically through dynamism, the dynamism of human attention. This happens all the time. Fictivity is crucial in science and mathematics. Blending makes them possible, because blending makes it possible for us to construct a network that we otherwise would not understand.

The question to ask is not whether the blend is true, but whether or not the network is useful, whether or not the network stands up.

Words are not a basis for good theories. We are all nominalists. Human beings tend to think that if we have a word for something, then it exists. If we have a word A and a word B , we can say, "What's the difference between A and B ?" as if there must be a difference. Remember words are just words. Words are little compressions. Words are not theories. Words are not true. Words are little prompts to us to think. We have many words that were invented descriptively a long time ago. Many of them are Greek: *category*, *analogy*, *metaphor*, *apostrophe*, *catechresis*, *metonymy*. Some of them are not Greek, like *counterfactual*. These are just words, so do not ask questions like "What's the difference between analogy and categorization?" Hold on. Step back. Maybe you are being misled by your words. And certainly do not say, "Oh, I see, here is a mental ability, here is a specific mental capacity that's called *categorization*. And here is a different one that's called *analogy* and I want to know which one is working, *category*, *analogy*, *metaphor*, *metonymy*, *association*, *abstraction*, and so on". There are just words, for little compressions that come up in consciousness, which is a thin little capacity. Don't assume that what is apparent to you in consciousness is sort of a basic level of science, for science to improve. Because we have a word for something doesn't mean that it exists or that there is a mental process that corresponds to that word. If we assumed that what we see in consciousness provided a good basis for science to improve, we would imagine that colors exist in the world and penetrate our minds. Thinking in such a way has not proven useful for cognitive science.

One of the things about which there is no controversy in linguistics is that human language and human linguistic ability are incomparably more sophisticated than

anybody ever imagined. I studied natural language programming when I was 17. We though were going to teach the computer how to understand English in a summer. This was ridiculously simplistic. It was absolutely moronic but we had been raised on Star Trek, in which Scotty says to the computer “Computer, fire up the engines”, and then the computer does it. Well, Ok, there are computers that do speech recognition now. I use one, in fact, this one. I dictate most of the email I do. It’s wonderful. It doesn’t work anything like a human being works. Human language ability is vastly more complicated than anybody imagined and it doesn’t work even roughly anything like our folk theories tell us it works. So we have to understand that we are in the funny position of trying to use our mental abilities to get past the little tiny folk theories we have about how we work.

The hardest thing about being a cognitive scientist or cognitive linguist is you have to be very perverse. You have to be very contrary. You have to be immediately suspicious of anything that looks obvious, anything that seems to you natural and factual. You have to wonder instantly what’s going on, because the things that seem easiest to human beings like vision and language and social cognition are the hardest things to explain. We thought when we started computer programming that the really hard things to do would be things like automatic theorem proving or chess. Why? Because they seem hard to us. So they must be the hardest things. No, no, no, playing chess is nothing compared to social cognition. Think of being able to stand up without falling down: your folk theory is that you lock your knees and you stand in the place. Completely wrong. You have a constant combination of flexion and extension of all your muscles. It’s all running in the background and if you had to do it consciously, you would fall over. So, the easiest things, the things that the two-year olds can do, are the hardest things.

We find a great variety of integration networks. Don’t focus on a four-space diagram or any other diagram with a fixed number of spaces. In research on blending, we find diagrams that look like all of these different things I am showing you in the slide show. And here I warn professor Li Fuyin that if you print out all my slides, you will be terrified. I’m just going to run through a bunch of these. We find all kinds of different integration networks. All of these are in research articles that are published, but what we are focusing on here is the mental operation of blending, not any particular diagram.

For any conception, you might have blending involved, but you might have very many mental spaces involved. Blending research never gives a complete characterization of the meaning that is going on in a particular example, because there are always many other spaces that have blending in them. People seem to be able to control many spaces in a conversation, focusing on maybe two to four, sometimes five to seven, but hundreds of spaces and blends can come up in the single brief section of discourse.

There is a continuity of principles of blending across many different ranges of thought and action. Money and symbols are conceptions — here in this slide is the blending network for the dollar bill — here in this slide are blends for gauges and

various kinds of controls. Here is a diagram for the way in which Hebrew uses morphological causatives on the verb in order to create a single verb that is causative. Here is a diagram from Scott Liddell on the way in which sign language uses blending to connect elements of the surrounding physical to the topic of discussion. Here is a diagram showing how, in the blend, one part of the signer's body is the cartoon character Garfield looking at the master, and the other part of his body is the master looking at Garfield. Scott Liddell has wonderful analyses of blending in sign language.

Here, in this slide, is an example whose analysis I will skip over. Blending works in jokes. Do you guys eat Toblerone chocolate? In Europe, they have chocolates that are shape like pyramids. This slide shows an advertisement for Toblerone. That arrangement of ancient pyramids in Egypt—that's "Ancient Tobleronism". And that's all you need to be prompted to make the blend. In the blend, perhaps Tobleronism has been a religion for millennia. Perhaps astronauts went back in time and took Toblerone with them to give to the Egyptians. You can construct all kinds of things in the blend, prompted by just these two words—"Ancient Tobleronism?" and the picture of the Egyptian pyramids.

I have an analysis of the way in which elliptical and hyperbolic geometry were invented by certain kinds of blends. If you look at *The Way We Think*, you'll find a discussion of the evolution of the concept of number, different kinds of numbers in the history of number. James Alexander — he is a professor of mathematics and chairman of the department of mathematics in Case Western Reserve University — has just published a paper called "Mathematical Blending" on the way in which blending is indispensable for the development of mathematics. It's in mathematics. It's in jokes. It's in Voodoo. Jesper Sørensen has a book on blending and magic. It's in counterfactuals. We say something like "*Put the tea in front of the missing chair*". The very fact that you can conceive of the missing chair seems simple. It's not simple at all; it depends on blending. We will see some of that. It's involved in various grammatical structures and so on. So I'm now coming to the end of this long part of the lecture, in which I have been trying to stop certain misimpressions from arising.

Blending is basic. It's indispensable for human higher-order cognition. You almost never see it. Everybody has noticed certain kinds of blends, so here is one, the earliest one I've found where somebody not only uses a blend — humans have been using blends for 50,000 years — but also notices that it is blend. It's from Aristotle. But until recently, the view was always that blends are very unusual things that come up very, very rarely. Those just happened to be visible; those were the only ones that anybody noticed. It's not that nobody ever heard of integration before. Literary critics, aestheticians, syntacticians—there is a thing called syntactic blending. Lots and lots of people have noticed little parts of this in their fields.

And here is another thing. It's very easy to think in blending that there is a set procedure. You start with certain inputs, you project to the blends, you get a blend. But you can go the other way! And you can work anywhere in the network. So, for instance, it may be that we have this understanding of knowing in terms of seeing. We say

“Knowing is seeing.” “I see that the square of two is irrational.” It may be, as Chris Johnson has analyzed, that when you start off as a child, you do not really distinguish knowing and seeing. If Mommy says, *“Let’s see if Daddy is in the driveway”*, the knowing and the seeing are not separated out. They are just one thing. But gradually, over time, you start to be able to separate out the knowing part from the perception part. And so, in this case, perhaps the inputs to the blend are developed after the blend, or we’d better say that what you start with is not regarded as a blend. You are not making inputs to the bend, but later on you decompress the blend into inputs. Work can be done anywhere in the network any time. Blends can recruit from more stable structures. Emergence is not only in the blend; it’s often in the network. You get new connections in the network that weren’t there before. So, emergence is not just a property of the blend, and there is a system of blending.

It can seem as if blending is supposed to be a magic key from a mythological story. It just explains everything. It’s everywhere. Well, hold on, for starters, in fact, for human thought, it is everywhere. Analysis of higher-order human cognition in any field that leaves out blending is, I think, going to be mistaken or inadequate.

But in fact double-scope blending is a very small advance. Remember that double-scope blending seems to have arisen only fifty, sixty, seventy, eighty thousand years ago, or something like that. The world didn’t need double-scope integration. All of the species did just fine. And we human beings cannot do double-scope integration without that two billion years of evolution and what it developed prior to double-scope integration. It’s just that double-scope blending is the most recent change. In the great scheme of things, blending is a very small advance.

Blending does indeed happen everywhere in thought, by which I mean, across all conceptual domains. This is not so surprising. Think of categorization: categorize everything. You categorize grammatical relations. You categorize furniture. Does that mean categorization is a theory of everything? That it does everything? That it is a magic key. No, categorization is everywhere, but it is not everything. The same is true of double-scope blending. Double-scope blending is actually a very small addition, but it makes a huge difference in what we can do.

Blending theory is not a reductionist program. It’s not the case that you say, “Oh, now that I know about blending, I can forget about all those other theories about polysemy, syntax, and so on”. No. You need all those other things as well.

There are many disciplines in which blending research is being pursued. Please visit <http://blending.stanford.edu> to see them. Just to remind you: Blending is basic. Doing cognitive science is hard. We have to re-conceive our notion of truth, and science, and conception, and what it means to have truth. Words are not good theories.

I am going to show you a blend that I think probably you have never seen before. It’s a double-scope blend. It’s going to look pyrotechnic—back we go to the exotic stuff. You may regard it as funny. Some of you may regard it as offensive. That’s the point: If you regard it as funny or offensive, that means you understood it. And it also means that it is powerful. So, that’s what I’m going to do here. And as scientists, we do not

just look what we like. If we were botanists, we would not say, “I’m only going to look at flowers. I’m just going to pretend weeds are not in the world”. We have to look at actual human cognition.

Here’s the blend. See this slide? It’s called “The Engineer’s View of the Sexes.” In the top picture, we have a control box with only an on-off switch. That control box is labeled “MEN.” And in the bottom picture, we have a control box with every kind of switch and calibrated gauge imaginable. This control box is labeled “WOMEN.” Let’s look at a diagram for this blend. It has one frame for control boxes, and another frame for the psychology of human beings. These are very different kinds of things. This is a double-scope blend, a new one for you. But you all understand what it is trying to say, that it is trying to say something about the relative complexity of male versus female psychology. It may be simplistic. It may be wrong. It may be offensive. But the point is that nobody has any difficulty understanding it, even though they have never seen it before.

Here is something that one of my children said “*We are lucky we are not chickens. If we were all chickens, you, William, would be about Elizabeth’s age; you, Peyton, would be about Dad’s age; and me, Dad, and Mom would all be dead of old age. We are all five alive. We are lucky we are not chickens.*” My child said this off the top of his head at dinner, when he was young. One input is human beings. The other input is chickens. Why can Peyton be alive if he is Dad’s age in the blend but Dad is dead in the blend? It’s because in the blend, *Peyton would be about Dad’s age* means Peyton would be at the stage of life that a chicken is at when the chicken has lived as long as Peyton, and this stage corresponds to the stage of life Dad is at now as a human being. Notice, we have language, “*Peyton would be about Dad’s age,*” for expressing this. (Kid’s love talking about “dog years.”) That’s not the part I want to focus on. The part I want to focus on is “*lucky.*” Suddenly, in the blend, we are lucky. That’s new structure. That’s emergent structure. Here are a bunch of chickens and they live the way they live. Here we are at the table, eating our spaghetti or rice and we are just the way we are. In neither space are we lucky. Lucky is not even in the conception. But now, in the blend, we are either old or dead. So, suddenly the situation we now inhabit at the dinner table has a counterfactual situation: one in which we are old or dead. The compression of that counterfactual link between those two spaces now makes a present feature for us at the dinner table; that feature is that we are *lucky*. We are now lucky inside that new space, the one where we are not just at the dinner table but now lucky to be at the dinner table as we are. We very much prefer to be in this space than in the blend. This outer-space vital relation of counterfactuality is now compressed to an inner-space feature. We are quite used to human-scale features: features like being yellow. *Lucky* is a compression of a counterfactual link in the network. There are other words that prompt for us to compress network relations. Words like *accident*. We say, “*This thing is an accident*” or “*You are safe.*” Ronald Langacker might say “If you have a verb like *dent*—you have *dented* something—it’s only a *dent* because in the network, compared to something else, there is a disanalogy. *Happiness, mistaken,* these kinds of words, which just seem perfectly normal, *lucky*—they all seem perfectly

obvious to us. Being lucky seems as normal, as human scale, as being five foot ten inches tall. It's like having brown hair or being mail. *Lucky* is now a feature. Network structure gets compressed into something that is in the blend.

From now to the end of this talk is about five minutes. In that time, I am going to present to you another pyrotechnic example. In that time, I'm going to go through the various features of blending and introduce some questions about the language. But don't be confused by the fact that this is a pyrotechnic exceptional example. I am using it so you can see. Remember that the features you will see in this example are basic in human cognition.

This is a passage from Racine's *Phèdre*, which was brought to my attention by Gilles Fauconnier. There was a king, Theseus of Athens, and he went down to Crète and he defeated the Minotaur. Do you know the story? I have to tell you carefully. He goes down to Crète and he defeats the Minotaur. The minotaur is a monster. Phèdre is the younger sister of Ariadne, and Ariadne falls in love with Theseus. And she helps him out with the Minotaur by giving him a thread. So when he goes into the maze to fight the Minotaur, he can find his way out again by following the thread. Theseus escaped from Crete, and took Ariadne with him, but he abandoned Ariadne on the island of Naxos, or so one version of the story goes. She became a princess of Dionysus. Theseus and Antiope, an Amazon, have a son, Hippolytus. Later, Theseus marries Ariadne's younger sister, Phèdre. Hippolytus is a beautiful young man, and he has grown up. Later, Theseus is away adventuring, and Phèdre is alone with Hippolytus. Phèdre has developed an absolutely consuming passion for Hippolytus. And she fights against it and she fights against it and she fights against it but you can not fight against the gods. One day, she is hanging out in palace, and Hippolytus notices that she is doing well enough, and he asks her, "Don't you miss Theseus?"—that's his father and Phèdre's husband.

Here you see the text from Racine's play, in which Phèdre talking to her stepson Hippolytus, and she says, "Yes, I miss him. I want him, I love him, not the way he is now, profligate, unfaithful, chasing all the girls, but the way he used to be, the way you look". This is an analogy and not a hard one to follow, because Hippolytus is the son of Theseus. Of course there's an analogy. Hippolytus is a prince. You understand this analogy between the Theseus of long ago and his son Hippolytus now. In one mental space, we have Theseus and Phèdre. In the other one, we Phèdre and Hippolytus, two separate spaces. And Phèdre is making an analogical connection between Theseus and Hippolytus. "He had your courage, your eyes, your way of speaking." "Why?" she asks Hippolytus, "Why couldn't you have come down to Crète and saved us from the Minotaur?" Of course, the answer is he wasn't alive at that point, but so what? You can imagine it. And now, we have a blend. There is one space in which Theseus saved them from minotaur, and it is being blended with the space that has Hippolytus. In the imaginative blend, it's Hippolytus who saves the Cretans—Minos, Ariadne, Phèdre—from the Minotaur. And Phèdre says, "You could have been the one that my sister saved," and then she says "No, actually I would have been the one". So suddenly, the analogy between herself and her sister is compressed, and now in the blend, she is the one who

is saving Hippolytus from the Minotaur. “I could have given you the thread,” she says. And then she says, “No, no, a thread would not have been enough. I would have gone into the Labyrinth with you.” And she says, “I would have gone in with you. My love would have made me think of this idea”. Now, because of the blend, she can talk about the love that Phèdre, back in Crete, being saved from the Minotaur, feels for Hippolytus. She can talk about her love for him.

This can still seem within bounds, in a way: she loves the hero who defeated the Minotaur, and in the historical space, that person is Theseus, her husband. But in the blend, it’s not Theseus she loves, it’s Hippolytus. Then she says, “I would have gone with you. I would have saved you. I would have walked before you. Phèdre would have gone in the Labyrinth with you, to be with you and to be saved or to be lost.” At this point, Hippolytus cannot miss what’s going on in the blend. This sounds to him like an avowal of love. He says, “Madam, have you forgotten that Theseus is my father?” And she backs up and says “What do you mean? Why do you challenge me?” She is trying to hide behind the analogy, but the blend is very powerful

There is extraordinary emergent structure in this blend. There are two separate mental spaces. There are connections across them. We have projections to the blend and emergent structure. In the blend, Hippolytus is in Crete, and Phèdre saves him. Presumably, they get married there. She is now able to say things like “*your lover*” to refer to herself, because she is his lover in the blend. She can utter a phrase like “*my love*,” referring to her feeling for Hippolytus. She can say of the blend things that she refrains from saying about the actual situation she inhabits.

When Phèdre objects that she has said nothing wrong, Hippolytus says, “Oh, I am sorry. You are right.” He is willing to pretend it is just an analogy, and he backs off. This is the famous avowal scene in Racine’s *Phèdre*. This is the scene in which she avows her love for Hippolytus. She replies to him, “Ah, cruel one, you understood me only too well.”

This blend looks pyrotechnic. There is a mapping between elements in the two stories. Hippolytus, Theseus, Phèdre, Ariadne. I won’t go through it. There are lots of mappings, identity, analogy, similarity, causality, change, time, intentionality, space, role, part-whole, representation. There is selective projection. Only some things come down from the input spaces into the blend. There is emergent structure, really spectacular emergent structure. Now, Phèdre and Hippolytus are lovers in the blend, whereas they are certainly not in any other kind of place.

When we talk about extremely simple expressions like *Paul is the father of Sally*, when we talk about the evolution of numbers, when we talk about Racine’s *Phèdre*, when we talk, as I will in this afternoon, about certain basic grammatical clausal constructions, it may feel as if we are talking about different kinds of things. But don’t assume that the mental operations underlying them are different. Of course there are differences between all these examples, and we must explain why they feel different to us, but underlying all of them is a basic mental operation that makes human higher-order cognition possible and indeed makes language possible and that will be part of the subject for this afternoon. That mental operation is double-scope blending.

Thank you very much!

Lecture Four Mental Packing and Unpacking

All of the speakers in this series travel here from somewhere else, so I'd like to talk to you about traveling. I am going to use traveling to present a metaphor for thinking about cognition.

You know when we travel, we take a little compressed suitcase of stuff. The little compressed suitcase of stuff is designed to fit in the world. That is, it goes in the baggage-handling machine in airports I've never been in. By the way, the baggage-handling machine in Beijing airport is the best I've ever seen. I think Beijing might be the biggest airport I've ever seen and that's really saying something. The little compressed rolling bag goes in the overhead compartment of the plane, including planes I've never seen before, and it rolls through airports I've never been in. It goes down runways and walkways that I've never seen before. It is built to fit into the world, wherever it goes.

And inside the rolling bag I have some stuff. So, for example, I have this little electrical supplier for my computer. It's a power brick. It goes in the bag, and it plugs into the world. It plugs into parts of the world that I've never been in. It plugs into places that have current of 110 volts. It plugs into places that have current of 220 volts and places where the current is 50 hertz and places where the current is 60 hertz. I leave my house with this little power brick knowing that somehow I would be able to make it plug into the world. And if it doesn't really plug into the world, I have another little thing that I carry in my bag. It's a universal plug adaptor. Sometimes I have to pump out this plug to make it plug into the world. And sometimes I have to pump out this plug to make it plug into the world. The plug from my power brick plugs into one end of the adaptor, and a different plug pumps out the other side, depending on where I am in the world. Some of the plugs needed in the world are very odd. British plugs, for example, I find to be extremely goofy. The British plug is very large—here, something like this. You also have this kind of outlet in Hong Kong.

So I put together this thing—a power brick—with that thing—an adaptor—and with the outlets that are in the wall, you see, and it's amazing: it works! See, I do this, and I plug this in, or I need the other one, see, Ok, this one, yeah, good. Where I go in the world, there are little microphones that plug into me. I bring my clothes, my voice, and we hook up the microphone, and it works.

And I take my tooth brush. I don't take water. I think there is going to be water there where I go. And I take my clothes but I don't take hangers.

When I travel, I can't take my whole house with me. I don't say "Oh, I need some place to live; I think I will take my whole house and its electrical and water supply. And I'd better take my car too. And I'd better take all of my office with my books." I don't have to do that. I take this small compressed packed bunch of stuff in my rolling bag.

And some of the stuff I take is in this little tiny thing, a wallet, little tiny thing.

And just the way the bag fits the airplane, this wallet fits in my pocket. And I have a pocket. And I can put all kinds of things in my pocket. I take them out of my pocket. When I leave my house, I have my pockets and I have my wallet and I have my passport. Amazing thing, a passport. Everybody all over the world looks at it and says: “Oh, Ok”. And I have my little rolling bag. And when I get somewhere, I unpack it. I unpack the cord, and I take out the money, and I say, “Hi, I have this money. Do you take this money? No, this is no good. Do you take this other money?” “Yes, we take that money. Ok, fine.” And if I don’t have any money, then I say “Oh, well, do you take this little piece of plastic?” They say “Ok, fine, the little piece of plastic.” And people say “Who are you?” and I say “I’m this person”. Ok, all these stuff, right here, in this little wallet. I unpack it when I get there. And I can work with what I find. And people can work with me.

When my trip is over, I pack my bag back up to go home. My repacked bag is pretty much the way I brought it, but not exactly. It changes a little, because when I unpack my stuff and use it and then I repack, maybe I find a slightly better way to pack. Or maybe where I went, I found something that was just a little better than the thing that I brought, so I replace it. Or maybe I use something where I go and I acquire it and take it back in my bag to go to the next place. So when I leave, I am taking a slightly different packed set of stuff.

I would like to suggest that this “packing and unpacking” metaphor is a better metaphor for thinking about how we work with the world, and how thought works, and how meaning is developed, than the metaphor in which we have all the stuff in our head and we just pull it out and use it.

Everything that I know, everything that you know, including all our linguistic competence, is not all preformed here in the brain. No, it develops when we unpack our constructions. We unpack our templates for thinking. We build a network, a conceptual network. We build expressions. We take in expressions when people say them. We use the constructions we have to try to make sense of their expressions. We build a network by unpacking the constructions that we have and using them, and we pack the network back up. When we pack them back up, we can take away with us little compressed mental frames, little compressed linguistic constructions. We then unpack them again, next time.

The next place we go, the money will work or we can convert it. The tooth brush will find some water. The electricity will plug into something. If it doesn’t, well, we have this credit card and we can buy some new stuff when we get there.

Packing and unpacking—this is, I think, what we are doing all the time. WE need a model of linguistic knowledge that follows this metaphor of packing and unpacking. Why do we need it? We have heard today a wonderful presentation on diversity and change of linguistic items, and these items are what we call open-class items. So for example, although grammatical categories are closed-class—noun, verb, adjective, things like that—the individual items that go into these categories are open-

class, which means you add them very easily to the category. For example, the word *fax* means ‘send something to somebody through the telephone line’. It comes from the word *facsimile*. “*Fax*” did not exist until I was at college. But that kind of change—adding a word to a grammatical category that is open, is easy. You just make a new verb. You make a new noun. You make a new adjective. If you discover a new planet, you make a new noun. The linguistic fact that this noun can go into only certain kinds of positions, that’s closed class. But the fact that you can make a new one that we can agree on easily, that’s open class. Everybody understands why open-class aspect of language would change—because there is new meaning and you want a new word. Fine.

But how about closed-class? Modal verbs in English, for example, are a closed class. We are talking about verbs like *will*, *shall*, *could*, *can*, *would*, *might*, *may*, *must*. The modal verb system in English is different now than it was in the 14th century. Why? Why would it change? If you are always just pulling out the knowledge you have and using it, why would what you pull out ever change? Why would the modal system of English have changed over the time? That needs an explanation, because languages change. They change gradually in their closed-class systems. They change gradually in their pronunciations. Why would that have happened? Why does it continue to happen? We must explain this.

If we think that our linguistic knowledge or our linguistic performance comes from pulling out static knowledge that we have and using it and letting it go, then there is no reason that closed-class items would change. It has to be that something about the way human beings use language results in the kind of changes that we see in closed-class items. It is true that we see variation and diversity. But it is not random. Things do not just change willy-nilly. They change according to various kinds of regularities. We are always trying to explain this feature of human cognition: How do we get diversity over regularity? Human performance always has these features: we see diversity and change, but with strong regularity underlying that diversity and change. I think the reason that closed-class dimensions of language change is that we are always unpacking and re-packing them, and in this constant packing and unpacking, we can recompress the closed-class dimensions of language slightly differently.

Mental packing and unpacking are the subject for today. What we are trying to explain here is language and why human beings have the kind of language that they have. In order to operate in the world, we must have knowledge that is portable. Our knowledge has to be portable. Our linguistic ability has to be portable. We must carry it with us. It has to be able to occupy the available bandwidth. What that means is it has to fit in to the ecological affordances of the situation in which we find ourselves. My suggestion is that language is portable and fits into the available affordances because language is an ability to unpack and repack basic patterns. In unpacking, we construct integration networks. We then repack from those integration networks to compressed patterns that we carry with us. Metaphorically, language is not so much a big estate with everything in it as it is a compressed, packed little rolling bag of stuff that we carry

around, a very portable and light little rolling bag of stuff, which we unpack to plug into the situations in which we find ourselves, and then we repack. And this is the basis of change in the guts of language.

I am contrasting two ways of looking at language. One way is to think of language ability and knowledge as “retrieval” and “use”. In that way of looking at things, you have a tool shed or you have toolbox. When you have a need, you pull out the tool and you use it. But when you use the tool, it is still in the shed. You can use it and let it go, get it from the tool shed again and let it go, get it again and again, because it is always in the tool shed, no matter how many times you take it out. You just go and get these things that you have. They’re all fully available, all the time. You pull them out and use them. The other way of looking at language is as a matter of unpacking and repacking basic patterns.

Here is an illustration of unpacking. Here, in the slide, is a sundial. If the environment you are in doesn’t have the shadow in the sundial, then the sundial is useless as a clock. It might be pretty, or memorable, and so on, but it’s useless as a sundial.

Where does the shadow come from? It’s not the case that the sundial has it. The sundial does not carry the shadow and put it in the right place. And it’s not the case that when I look at a sundial, I pull a shadow out of my head and locate it in the right place. The shadow comes from the environment.

I know how to read the sundial. I carry that part with me, mentally. And I run across this object in the world, this sundial, in its place, and there is also the sun, and physical system that produces what looks to me like a mark. The system of the sun and the sundial produces a mark. Was the shadow there without the sundial? No. Was it contained in the sundial? No. The chronometer emerges when you put together the environmental system with what I know about sundials. I unpack what I know about sundials and plug it into the world I am perceiving, and as a result, there is meaning. You can think of your knowledge of a sundial as one of the little things that you carry around in your rolling bag, your suitcase. What you are carrying around in your mental suitcase is your knowledge of how to read the sundial. In fact, what you are carrying around is probably something more like the ability to infer how to read a sundial. You know that the shadow, the mark, is a variable part of the frame. There are static parts and variable parts. There are regularities. You unpack what you know to fit the world, and you end up knowing what time it is. The sundial doesn’t really tell you the time. Knowing the time is a result of your unpacking your knowledge so as to fit the world. But sundials are highly variable, and when you repack your knowledge of the sundial, perhaps you leave with something a little different. For example, I once saw a sundial with no gnomon—that is, no mechanism for casting the shadow, and that was because the observer was supposed to stand in the middle of the sundial and be the gnomon! The observer’s shadow told the time. That’s neat. Instantly understandable, but creative. It was not previously in my frame for sundial that the observer and the gnomon could be the same thing. But when I repacked my frame for sundial, now I had a slightly different structure.

Languages change. The way I'm going to be talking for the next six lectures about language is to take the view that knowledge of language consists of a rolling bag of small, packed constructions that you can activate in a situation in such a way as to construct a blending network suitable for the situation. But once you have constructed the network, you do not keep it. It vanishes. You let go of all these networks, but the construction of the network can change the elements that went into the network. You can make a network that can change the knowledge that you have, just as repacking your rolling travel bag can change it a little as you go through the world.

I came here to China from Portugal. When I land here in Beijing, coming from Portugal, I unpack what I have as I need it, and then I fold it back up. My next destination is Ohio. After that, I am going to go to Berkeley to teach in the LSA Summer Linguistic Institute. I do not take with me everything that I generate in any one of those places. What I do take is a set of constructions that I can put together to make a new network in the next place I go.

We are talking about the way we pack and unpack things. We use the principles of conceptual blending to pack and unpack mental arrays. We pack up what we do into something that we can carry, so we can deal with the next situation. The utility of what we carry in our rolling bag is that they can assist us in constructing a conceptual integration network.

Think of an expression like *In a leap year, we add a day to February*. It seems perfectly normal, uncreative. Does everybody understand that phrase? Every four years, we have to add one day to the year. Now think about that for a minute. What is the year? What is February? In your experience, there is just one day and then another day and then another day and then another day and then another day and then another day and it never ends. There is no *February* out there—that's one thing. And there is no *year* that's out there—that's another thing. Of course, there is one day and another day and another day and another day and when you add them all up, you think they make a year.

There are analogies and disanalogies across all the days in our experience. We pack those mental spaces for those days, and those analogies and disanalogies across them, down into a thing, one thing. We transform many into one, by blending. In the blend, there is one packed thing, *the year*. And then it repeats.

I ask you to focus for a minute on the fact that years do not repeat. Days do not repeat. Life is not like the movie *Groundhog Day*, in which one particular day really does repeat. It's never happened that on January the first, I'm awakened and it's the same as the previous year. I will never be the same age as I was a year ago. My children are never the same age as they were a year ago. Every year is different.

It's not true that the year repeats, and we know this. But we pack all those different years, with the connections of analogy and disanalogy between them, into one year, one little package. This concept of *the year* which *repeats* is like my electric converter or my power brick. It's a really useful thing to have in my rolling travel bag. I can unpack it and use it to understand new stuff, new years. I understand this year and

I understand the years before and the years after, because I have this concept of a year. It is something to unpack and plug into the world.

Now that I have this one thing—*the year*—I have what I need for a human-scale scene in which I take something and add a little something to it. I have my coffee, and I add a little milk to it. I have the year, and I add a day to it. Adding a day to the year is an extremely creative conception—impossibly beyond the mental capacities of other species. We can say to ourselves and each other, “Ok, I’ve got this one year and I am going to add something to it”.

Now of course, in the day-to-day view of time, I’m not adding anything to anything. In the day-to-day view, there is one day, then another, then another, then another, and so on. But I can compress the analogies and disanalogies across all of these days. The analogies compress to make a cyclic year, made up of 365 days. The disanalogies compress to a change for the year: this year, this leap year, I am *adding* a day to the year. Of course, I am doing no such thing, but in the blend, that is what I am doing.

This is a very standard form of packing and unpacking. Look at this slide. Out here, across all of these mental spaces, we have analogies and disanalogies. These are vital relations connecting spaces. One standard type of packing is to pack outer-space vital relations that connect spaces into structure that is inside the blend. One very common form of packing is to pack the analogies across spaces into identity in the blend—so analogies across days and seasons are compressed into the cyclic year—and to compress the disanalogies into change for that identity—so, for example, in this leap year we change the year by adding a day.

To repeat: there are analogies across time. That season was like this season. That’s an analogy. Those analogies get compacted into identities: *Fall*, *Spring*, and so on. Packing those analogies gives us the cyclic year.

And there are disanalogies. This year has 365 days. That one has 366 days. This disanalogy is compacted to change for the identity: we *changed* the year by *adding* a day. Of course, we didn’t change this year. We didn’t add anything here. But I compress the analogies to an identity—*the year*. And I compress the disanalogies to a change for that identity—*adding a day*. The fact that this year is different from that year now becomes, in the blend, a change in *the year*. But in fact nothing changed in the previous year.

Down here, in the blend, I have an identity and changes. No one is fooled. We don’t now think that the year before the leap year has an extra day, or that we add a day and now we have changed the year. No, the identity down here has a change, but in the network, we can expand the notion of the cyclic year to understand there are all these analogies and disanalogies in the network. We pack our experience into a blend, and we unpack the blend to plug into our experience.

We say something like *kick the habit*. Notice what it means to have a habit. I have a habit of drinking coffee. What it means is that after lunch I go get coffee one

day, and the next day I go get coffee, and the next day I go get coffee, and the next day I go get coffee, and the next day I go get coffee, and I do it again, again, again, and again.

But where is the habit? What is the habit? The habit is a compression over events. It's a packing of analogies across all those days.

Now, when I pack analogies over all those days into the blend, I have something new, I have a *habit*. Now that I have this one packed thing in the blend, I can do things with it. I can kick it, for example.

This is really very interesting. How is that if one day I don't have coffee, I can say things like *I'm kicking my habit* or *I'm letting go of my habit* or *I'm banishing my habit*. Down here in the packed blend, I have one thing, and I can change it. I can do things to it. The disanalogies across these spaces compress to change. They are packed to change. The analogies are packed to identity—my habit.

Your agent is waiting for you. This was the phrase I saw in an advertisement for airlines, a worldwide airline. In the ad, there were pictures of various gate agents and ticket agents, one from each country: China, Thailand, Germany, France, the Philippines, Mexico, Sweden—one agent for each country. *Your agent is waiting for you.*

What is this *your agent*? It's a compression. All the places you go, there is somebody there willing to help you. He doesn't know you—you haven't even met this person. But now in the blend, it's *your agent*. There is a role in the blend, your agent. For you, all those possible agents are compressed into one per country—your agent. And for the agent, it's a compression that there is a *you*, because up here, in all the mental spaces with all the possible real agents and all the possible real customers, there are many, many agents and even more customers. Many different agents are waiting for many different customers in many different places. But down here in the blend, we have a person waiting for another person: your agent is the person waiting, and the person for whom they are waiting is you.

In what sense is this *my agent*? This agent doesn't know me. But when I interact with an agent, the agent knows me. So now all of my interactions with other agents are compacted down to a possessive for this role. This is now my agent. So I can say, "Well, my agent in Russia will handle this" *even though I have not met the agent*. I know that there will be one when I go to the counter.

In a case like this, I am unpacking something—not my electrical adaptor, but instead my frame for a customer at the counter with an agent. When I am at that counter, I can say, "I'm looking for my agent", because I expect it to work, I expect to be able to unpack my portable frame and apply it to the local situation. The customers are always right.

His girlfriend gets younger every year. Now that could in principle mean that she is aging backwards. She used to be 30 and now she is 29. It would be false in our world, because people don't do that. But the statement is intelligible, and you can imagine its being true of a science fiction novel. Perhaps, for example, his girlfriend gets younger every year because she has a special disease that causes here to age in the opposite direction. She could make a lot of money by patenting that disease, I'm sure.

But, except for the science fiction novel, that's not the way we interpret this phrase. We mean there are different girlfriends, and there are disanalogies across them. Of course, there are analogies: it's a girlfriend, his girlfriend, and the relationship began recently. And there are disanalogies, such as in her age. The analogies are compressed to identity—his girlfriend—but the disanalogies are compressed to change for that identity: she gets younger. *His girlfriend gets younger every year* makes perfect sense.

Nobody thinks that human beings get younger, actually. It's false if you are focusing on the blend, but it's true if you focus on what the phrase means for *the network*. We are not deluded about the blend. In the blend, there is *getting younger*. *Getting younger* is not in any of the input spaces for any of those individuals. But we pack down our knowledge of this man and his relationships. Now we have a packed scene, in which his girlfriend gets younger. We can expand this packed conception and plug it into the environment the next time we meet him with a girlfriend.

You want to get your tennis serve back. I know Sculpting in Time cafe is next to Deep Tennis, which I love. Everybody is playing tennis. How do we interpret “You want to get your tennis serve back”? Suppose that there is something you can do this day, and this day, and this day, and this day, and you start not to be able to do it and pretty soon you are not doing it. Where is the thing? This is just a set of analogies and disanalogies, across your actions on all of those days. But we pack those analogies and disanalogies into a compressed blend. There is no “thing” in the various mental spaces, but in the blend, you have a *thing*: your tennis serve. Your ability to do this now becomes one identical thing. Change comes from compacting the disanalogies: if you could do it, you could do it, you could do it, you could do it, and then you start not to be able to do it, and then later you start to be able to do it. These are just differences. They are not changes for a thing. They are differences. Every day there is a different action. It's not a changed thing. But it's so natural for us double-scope blenders to compress this down into some one thing for which there is a change. Now the change is that you “get back” the thing as if it were an object.

Consider *Your French has disappeared*, for example. If you pack analogies and disanalogies across all the spaces in which you want to speak French, then you have in the blend something, your French, and when you lose that thing, it disappears. It makes a nice little compressed scene. We can use these congenial packed conceptions to grasp networks that are much bigger and more complex than we could otherwise store and manage. We construct these networks on the fly. We make sense of things by unpacking what we have and plugging it into the environment to make a conceptual network, and then we repack our conception, and we move on. I propose that we think of this not as retrieval and use but as packing and unpacking..

Gilles Fauconnier came up with the following example. He received a bill in an envelope. The bill was for his use of electricity in his house. And the back of the envelope said *make this envelope disappear*. It was a request that he sign up for electronic reception of his bill instead of receiving a paper bill in the mail. So what the electrical company actually wants is a history in which he received a paper bill this

month, and this month, and this month, and this month, and this month, and this month, and this month, and then not this month, not this month, not this month Across all these months, there are analogies, that is, he receives a bill every month. But then there are disanalogies: the bill is paper this month, but not that month. This array of analogies and disanalogies across all these events gets compressed: in the blend, there is one thing—the envelope—and it changes: it disappears.

Think of it: it's impossible to make this envelope disappear. It is here. The envelope that is printed with the phrase "Make this envelope disappear" is not going to disappear. It's not going anywhere. The recipient of this envelope is not supposed to say "abracadabra" and make this envelope disappear. But in the blend, the analogies have been compressed to one envelope, and the disanalogies are compressed to change for that envelope—it disappears. The disanalogy that previously the bill was paper and in the future it will be electronic is compressed to a change: the envelope disappears. No one is deluded. No one thinks that the envelope is actually going to disappear. But this packing is a way we can grasp this network and hold down to it. We pack the network to something at human scale and we carry these things around and unpack them.

Let's talk about packing and unpacking for language. You assemble many mental spaces, and you have many experiences. People say things to you. Often what they say is like something else you have heard. You are in a different room, you are in a different airport, and you hear something like something you have heard before, and there are analogies across all these utterances. All those utterances constitute a lot of information, a lot of stuff. There are also disanalogies across these expressions.

You can compress the analogies into identity, and the disanalogies into change. Further, you can compress all the identities one more step into uniqueness, one unique thing, inside one mental space—like the envelope, or the girlfriend. This is an extremely common pattern of packing—analogy to identity or uniqueness, and disanalogies to change—and we will continue to consider it for a while. It comes up again and again and again and again in human cognition.

It's not that you lose the analogies and disanalogies. It's not that you lose the network. You can always unpack the change for the unique thing into a much larger, more diffuse network that is grounded in the human-scale blend.

The analogical vital relations can be compressed down into a unique thing conceptually, and the differences can be compressed down to change.

We do just this in language. When we hear different people speak, there are strong analogies across the ways they speak, but there are also disanalogies across the ways they speak. We can say "Ahha, they're both using the same kind of construction, but one of them pronounces it or deploys it just a little differently than the other".

When you encounter a speaker of a language, you try to make sense of what they say. You have this little packet and you expand it to take in what they are doing. But when you make a network, you can repack it. So everybody you hear has the chance to change just a little your linguistic knowledge.

Theory of language must take on board the actual facts of language, so we need to include in our ideas of how language works mechanisms that produce these gradual

changes in closed-class items. We know that languages do change in this way over time.

Eve Sweetser has looked at such change in an article on change predicates. She investigates sentences like “*The cars get three feet bigger when you enter Pacific Heights*”. Pacific Heights is a rich neighborhood. “*The cars get three feet bigger when you enter Pacific Heights*”, is interpreted as meaning that the cars that you find parked in Pacific Heights are larger than the cars in the other neighborhood.

Of course, no individual car got bigger. All the cars stayed exactly the same size as they were. But there are analogies and disanalogies across all the cars you see. The analogies are compressed to *the cars*. The disanalogies are compressed to a change for *the cars*, namely, they *got bigger* at a certain point. First you are in one neighborhood, and you see cars. Then you enter Pacific Heights, and you see cars. In the blend, *the cars got three feet bigger*. You have language for such a blend, the language of change for an element. You can use this language of the blend, and since the blend anchors the network in various ways, we can figure out what is meant for the network by such an assertion for the blend.

Another example from Sweetser is “*the fences get taller as you move westward across the United States*”. This is the normal way to say it. Again, of course, no fence got taller. As you go down the road, the individual fences do not grow upward. Note that the sentence would be grammatical if that is what you intended to mean, of, say, a science fiction world, where fences are living and grow, and grow taller for some reason as you drive westward.

Instead, we interpret this sentence as saying something about the blend, the one in which the fences are compressed, and change. We unpack this conceptual structure for the blend by creating analogies and disanalogies across the network. Language exists for talking about the blend, and we know how to connect the blend up to the network, so talking about the blend conveys meaning indirectly about the network. This is a great virtue of conceptual integration networks.

We pack conceptual integration networks to human scale through blending. Today, we are talking about mental packing and unpacking as a way of thinking, and as a way to have language. You carry your little pack of stuff with you, your portable stuff, as you move through life, and unpack it to plug into experience, then repack it, and so on. It’s not the case that everything you could know about language has already been set in your head. Instead, you have been provided with the ability to construct conceptual integration networks. And your culture has provided you with constructions that work as prompts for packing and unpacking.

Let’s give another example of packing and unpacking, and of how language prompts for packing and unpacking. There is a place called Lake Tahoe in California, and they have wild bears there. The problem is that Lake Tahoe is a resort area. The local people living in Lake Tahoe know that you should leave the wild life to be wild. But many people from San Francisco and surrounding areas go up to Lake Tahoe and many of them do not understand what a wild bear is or how to treat it. So they try to do what they think of as very nice things to feed the bears. Or they do other things

without meaning to feed the bears but end up feeding the bears. That is, they have packed knowledge—their little mental rolling bag of mental stuff—and they take to Lake Tahoe and they try to unpack it to fit into Lake Tahoe. One of the things they have in their packed mental stuff is the idea of feeding animals, especially cute animals, especially furry mammals. They have dogs and cats and they go to the zoo. At the zoo, they buy food to feed the animals. You are permitted to feed the zoo animals only certain kinds of food, which you buy at the zoo. Ok, fine. And they go to Lake Tahoe. Nobody is actually crazy enough to feed a grown bear. I don't know whether you have seen a grown bear. There are certainly plenty of them in China, and they are scary. But what you can do is throw food to the cubs. You can also leave food outside when there are no bears and then go inside and maybe the bears will come later and eat the food. There are other things you can do that have the effect of feeding the bears. You can throw food into the garbage can that you mean to be taken away, or you can leave food in your car with the windows down, or you can leave food on a picnic table. It turns out that the bears are very good at getting into the garbage cans and the cars and the picnic baskets. In these cases, you didn't mean to feed the bears, but in fact you do.

Or perhaps you own a vacation home in Lake Tahoe, and you plant fruit trees. Why not? You unpack what you know into Lake Tahoe, and it lets you know that your vacation home is your fun area, so you do what you want, within normal limits. You plant six fruit trees. But then the bears come to eat the fruit.

If the bears get used to eating food you have provided—intentionally or unintentionally—near human beings, then the bears become accustomed to being near human beings, and their offspring become accustomed to it, and pretty soon you have bears prowling around human beings, and some of those human beings may call the authorities to come shoot the bears.

All over Lake Tahoe, there are signs like this one you see in the slide show. They are posted in bookstores, on walls, and so on. Explaining to the tourist populations not to feed the bears does not seem to work very well. Somehow, the tourists don't repack their networks so as to include that knowledge. But this sign seems to be working better. It says, "*A fed bear is a dead bear*" In the accompanying picture, there is only one bear, not all the bears. There is one bear, and it is being fed by one agent. The one bear is a compression of all the bears, and the one agent is a compression of all the agents who feed. No one would actually feed a bear by hand, but that is the picture we see here.

Do you recognize this agent who is feeding the bear? Are you familiar with this agent? It is a cultural symbol. I see, you do know what it means. Good, because it's a conventional compression. It's something we call the Grim Reaper. The Grim Reaper is a mythical being who comes along and takes you away when you die, in fact makes you die. Reaping is what we do to plants like wheat when we harvest them: we cut them down. The instrument that The Grim Reaper carries, here in his left hand, is a scythe, a blade for cutting down plants.

The Grim Reaper is an amazing compression. It requires many steps. To begin, we need to create a general cause for death. There are many different individual

deaths, which happen in many ways: old age, accident, infection, cancer, murder, and so on. There is no one cause for all these deaths. We compress all these events of dying into a general event—death. We compress all those different causes into one cause—Death. Death causes death. This is the result of a general pattern of compression. What causes hunger? Hunger. Hunger makes you hungry. Hunger causes hunger. What causes lust? Lust. Lust causes lust. Sleep causes sleep. And so on. When you think about it, it's a little crazy to have these general-cause compressions, but it is crazy in an extremely inventive and useful way. We pack networks down to our portable mental stuff, and in that portable mental stuff, there is a concept of Death as a general cause. Death-in-general comes to everyone. We can unpack that concept as we go through the world to plug it into experience and construct conceptual networks for understanding individual deaths. The creation of a concept of Death as a general cause is the result of a general pattern of packing that produces causal tautologies: Fear causes you to be afraid; Desire causes you to desire. Many different things can cause an act of anger, but we have a packed notion that there is a general cause: Anger causes you to be angry.

Now let us consider the reaper. Reapers cut down plants. But notice the reapers don't kill the plants, in general. The plants can grow back up. And there are many reapers in a field, usually, not just one. And they are not grim. Instead, they sing, maybe they smile.

Actual people who are reaping out in the field, cutting down the harvest, are not wearing heavy robes with cowls. So the projection from the input mental space of reaping to the packed blend of The Grim Reaper is very selective.

There is another mental space that is an input to The Grim Reaper: a killer. Killers can be grim. The killer knows that the victim is there and focuses on an individual, typically.

So, we have three input spaces: Death-in-General, the reaper, and the killer. The killer knows that you are there. So you have three things go in one here, the killer, the reaper. Blending creates a personified Death-in-General in the blend, who is a reaper and a killer, who is grim, recognizes his intended victim, and is, like the harvest, an inevitable part of the cycle of life.

Let's review what goes into this blending network. First, we have a compression to make a general cause of death, Death-in-General. Death-in-General must be blended with a person to become a personified death. That can happen in various ways, but in this network, it is done by using as inputs the reaper and the killer. Down in the blend, we have one agent, as indicated by the definite article, "The." The uniqueness of The Grim Reaper comes from the mental space of Death-in-General, because there is only one such general cause. The grimness of the Grim Reaper is projected from the input space with the killer. And the reaping is projected from the input space with the reaper.

Now, what about the bear? In this slide, showing the poster, The Grim Reaper is feeding something to the bear by hand. In the unpacked networks, there are lots of bears eating, and there are analogies and disanalogies across these bears. In the blend, there is one bear. It might be that for the real bears, it is the mother who becomes

accustomed to human beings but the cub that gets shot. But in the blend, the bear that is fed is the bear that is killed.

In the blend, the complicated causal links that run from your planting fruit trees or leaving picnic food out all the way to the shooting of the bear is packed down to a direct causal link: feeding causes death. A very long chain of cause and effect across all the bears gets packed down to a direct cause in the blend. All the bears are one bear. All these different actions that result in the provision of food to bears are packed down to one action. A fed bear is a dead bear. People who intentionally or unintentionally provide food to the bears are The Grim Reaper. The inference is that if you provide food in any way, you are Death, you are grim, you are a killer of the bear.

There is another possible input to the blend of The Grim Reaper feeding the bear. It is the frame of poisoning an animal by feeding it something laced with poison. Poisoning is not mentioned in the language in the ad, but the input might be activated for some people. If it is, then the person who provides the bear with food is The Grim Reaper who poisons the animal.

The result is a packed conception of food and bears that you are meant to carry around Lake Tahoe. The packed conception is intended to guide your actions, because you pull it out to plug into the world, and when you do, it should cause you to act so as not to provide food. The ad is an attempt to change your conception of who you are and what the bear is, so that when you go out into the world and unpack your conception, you will behave differently. You won't put your trash in the garbage unless it's a special locking garbage can designed to defeat bears. You won't leave you food in your car without at least thinking about it. You won't plant in the orchard.

When you look at this compression, you see in the ad, running down the arm of The Grim Reaper, the arm at the end of which is the food offered to the bear, you see some language. Actually, it's not the arm, but the clothing—clothing associated with funerals and religious rituals—that has writing on it. The Grim Reaper is a skeleton in a monk's habit. Why a skeleton? Well, up here in the uncompressed network, somebody dies, they get buried, gradually they decay, and many steps down the line, there is a skeleton as a consequence. A long chain of time and a long chain of causation run from the death to the skeleton. But in the blend, this long chain of causation is compressed: Death-in-General is a skeleton. The long-distance result becomes, in the blend, the form of the cause. That is a metonymy compression. It is a very strong conceptual move to take all those diffuse and extended outer-space connections and compress them down to a relationship of form.

If you look at this arm, or rather this robe, all along it you see that there is writing: "birdseed, orchards, garbage" and so on. In other words, all the food that can be provided to the bears becomes just the little biscuit handed out to the one bear in the blend. No one is deluded by this blend, but it is a packed conception you can carry around for unpacking in order to make it possible for you to deal with the world in a certain way. When you roam around in the world with this packed conception, it is supposed to change your behavior by providing you with guidance as you unpack it.

A fed bear is a dead bear. The rhyme is an attempt to create a greater

compression. If it said *a fed bear is a killed bear* that would probably have been more accurate because the California State Game officer is going to show up and kill the bear. But *a fed bear is a dead bear* is a much tighter phonological equation: fed, dead. This may seem strange, but a great deal of the world's poetry and in fact insults and things that children say to each other depend upon phonological equations as compressions of conceptual equations. The phonological equation is meant to be unpacked into a conceptual equation. The similarity in the blend on language is to be unpacked to different conceptual relations in the network. If you hear children in the United States taunting each other by saying "Mary, Mary, Mary's really scary", you will see that kids hate it. They hate it when you do something with their name, because the name is connected to their essence, and if you equate the name with the description, somehow the universe has provided an equation, somehow the very name now carries this conceptual association.

Once, when I was in the Sangre De Cristo Mountains, I heard somebody say to the tour driver "*at what attitude do the deer turn into elk?*" Now that's very funny because you can hear it: it sounds as if a deer is going to turn into an elk. What the speaker meant was that first you go up to the mountain on the road, and you see deer, deer, deer, deer, deer, deer, deer, then elk, elk, elk, elk, elk. No deer turns into an elk. But, once again, there are analogies and disanalogies running across all these perceptions. The analogies are compressed to identity and the disanalogies are compressed to change for the identity. Consequently, although nobody thinks deer turn into elk, you can say "*at what attitude do the deer turn into elks?*"

Here is a similar example where the packing of analogy and disanalogy seem to go unnoticed: *The new theory is that dinosaurs turned into birds*. In the blend, the dinosaurs turn into birds. It doesn't mean that any dinosaur turned into a bird or that there is any "turning" at all.

My tax bill—this is from Gilles Fauconnier—*my tax bill gets bigger every year*. This is a perfectly normal way to convey this meaning. In the United States and many western countries, we citizens receive a tax bill every year. We must figure out what taxes we owe to the federal government and we have to pay them. So an individual sees analogies and disanalogies running across all the tax bills he receives. In the blend, the analogies are compressed to identity—you have a tax bill—and the disanalogies are compressed into change—it gets bigger

The President changes every four years. That can mean that the same president changes every four years, but that's not the way it is usually interpreted. It is interpreted as indicating that there is a constant role of president, and that the value of the role changes every four years.

Let's go back and look more closely at those dinosaurs that turned into birds. *Dinosaurs evolve into the birds*, we say. Of course we know that in the theory of evolution, there is no teleology. There is no intention of the organism to evolve. We know that what happens is instead that there are some organisms, and some of them die, some of them live, and they all die sooner or later, and they have offspring and then those offspring die and so on. There are differences, disanalogies, across all these organisms

as you go through time. It's just irresistible to say things like "These birds developed bigger wings," when what you mean is that evolution gradually produced animals with bigger wings than their ancestors had. The analogies are compressed to identity—these birds—and the disanalogies to change for the identity. This is, again, an extremely common pattern of conceptual packing.

We say *dinosaurs turned into birds*, *dinosaurs became birds*. Strictly, scientifically—by which I mean, when we use special knowledge to investigate the full network—it's not true that dinosaurs turned into birds or that they became birds. And we know that. But it is true for the blend, and we can say it of the blend, and when we do, people understand how the blend connects to the network. The utterance is not true of the top-level of all those mental spaces containing all those individual organisms, but that's not why anybody makes such an utterance. The utterance is true of the blend, and the speaker expects us to be able to unpack the blend to connect to the network, so the utterance is, cognitively, true of the network.

There are many things you and I know about language that are false about actual performance. But they are good packets to carry around for unpacking into actual language usage. They are true of the packed constructions, and we know how to unpack those constructions during performance.

This investigation into how dinosaurs turned into birds began with an illustration I found in *Zoobooks*, which was a magazine to which one of my sons had a subscription. This boy received a new issue of *Zoobooks* every month. Before we rush by it, notice how easy it is to make a compression here. I open the mail, and I say to my son, "Here is *your monthly Zoobook*". In the blend, there is a cycle, repeating monthly, and your *Zoobook* shows up every month. All these different things are compressed into *your monthly Zoobook*. It's like *your afternoon coffee*.

Here, in the slideshow, is the image of the dinosaur turning into a bird, from *Zoobooks*. One dinosaur runs along a single curving track, and we see drawings of the dinosaur at different spots in the track. The one path and the consistent shadows suggest that this is a uniform scene. This dinosaur is chasing a dragonfly. With each new position, it becomes more bird-like. At the end it is a bird, and it catches the dragonfly. The dinosaur turned into a bird.

Of course, in reality, there were many different dinosaurs. They lived and died, and there were analogies and disanalogies across them over time. The analogies are compressed to identity—the dinosaur, or even a group identity, dinosaurs—and the disanalogies are compressed to change for that identity—the dinosaurs change. In the blend, they change as they go down the path.

Notice something else that has happened in this packing—something that happens very often in blending: we create new vital relations in the blend to make it stronger and more memorable. In this blend, "the dinosaur" wants to catch the dragonfly. It seems to want to evolve. That's false for evolution, but it's true of the blend, and we know how to hook the blend up to the network, so, in a way, the utterance is true for the network. We can carry this packed notion of *intention to evolve* around, mentally. It is very useful even though we know it is false of some of the mental

spaces in the network. We know that when we unpack this packed notion, we should not think that these individual organisms have an intention to evolve.

My favorite book when I studied mathematics was a small book titled *Counterexamples in Mathematics*. There are many things in mathematics that you are certain must be true in general, because they seem so right, but that are not true in general, and this book helped remind us of that. The packed compressions were very useful, so long as one remembered to take care in unpacking them. You say to yourself as you unpack it, “Wait, wait, wait, wait, wait, is this actually integrable over all these different topological spaces?” You stop yourself a little as you unpack, and check.

We do the same thing when we run financial accounts. We have packed notions of finance that we carry around, but we also have cautionary knowledge, to get us to be careful about the unpacking.

Once you have noticed the packing of analogy and disanalogy to identity and change, you will notice it everywhere. It is the main way in which political cartoons are constructed. The cartoons are instantly understandable. Or rather, they are understandable if you have the right cultural knowledge. If you don’t know who The Grim Reaper is, you won’t recognize him on the Lake Tahoe poster. I was in Berlin to give a talk, and the day I was going to give the lecture, the political cartoon I am showing in the slideshow appeared in a newspaper. The main headline was “World Food Crisis”. In this political cartoon, there is a pudgy American who is filling up his car with gasoline. It’s a green car and it says “GO GREEN” and he’s filling it with bio-ethanol. He has all the trappings of caring about the environment. In fact, part of what we see here is the cartoon is editorializing by the cartoonist. On the gas pump, the bio ethanol is called PURE CORN. Perhaps you wouldn’t know what this means—it’s a bit of cultural knowledge. In American English, “*pure corn*” means “nonsense”. You would never find a gas pump of bio-ethanol, which is a corn product, labeled “pure corn”. Ethanol is added to gasoline, and this is supposed to help the environment a little by not burning so much gasoline. The editorializing in “pure corn” is to indicate that the guy at the gas pump is stupid, that he has the wrong idea.

In the cartoon, you see over here two characters, whom we take to represent the hungry of the world. Notice that they are sort of standing on the other side of the curve that could be sort of taken as the edge of the globe, so they are in continents far away. This is completely spectacular. One of them is vaguely Asian, wearing a peasant horn hat, and holding up a rice bowl, in the classical gesture of supplication. The other is vaguely African, with the edema associated with malnutrition. The American is saying “Sorry, I’m busy saving the planet.” It’s completely impossible that the world’s hungry could show up, two at a time, in the gas station, in the American mid-west, and make a gesture of supplication like this. It’s completely impossible. The American cannot see the world’s hungry who are in different continents. They probably don’t understand English. But it’s a basic human scene. There is need. There is somebody right in front of you, who is asking for something, and you say no, and you give an excuse, and your excuse is silly, because it is so self-serving. You see this, and you are

now supposed to carry around this little packet to unpack when you think about the world food crisis. The political cartoon gives us a packed, human-scale scene that is supposed to be unpacked to understand American and the world food crisis. In the network, this isn't just one guy at the gas station in the American Midwest. This is now representative of America.

I'm going to pass over all the complexities of packing and unpacking we see in this political cartoon, because I want to give, for the remaining seven minutes, an application of this idea of packing and unpacking to the nature of small linguistic packets, small grammatical constructions, that we carry around. Let's take as an example something like the Caused-Motion clause in English. The caused-motion clause is one of the basic clause types in English and many other languages. Caused-motion is a very basic human scene: an agent performs an action on an object that causes it to move in a direction. For example, you pick up a stone and you throw it. You push the wagon. You pull the cup. You lift the cup. You perform an action on an object and cause it to move in a direction. This is a basic embodied scene. All human beings know it. It's in your bones.

Now there are certain verbs in English that are direct caused-motion verbs. A verb like *throw* is an example. *Throw* suggests an agent, in a frame with arguments, which means places for other elements. *I throw the ball over the fence.* *Throw* already carries an agent performing an action on an object to cause it to move in a direction. *Throw* probably also indicates something about the manner of the action, the means, and the manner of the motion that is caused. But there is also a clause — Noun Phrase, Verb Phrase, Noun Phrase, Prepositional Phrase — that is attached to this packed, human-scale scene of Caused Motion. You can use that clause to indicate to people that you would like them to make a blend, a caused-motion blend.

Suppose, for example, that there is a diffuse scene in which some stuff happens. There are some tanks; they want to come in to the compound. It's not so easy for them to come in to the compound. Some effort has to be taken in to get them to come to the compound, but they finally end up in the compound. So there is motion of the tanks. And I say "*How did the tanks get into the compound?*" And somebody says "*Oh, the officer waved the tanks into the compound.*" Notice the clausal structure here: Noun phrase, verb phrase, noun phrase, prepositional phrase. What that means is that there is a bunch of diffuse stuff up here you need to manage, involving tanks and compounds and motion and permission and so on. And I want you to take the caused-motion scene and blend all that diffuse stuff about the tanks with the caused-motion frame. The Caused Motion frame is a small, human-scale packet. When we use it as an input to the blend, we can project its compression. As a result, we get in the blend a tight compression of stuff about the tanks that was diffuse. Now, in the blend, we have an agent who performs an action on an object that caused the object to move in a direction. Even though *wave* is not a caused-motion verb, it now becomes a caused-motion verb, or a prompt for a caused-motion packet, to help us to conceive of the situation. In order to understand "the officer waved the tanks into the compound," did I need to have in

my knowledge that wave can be used in the caused-motion clause? Did I need to have heard that usage before, in order to understand sentence or even to say it? No, what I have is a little packet. I have a caused-motion clause and over here I have a verb *wave* for a communicative action that somebody makes. I put them together, along with many other elements. I blend them. I use this language as a way of asking you to blend them. And down in the blend, you now have a human scale packet that you can use. So it goes like this. Here, in the slide, is one input. This is your caused-motion frame attached to a caused-motion clause. Then there is a bunch of stuff over here in the slide, for lots of diffuse information.

Using this blending ability, we can say things like—this is Adele Goldberg’s example—*Paul sneezed the napkin off the table*. *Sneeze*, as she said, is a parade example of an intransitive verb. But now it is in a clause that takes three arguments. *Paul sneezed the napkin off the table*. When I say this, I am prompting you to make a certain kind of blend, a certain kind of packet. Suppose I say, “*Paul sneezed the napkin off the table*”, and you say, “what do you mean?” And I say, “Oh, Paul sneezed, so I thought he might want to blow his nose, so I lifted my napkin off the table and gave it to him,” and you say that’s not what I can mean because nobody is going to understand it that way. And I say “Why not? My expression, my story, my explanation, has all the elements: Paul, the sneeze, the napkin, and the napkin’s coming off the table”. But you do not accept my explanation, because you took me as prompting you via this caused-motion clause for a story in which Paul is the agent of the caused motion.

Using the caused-motion construction, I can say, “*Junior sped the car around the Christmas tree*”. Where does *speed* come from? It’s the manner of the caused-motion. Now I can take all kinds of things, all kinds of words from all over the diffuse network. And I can put them into the caused-motion clause, prompting you to come up with a network that you can manage because you’ve got a nice caused-motion blend.

Here, in this slide, is a representation of the cause-motion network. Notice that I can say things like “*I walked him into the room*”. *Walk* is not caused-motion verb. *He sneezed the napkin off the table*. *I pointed him toward the door*. These are all noun phrase, verb phrase, noun phrase, prepositional phrase. *They teased him out of his senses*. We’ll say this is metaphoric, and indeed it is, but the action and the causation are not metaphoric. It’s a caused-motion blend. *I will talk you through the procedure*. *I read him to sleep*. *They prayed the two boys home*. *I muscled the box into place*. *Muscle* now is part of the manner of the causal action but it can come in as a verb. *Hunk choked the life out of him*. Now pay attention to this. Hunk did not choke. And Hunk did not choke the life. What happened? Because of the caused-motion construction, we understand that Hunk did something with a result for an object and the result is that the object was dead. Metaphorically, we understand this as change of state; it’s change of location. So Hunk performs an action on something with the result that *the life* moves: it goes out of him, metaphorically. Think of *He floated the boat to me*. *Float*, all by itself, doesn’t have to mean there is any motion at all.

The caused-motion construction can also be used when you stop the motion. Consider *We blocked him from the door*. That’s pretty interesting. The point here is that

as I run around the world and you run around the world, I do not need to have complete command of all of the things that I might produce or that I might have to understand. What I must have is little packets of stuff in my mental and linguistic portable rolling bag that I can lock together with the context, with the environment, so as to create a network that is grounded in a human-scale scene I can understand, a scene that will let me understand the causal structure, the modal structure, and so on, that are in the network. To make and understand these expressions and construct these new conceptions, I do not have to have new language. None of these examples requires new linguistic constructions. These are all unifications—selective unifications leaving some things behind—of linguistic constructions that we possess and that we can use to prompt for the construction of elaborate networks.

It's the same general story for the Resultative construction: *I boiled the pan dry*. I can say, "No zucchini, tonight, honey. I boiled the pan dry". Think of the long causal chain involved in the result that the pan is dry. Notice that I didn't boil and I didn't boil the pan. How do we understand the resultative? We can say "*Cathy painted the wall white*", which means Cathy performs an action. You don't know what action she performed. Perhaps she used a paint sprayer. Perhaps she used a paintbrush. Perhaps she ordered somebody to do it. Perhaps she wasn't even there. But the result of the action she performed is that paint was applied to the object and turned the object white. *Cathy painted the wall white* doesn't mean Cathy painted the wall because she is white, or that Cathy painted the wall although it was white. It doesn't mean any of those kinds of things. It means Cathy painted the wall with the result that the wall became white. It's a resultative. And I can use it now for *I boiled the pan dry*. You don't know what I did, but you know I did something that resulted in the liquid's boiling. The liquid, the water, is not even mention. The result is that the pan became dry.

So grammatical construction like *I boiled the pan dry*, or rather like the resultative clause, or verbs like *boil*, and *dry* and so on, are little packets you can unplug and blend and hook to the world. You roll along a portable bag of mental stuff, which you can unpack and put together and use to construct a network to capture that situation.

Just as in the caused-motion construction, look at the compressions you can get with the resultative construction. *Catherine painted the wall white*. *She kissed him unconscious*. That's Goldberg's example. *She kissed him unconscious* doesn't mean she kissed him although he was unconscious. It means she kissed him so wildly or effectively that he just passed out. She kissed him with the result that he went unconscious.

Last night's meal made me sick. This causation runs over several hours. *He hammered it flat*. *I boiled the pan dry*. *The earthquake shook the building apart*. Here is one I love: *Roman imperialism made Latin universal*. *Latin* is not a thing and *universal* is not a feature. But in the blend it is. Roman imperialism is not an agent, but in the blend Roman imperialism becomes an agent that works on an object, namely Latin, with the result that Latin becomes universal. Latin has a feature. This is the same

general resultative network we saw in *Catherine painted the wall white*. But now it runs over centuries, and hundreds of thousands of people, and vast causal connections: *Roman imperialism made Latin universal*.

This is the topic we are going to be exploring in future sessions—the way in which linguistic structure consists of small construction that you can blend to prompt for blending in mental networks as you pack and unpack in order to understand the world. Thank you.

Lecture Five

Big Ideas

I thank Professor Dirk Geeraerts for the presentation he gave on the importance of historical, diachronic linguistics for the analysis of metaphor. This afternoon at 3:30 I will give a talk on further advances in metaphor analysis and metaphor theory. So you start today with new advances in metaphor analysis and you end the day with new advances in metaphor analysis, and they will be supplementary.

Some of you have asked me what my email address is. If you go to my web page, <http://markturner.org>, you will find my email address. All you have to remember is “Mark Turner”. Use any search engine, and you will find that page, and all questions about how to contact me or what I am working on will be answered. You can find the Cognitive Science Network by going to that page. \

Yesterday, we talked about packing and unpacking, and how concepts come from packing and unpacking. The idea that words mean or that they have meaning or that expressions mean or expressions have meaning is one of these packed conceptions. It is completely false if you are talking about the distributed conceptual structure throughout the network, but it is true of the packed blend, and we know how to connect the blend to the network. Words do not mean. Words do not have meaning. Expressions do not mean or have meaning. These notions result from mental packing.

Packed conceptions are very good things. We carry them around and we use them, and sometimes we believe them because we think they are true. And when we do that, we are living in the blend; we’re living in the packed meaning and we are not seeing outside it.

The idea that words mean has led us to mistaken notions about the truth or falsity of an expression. If I say something that is true of a packed blend, and I can rely on you to connect the blend in the right ways to the network that it anchors, then my language could be called “true” even though it is manifestly false as applied to the parts of the network on which I am trying to comment. Expressions are prompts for people to construct meaning. If I say something of a blend, and expect you to connect the blend to the network, then what is important is that what I say be accurate for the blend, not for the rest of the network: you will do the work of connecting the packed blend up to the network.

I want very much to try to get you to unpack these notions having to do with expressions, meaning, and truth. Words do not mean. Words are prompts that human beings use to try to get someone else to construct mental meanings. Words and expressions are tools that human beings use to get someone to use mental operations that they already possess to work on knowledge that they mostly already know to put together a meaning. The words do not mean. The words do not carry meaning.

There is an elaborate causality here. I have little tools right now. I’m trying to get you to make meaning. And I’m using these little tools. It goes from my use of tools to the work that you’re doing to the integration networks that you’re putting together and to the meaning that you understand. That’s a long causal chain with a lot of agency. And you can compress that down and think that words have meaning the way things

have color or feature. That's very much at human scale. This wood is hard. It has a red color. And I can think, following the same compressed lines, that words have meanings the way objects can have color.

Similarly, we think that agents do things. That's a thoroughly human-scale notion. I'm an agent. You're an agent. The pilot can fly the plane. The painter can paint. That's agency. When I talk, I am an agent engaged in trying to get you to construct meanings. We can compress down that whole network into the idea that the word itself is an agent, that the word is doing something, that words mean. That's true of the blend, but it's false as a scientific theory. Words don't mean. Expressions don't mean. Meaning is not in words or in expressions. Meaning is in the mental operation that you conduct. The little prompts that I use to try to get you to understand something profoundly underspecify the meaning. There are many things that you can do with those prompts and almost everything about the meaning is not anywhere carried in the words.

We saw this yesterday, when I gave a list of expressions containing the word *father*. If you had said, "Well, what is the meaning of *father* that is in these expressions", that would be the wrong place to start. That is a mistake. When I use the word *father*, you understand me as thinking that I would like you somehow to use the frame of kinship relation. You have that frame, and in using the word *father*, I am asking you to do something with that kinship frame, as you are building an integration network.

In fact, in all those cases, I wanted you to do quite different things with the frame. I wanted you to do selective projection from the frame, and I wanted you to develop emergent meaning in the blend, very different emergent meaning. So, consider again

- *Paul is the father of Sally.*
- *Joseph is the father of Jesus.*
- *The pope is the father of all Catholics.*
- *The pope is the father of the Catholic Church.*
- *Newton is the father of physics.*
- *Fear is the father of cruelty.*

It's just a mistake to ask what *father* means in these expressions. It is a mistake that comes from living in the packed blend. In that blend, words have meaning—features. And words are agents. But the blend is false as a scientific theory. Words do not mean. "Father" does not have a meaning. It does not carry a meaning. A word is a sound. It is a tool I use to prompt you to construct a meaning dynamically.

It may be that there are some meanings that you construct again and again and again and again and again and again like *father* for a particular kind of application of a frame to a domain. You do it again and again and again and again. Fine, you're doing it and it's frequent. Don't mistake that for the word's having a meaning. The word is a tool I use for you to construct the meaning and I can use it in many different ways. In *The Way We Think*, Gilles Fauconnier and I consider the word *house* and we point out that the use of the word *house* as a tool may prompt people to construct certain kinds of integration networks and some of them we construct again and again and again and again and again and it becomes quite frequent and entrenched and we're happy with

them. That doesn't mean the word means something. It's a tool to get you to construct that network, even if it's one that's very familiar to you. And we can use it in lots and lots and lots of different ways. Already there are lots of data indicating different uses of the word *house* to get you to construct very different meanings, and we can come up with all sorts of novel meanings that don't even feel novel.

Why is this? The answer is: because you are able to construct conceptual integration networks, and I can use these tools to prompt you to do that. You construct meanings. I construct meaning. A word does not know that it has meaning. It does not even know it's a word. It is not an agent. Thinking that the word does something, that it is an agent, is a compression. Thinking that it has a feature, namely a meaning, is a compression. It is a very useful compression to have in our culture. We can say, "Oh, in Chinese, what does this word mean?" and someone can say, "Oh, it means beef, cow meat, or pig, or whatever." And that's very useful. It's very useful. When you say such a thing, you are helping me construct a meaning. This is a useful compression to have, but it is not a respectable theory of language. Don't live in the blend.

We are going to investigate these issues today. Let me give you an analogy. Beihang University, where we are today, is a university that has historically specialized in aeronautics and astronautics. I have seen wonderful warehouses of airplanes in parts here. Suppose I have a toolbox. Suppose there is a terrific Chinese graduate student here who has a tool box, and we take our tool boxes to the museum, the warehouse, the junk yard of all these old planes, and we assemble a plane and then we fly it. Yeah. We do it. That's what we do with our tools. We act.

Our toolboxes do not contain an airplane! Our toolboxes do not fly airplanes! Our toolboxes do not know airplanes! No, our toolboxes are things that we use to make the airplane. And we can use it to make a car. We can use it to make a boat. We can use it to make things we want. But the tool box is not magic. It is not an agent. It does not have a plane. It does not have a boat. It does not have the ability to fly. We do. We use these tools to make these things.

Let's look at some of the mental packing that we observed yesterday. I pointed out—hold on, how much time is left? There is no clock in this room, but luckily, I have another tool, my international SIM card telephone, with its China Mobile SIM card. It definitely does not know the time. But I know what time is—meaning that I have, just as you have, some incredible complex conceptual integration networks for time, and with this tool, I can determine what will count for us all as the time in Beijing.

Yesterday, I pointed out that I can use the caused-motion clause structure in English to get you to construct meanings. So for example if there is a diffuse set of stuff, I can suggest that you use the caused-motion frame to blend with all that stuff so as to come up with a meaning. We looked at the example, *we blocked him from the door*. In the caused-motion frame, the agent is causing an object to move. But *block* is another tool that we use to indicate a kind of force-dynamic stoppage. That's not what *block* means. It doesn't have that meaning. *Block* is something I say to you so you can construct a meaning. The word doesn't understand force dynamics. It's just a sound. You use this to say, "Ah, I should somehow use the force dynamics of blockage, of

stopping, and I am supposed to blend that with the caused-motion frame.” And you come up with a wonderful bit of emergent structure. In the blend, the agent—*we*—is causal: *we* take an action that affects the motion of something. *We* affects the motion, not by causing the motion, but by stopping it. This is emergent structure. You take the caused-motion frame. You take your force dynamic understanding of stopping. *Block* is a good tool for getting you to call this stoppage up. You put them together. You blend them, and in the blend you now have a meaning. The blend is mental. In the blend, you now have a meaning in which *we* perform actions that affect the motion of something—*him* in this case—and the way the agent—*we*—affects that motion is by stopping *him* on the path on which he was going to the door. That’s emergent meaning. We are packing this situation into something like caused-motion.

Similarly for *I boiled the pan dry*, there’s a resultative frame, a resultative clause that I can use to say I want you to use the frame of someone performing an action that causes a result for something; I want you to use that understanding, and I want you to use that understanding on words like *I* and *boil, the pan, dry*. And you construct a scene. The point I want to get to is just this. This is where we ended yesterday. *Roman imperialism made Latin universal*. Think of the vast range of meaning that this expression prompts for. Roman imperialism took centuries. It had lots and lots of agents, thousands, tens of thousands of them, lots of actions, lots of events. This is not something that human beings are built to understand. It is beyond our scale. It is beyond the basic mammalian scale of the visual field, of interaction with agents. But no problem, because the resultative frame is something we do understand very well and we can blend that basic resultative meaning that we know with diffuse structure to achieve a compressed understanding at human scale of something that is beyond human scale. There’s an action that causes a result for an object. We can blend that frame with all this diffuse structure to get a human scale blend we can grasp. In the blend, Roman imperialism is now an agent, and it’s a causal agent, and it causes a result for the object, and the object is Latin. Latin is now an object in the blend. And the result caused for the object is that the object becomes universal, which is kind of a funny feature when you think of it. If you paint the wall white, it’s white. That’s the result. There’s a wall and you do it and it’s white. Now there’s this thing called Latin, and Roman imperialism can perform actions that cause Latin to have a feature and the feature is that it is universal.

We can have very, very big ideas that run across time and space and agency and causality that are very far beyond what we are built to understand, because we can construct integration networks that have anchor spaces in them, spaces that we are built to understand, and from that platform we can grasp and hold on to and work in an integration network that otherwise we wouldn’t understand.

That’s the topic of today’s discussion—how we can have such big ideas by anchoring them in packed blends. I said in my opening lecture that this was a major question for cognitive science. For all other species, the mental understanding that members of the species can manage seems to be very tightly limited in time and space, very tightly limited. We can talk about those limits. I love animals, and we can talk

about the ways in which some animals for some activities with certain kinds of support seem to be able to understand just a little, just a little beyond very local scale.

The issue is *not* how evolution can build things into the local scale, into the local organism, that have long-range causality. For example, you don't need to be able to think about great-grandchildren in order to be attracted to a member of the opposite sex, because evolution has built desire right here into the local scale, and desire will run a local scene that ends up having reproductive causation. So, let's be clear, the question is not at all how you can operate or evolutionarily how you can be designed so as to operate in the local scale and have long range consequences.

The question is instead how you can *conceive* of long-range causality. How can human scale thinking and feeling be involved in long-range causal patterns? As I said, here, in asking this question, we need to avoid "living in the blend". We must unpack our own sense of our own thoughts to see what is going on.

In our own sense of our own thoughts, it seems to us as if we have very stable thoughts, very stable ideas, very stable conceptions in the head, and they just stay with us—memories, things like that. But in fact, biologically, your brain is always changing. Your brain is always activating different things. It comes and it goes, and nothing stays. All you have is the biology that you are running right now, biology interacting with environments right now. That's all that is here. The only thing you have for these long-range ideas, which are not here, which are not now, it's the biology that you do have here and now. The English neurophysiologist Sir Charles Scott Sherrington said the brain is

"an enchanted loom where millions of flashing shuttles weave a *dissolving pattern*, always a meaningful pattern though *never* an *abiding* one; a shifting harmony of subpatterns."

With what you have in the present, right here, right now, you are able to think of ideas that are not restricted to the local understanding of agency, causes, and scenes.

Instead, you can think big ideas. How can you do that? You might object, saying, "Well, wait a minute, wait a minute, doesn't memory bring the past into the present? Don't I have long-range ideas from the past that are just available to me?" The answer is of course NO. Memory is right here, right now. Memory is your brain running right here, right now. It feels, when we remember something, as if a window opens from the past, and the past comes into the present. But these different times never coincide. The past never comes into the present. That never happens. It just feels that way because that's the way we think. But don't live in the blend. What's happening is that you are running right here, right now, processes that present to you what feels like the past, but it isn't the past. It's the present. Far from memory's solving the problem of how we can have big ideas, it presents an extra problem: How can we have such memories? How can we feel at human scale about things that are beyond human scale? How can human-scale thinking and feeling be involved in long-range causal patterns? And what's the role of language in that?

I want to introduce to you a wonderful machine. There's a character named Doctor Who on British television. He has a machine called a Tardis, which is an

acronym for Time And Relative Dimension In Space. He gets in his Tardis and it takes him through time and to different places in space. Here in the slideshow, you see the Tardis. It's a blue police box, and it flies through time and it flies through space. So he can put together big ideas that run across the entire universe throughout the entire chronology, because for him, the big ideas are often at local scale. For him the past and the future really do intersect with his present.

Make sense. It seems to us as if it's possible mentally to move around in time and space. And indeed we do have a Tardis, each of us. Our Tardis is the ability to do advanced blending, advanced conceptual integration, to compress things so that we can have big ideas that run across time and space and go very, very far beyond the scale of thinking that any other member of any other species can even begin to conceive. We are extremely distinctive in this way.

Here, in the slideshow, is Doctor Who's Tardis, and it seems to us as if we have a Tardis. Here, in the slideshow, is a quote from Charles Fernyhough, in which he uses just the same metaphor that I want to use for escaping human scale. He says,

“A self can feel such a singular fixture, hugging one's here-and-now like a twenty-four-hour undergarment, but actually it's a string, looping back and forwards in time to knit together our past and future moments...”

That's the problem. How can we do that? No other species does that. In a sense, we are not built to do it, in the sense that it is not at human scale. But we are built for it in another sense: we can do advanced conceptual integration, which allows us to manage network conceptual scale far beyond the usual human scale. Fernyhough writes,

“A self is a Tardis, a time-machine: it can swallow you up and spit you out somewhere else.”

Here, in the slideshow, is a picture of Dr. Who's Tardis, and here is a picture of our Tardis: the human brain. I am going to give you an example of how our Tardis works, how it lets us go beyond human scale to have big ideas at network scale, which is so easy for us, but only us. Let's look at some big ideas.

In Al Gore's movie, “An Inconvenient Truth,” in which he's trying to get everyone concerned about climate change and global warming, he says that the earth is a “pale blue dot” if you look at it from four billion miles away. In fact, he gives you a picture, which is here in the slideshow. I think it was Carl Sagan who made this picture famous. You can see the Milky Way, and you can almost see the pale blue dot in the Milky Way. Here, in the slideshow, is a close-up of the pale blue dot. That's the Earth. He's going to do something very interesting. Notice, by the way, that I can point at this single blue pixel and say, “that's the Earth,” and you have no trouble following me. What you have just done is an example of blending. The pixel you can see, right here, right now, is blended with a conception of the earth in space that is certainly not right here right now, that is certainly not at human scale. But in the blend, it is human scale: it's right here, in your visual field.

The picture was taken from a distance of four billion miles. Gore says,

“Everything that has ever happened in all of human history has happened on that dot. All the triumphs and tragedies, all the wars and all the famines, all the major

advances. That is what is at stake—our ability to live on planet Earth, to have a future as a civilization.”

Now what’s going on here? You cannot see the earth from four billion miles away, not least because you cannot travel four billion miles. It’s not even clear independent of investigation that it would be true that if you were at a distance of four billion miles, your vision would work the normal way. There are many things in physics that do not work at great distances or at small time scales in the way we expect them to, on the basis of our experience at human scale. So why does this picture seem so natural. It depends on a mental blend: you know that if you back up, something stationary in your field of vision subtends an increasingly smaller angle as you back up. In everyday language, it gets smaller. You can take the idea of the Earth and your human-scale knowledge of backing up and blend them, and in the blend, you back up four billion miles. This is completely impossible of course. You can’t do that. But you can imagine it, by blending. And now in the blend, what’s the Earth? It has gotten much smaller in your field of vision, until it is this one little tiny thing. And you can see all of it in your visual field.

Let’s think of our experience of visual field. What is in our visual field is at human scale. It can affect you. You can affect it. Think of a child’s falling down in front of you and crying. You feel some responsibility for dealing with the situation. You feel something immediate, right there. Where is a child who falls down 10,000 miles away? You don’t even know. And if it does happen, you don’t feel the same responsibility in the same way. But now in the blend, the Earth is something small in your visual field. You are used to having some kind of authority, some kind of responsibility, some kind of engagement with your local environment. So suddenly, in the blend, if you follow the projections that Al Gore wants you to follow, you feel some responsibility for this thing that is in your environment, in your visual field, this Earth. Suddenly, in the blend, you are now responsible for the Earth, and you might even be able to do something because it’s small enough and it’s at your scale.

So Gore creates a spatial blend that compresses great huge distances down into something you can see. He also invites you to compress history: he takes all of the events of the Earth and compresses them down to one thing – human civilization. All history, all civilization, is now one thing, and it is right here on this little dot.

This is an amazing compression. Everything becomes one thing, and it’s your responsibility. And it’s small enough for you to have power.

Next, Gore is going to perform a time compression involving the future, because he wants you to be concerned about the future of the Earth, when none of us will even be here. In order to get you concerned about the Earth, he has to compress it in both time and space. Otherwise the things that he cares about would be too far away for you to be concerned with.

So he says — I’ll make the text in the slideshow bigger — he concludes the film with this blend:

“Future generations may well have occasion to ask themselves, ‘What were our parents thinking? Why didn’t they wake up when they had the chance?’ We have to hear

that question from them now.”

So think about what he’s talking about: future generations that are not here. He’s talking about things that don’t exist and that we are not going to see. They are not in our time scale at all. They are not even close. And there are billions of agents in those future generations all over the Earth and they don’t all speak English. And we are not only talking about their asking questions. We are also talking maybe about what they think or how they feel or what they write. All that, by all those people, in the future, all over the world, in lots of languages, in thinking, writing, and speaking, all that gets compressed down into a few agents in the blend who are at human scale.

Everybody knows the human-scale scene in which someone asks you a question that you must answer. In that scene, the questioners are right there in front of you. And everybody knows about children scolding their parents or being disappointed in their parents and asking them this question “Why didn’t you do this for me? Why did you do that to me?” Parents have responsibility. All those future generations are compressed in the blend down to some children saying to their parents, “Why are you doing this bad thing? Give me an answer!”

Notice that this packed blend works even if you don’t have any children. In the blend, you have children. You are not deluded, but now in the blend you have children. You can imagine having children. Having children who speak to you is very much at human scale.

Gore says, “We have to hear that question from them now”. You take the scene in which somebody asks you a kind of challenging question, maybe a scolding question. That’s at human scale. You take the frame of responsibility for children. And you blend this into a scene into which you compress all of that stuff in the future that is at a huge temporal distance from now.

The purpose of this blend is to pack into one scene, to guide you to pack into one mental scene, a portable conception that you will carry around with you and unpack as you operate in the world. Maybe you will accordingly prefer a car that burns a different kind of fuel or maybe you will want to plant certain kinds of vegetables, or maybe you will want to contribute to a certain kind of community effort or something. *None of these actions is answering your children who are not here yet*, but you have this responsibility in the blend, in the time blend and the space blend, and you unpack it to actions that you do take here and now. Gore’s presentation of this blend is a rhetorical tool to lead you to have a new conception that is at human scale, memorable, portable, one that you can hold on to, one that you can express using the language you have for human-scale scenes. With these expressions, you might lead other people to manufacture and retain and carry this packed conception around, to unpack as they engage the world. This is an extremely common method in rhetoric.

Here, in the slideshow, you see an advertisement. It says *Joey, Katie, and Todd will be performing your bypass*. There are Joey, Katie, and Todd, and the patient that you see—that’s you. Actually you seem to be in two spots in this advertisement. You are on the table, about to be operated on, but the children are also looking at you. You are on the table, and these children look about eight years old. Katie is just about to cut

into you with a knife. Now this is a scene you will never see. You will never see eight-year-olds performing heart surgery.

What is the point of this advertisement? Well, what it says in the smaller print is

“before you know it, these kids will be doctors, nurses and medical technicians, possibly yours.”

If they are going to be any good at these things later on, they need to be educated now. The purpose of this advertisement is to get you to change your action now, to contribute personal money to schools—that’s the way we do it in United States; we contribute personal money to schools so that those schools will be able to educate those children. The point in the future where their knowledge is important is so long from now that it’s hard to be motivated, hard to think about it, hard to get up the energy to do anything about it. It’s on a scale that we are not built for. But good things will happen if we can work on these things now, and in order to motivate ourselves, we need a time compression. It’s very easy to make these time compressions, once you think of it. After you leave this lecture hall, you will see several such time compressions, and hear several such time compressions, before you go to sleep tonight. These time compressions use identity connectors. Here are these children, in one mental space, and they are eight years old. This is already a compression of the sort we have discussed—many schoolchildren in the network are now compressed to these three in the blend. Later on, in a different space, they are going to be adults; let’s say, they will be 35 years old. There is an identity connection between the children and the adults. And there is a time connection. You compress the children and the adults into one, and compress all that time scale to right now. Now, in the blend, the eight-year-old and a 35-year-old doctor are the same. You also compress down the time, so that, in the blend, you have maybe ten minutes to teach Katie how to do surgery before she cuts into your heart. This is a scene that you can really respond to. You can carry around with you, in your portable little mental rolling bag. And it will change your behavior for long-range effect. No one is deluded into thinking the blend is *true*, but again, we know how to connect the blend up to the network, so it is a mistake to call the blend false. The blend is telling us something that is true about the network. The blend is not something we *reify*, but it is delivering truth in virtue of our connecting it to the extensive conceptual integration network that it anchors.

The language in the small print also uses these time compressions. It says, “before you know it”. That’s a time compression. It asks us to take something that’s thirty or forty years away and move it down to right now. The language in the smaller print of the ad says, “If we want children who can handle tomorrow’s good jobs . . .” Notice that “tomorrow” here means *tomorrow* only for the blend; for the network to which the blend connects, tomorrow means decades from now. The ad is using “tomorrow” to mean, for the network, “far from now, maybe 30 years from now.” I’m using tomorrow to mean far from now, 40 years from now. We use these time compressions in language often: *Just yesterday it was not possible to visit China, maybe tomorrow.* We use these time compressions all the time in language.

Here, in the slideshow, is another example of a big idea, ranging over vast networks of time, space, agency, and causation. I heard this advertisement on National Public Radio. The purpose of this advertisement is to lead people to change their ideas about consuming fish.

Now, I don't see the fish out of the ocean. I see the fish off Del Mar when I surf, but I don't see all those fish out in the ocean that fishing fleets pursue. I don't see where the fish come from. I don't have any experience of those fish. But the makers of this advertisement want me to think about all those fish. The advertisement goes,

"We are eating the food off our children's plates. When we overfish, we eat not only today's fish but tomorrow's fish, too".

Let's analyze this. Notice, by the way, that "tomorrow" here means *the day after today* in the blend, but for the network to which the blend connects, it means years and years.

What's going on here? Let's look at the language and the concepts that we construct in response to this language. Well, there is a long chain of stuff. Think of fishing. Right now, in the ocean, there is a certain amount of fishing. In the slideshow, you see a mental space with a line whose length indicates the amount of fishing. The amount of fishing is causal, over time, for the amount of fish in the ocean and for the amount of fish that get eaten. In the slideshow, you see two more mental spaces: one for the amount of fish in the ocean, with a line whose length indicates the amount, and another for the consumption of fish, with a line whose length indicates that amount. These three spaces are connected by vital conceptual relations of cause-effect and time. The amount of fishing is causal for the amount of fish I eat or people eat, the consumption of fish.

It's important to recognize that fishing for fish and consumption of fish are two very different things. It's crucial to hold that distinction in mind as we do this analysis. Fishing and eating are very, very different, and they are far apart. Then, of course, there is the a long-range causation from fishing and consumption to the stock of fish that is available later on in the ocean—how many fish there are later on. You see the time relations and the causal relations. This ad is performing a time and causation compression much like the one you saw in "An Inconvenient Truth".

But now there is a different set of three mental spaces. In one of them, there is a desired, counterfactual, lesser amount of fishing. In another, there is a lesser amount of consumption of fish. And in the third, there is a higher number of fish in the ocean. You see these three lines.

Now you see in the slideshow a comparison of the two different spaces for the amount of fishing—you can compare the two lines. In the second one, there is a smaller amount of fishing. In the first one, there is a larger amount of fishing. You can also see in this set of comparisons in the slideshow that if there is less fishing, then there is less consumption and so more stock. You also see all the temporal and causal vital relations across these spaces.

But now we can blend each of these two contrasting spaces. Take the one for fishing: when we blend the two, the difference in the two lines becomes *overfishing*. What is *overfishing*? It's a compression. It's just fishing in this space, but now, when

we blend the two spaces, it becomes a new kind of thing—overfishing. Similarly, if you blend the two spaces for consumption, now you have *overconsumption*, that is, the amount of our consumption that happens that is greater than the amount in the second, preferred mental space. And now let us blend the two spaces for the amount of fish. In the blend, all the fish from the second, preferred space are conceptually there, but they have a feature: they are *missing*. The disanalogy and counterfactual connectors between the space are compressed to a feature in the blend: *missing*. The same thing happened in creating the *over* feature for the fishing and the consumption.

Because of blending, we can say, “Oh, if we only fish less, there will be a lot more stock.” You are running the blend and doing the inferences in the blend. The blend in this case is elaborate, working on analogies and disanalogies across two different networks: the network for actual fishing, consumption, and stock, and the network for the preferred amounts of fishing, consumption, and stock.

Let’s go back to the blend of the actual stock and the desired number of fish. The idea is that the desired number of fish would be there if we changed our amount of fishing and consumption. But they are not there in the mental space for the actual stock. But in the blend, all those fish that would be here are conceptually present, with the feature *missing*. In the blend, the amount of fish in the ocean comes in from this big line in this mental space, the one for the preferred network, but some of those fish have the interesting feature of being *missing*. They are *missing* fish.

Think of what the word *missing* is used to mean. When I use it in this way, it means that I want you to put together an integration network in which there is an element that has a certain feature, and that feature is not *white* or *brown* or *wood*; the feature is *missing*. We are built to understand that there are things in the world and that they have features. So, you take this extreme long-range dependency of contrasting hypothetical situations, and pack it down into something that’s at human scale. Now the blend has objects with features. It has classes of objects with features.

Let’s back up a second to make an important point. When you conceptualize a scene you are in, you are not conceiving of everything that is not in it. You could work that out, but you are not conceptualizing all of that. I’ll give you an example. We are in a lecture hall in Beijing, China. Look around. Do you have a good grasp of this scene? Good. Now, suppose I say, “*President Barack Obama is not in this room*”. Once I say it, you recognize quickly that it is true. You can just check. But it is extremely unlikely that your conceptualization before I said that sentence including the absence of President Barak Obama in this room. There are many other things not in this room: the planet Mars, tomorrow’s *New York Times*, a shark, even a tree. Once I point out that these things are not in the room, you understand that. But it wasn’t there before. There is an infinite list of things not in this lecture hall. When we point out that one is not here, we must make a blend of this room with what is not here, so that in imagination it is here. You recognize that this imagined scene is counterfactual with respect to the present scene, and we compress the imagined scene and the present scene to a new conception of the present scene, so that now a conception of the element is active, mentally, and in the present scene it has the feature *missing*. Now, you can say that you know that President Barak Obama, tomorrow’s *New York Times*, the shark,

and the tree are *not here*.

This is the kind of mental blending involved in conceiving of the stock of fish in the ocean in a certain way, namely, there are lots of *missing fish* because of *overfishing* and *overconsumption*. And if we keep it up, even more fish will be *missing* later on.

In all these cases, we have emergent structure. Before, we just had consumption, but now some of it is *overconsumption*. We use these words, such as *missing* and *over-consumption* and *over-fishing*, to prompt for a large conceptual integration network that has human scale blends anchoring it.

Recall that I said that words don't have meanings. They are prompts to us to construct meanings. *Missing* and *over-fishing* and *over-consumption* are prompts to us to construct certain kinds of conceptual integration networks. Words are prompts; they do not encode meanings. They prompt us to construct meanings.

This kind of double-scope blending is not cognitively costly. You're doing it all the time. 99.99% of the blending you do, you never even recognize. It happens in backstage of cognition. Most of it does not have any access to behavior. Certainly the great majority is never recognized in consciousness because it doesn't need to be. Consciousness is just a thin little reed. It gives you some little products that pop out of this elaborate construction of meaning. Blending is not cognitively costly and you can't shut it down.

But there's more. *We are eating the food off our children's plates. When we overfish, we eat not only today's fish but tomorrow's fish, too.* Look at the compressions. There is another compression, of the fishing and the eating. In the blend, overfishing is eating. That's a compression of time and causality. The overfishing is eating in the blend. That causal connection has been compressed so the cause and effect are now one action. Overfishing is eating tomorrow's fish.

Similarly, the harvesting of the fish and the eating of the fish are collapsed. Harvesting present fish and eating future fish are collapsed, condensed, compressed in the blend. In fact, we're bringing into this compression fish that our children would eat in the future. They are not even here yet. They haven't been born, those fish. But now our overfishing right now is eating fish 20 years from now. In the blend, the present overfishing and the quite distant future eating are fused. So overfishing takes on the moral value of eating this food off your kids' plate. The present day children are blended with their adult versions, just as we saw as in *Joey, Katie, and Todd will be performing your bypass*. Overfishing now is taking kids' food away from them now, because they won't be eating it when they are adults.

There's more compression. You scale the temporal distance from childhood to adulthood down to one day: *tomorrow*. Just as we saw in *tomorrow's good jobs*, in the network it isn't tomorrow's jobs.

There is yet more blending. We are asked to call up an indulgence frame by this ad. If you're eating more than you should, if you're eating food right now that you were supposed to save for tomorrow's meal, that's socially disapproved and it's

indulgent. You shouldn't do that. That's naughty and in particular you shouldn't eat somebody else's food.

But there is yet more blending. Taking food needed by your children off your child's plate and eating it yourself is very bad behavior by a parent. So in the blend, fishing now is the same as depriving our children by stealing the food off their plates and eating it ourselves. This is bad. So now you have a compressed little packet, a portable packet, to carry around, and to unpack for understanding the world. Now, if this ad has been successful, any of your actions having to do with fishing and eating fish will be understood in part through this packet having to do with overfishing, indulgence, depriving children, and missing fish.

By the way, I should remark that in America, for this ad to work, we probably have to say that when we overfish, we're eating the *food* off our children's plates, because in the United States, but probably not China, it is a stereotype that fish is not a favorite food of children. They eat fish sticks, which are mostly unrecognizable as fish. One of my children eats mussels, and all three of them are fond of salmon, but they are a little unusual in that regard. Notice that there is another compression going on: the fish are probably not food for the children, but in the blend, overfishing is eating fish which is taking your children's *food* off their plates. Since what you are, according to this logic, depriving children of is not food they would eat now but rather food that they would eat as adults, there is a compression of *fish* from the space of our fishing and eating, *fish* from the space of our grown children's eating, and *food* from the space of our children's eating food children like. This blend of missing fish and overconsumption and overfishing provides a human scale scene that can be used to understand the complicated network.

Once again, we have a compression for understanding big ideas, as we saw in "An Inconvenient Truth," and "Joey, Katie, and Todd will be performing your bypass". It's not that human beings have infinite minds that are not restricted to the scale of our species. We are restricted, except that we can use conceptual integration to bring ranges of stuff we are not built to understand to something we are built to understand, so that we can grasp it. We can understand network scale by anchoring it in a human-scale blend.

We saw compression in the "Mythic Race" mirror network where the fastest milers over the last six decades were placed on the track: in the blend, those six decades are compressed to one moment, in which they are all racing against each other. When we see such a compression for the purpose of understanding big ideas, we are not deluded. We do not think the blend alone is true, but we think truth is given by connecting the blend to the network. We saw such compression in the advertisement with The Grim Reaper feeding the bear. In that network, you feed the bears—e.g., you leave food in a garbage can, or plant orchards—and gradually over time you change the behavior of the bears such that they intrude on human environments and get shot by the state ranger. The network is compacted down to one scene where you feed the bear, and you are the grim reaper. Nobody's deluded. We saw that integration network.

We saw these kinds of compressions also for the dinosaurs. The dinosaurs turn into birds. This is compressed into one dynamic scene, with an intentional plot involving direction to a goal. In this network, we are using the blend to understand change, analogy, and disanalogy across millions of years.

We almost never recognize that we are performing these kinds of compressions, to create packed, portable conceptions for unpacking to plug into environments. Vera Tobin, in her brilliant, dissertation, quotes Hugh Kenner, a literary critic, discussing a poem by Marianne Moore. The poem in question is called “Poetry.” Marianne Moore revised it many times. Notice that I say *it*, but what is this *it*? Think of what Marianne Moore did. She wrote some words on paper, and she did this very many times. There are analogies and disanalogies across all those events, all those words, all those pieces of paper. Literary historians still have all of the relevant records of those events, and of the events of publication. So what is this *it*? What is *the poem*? Well, by now, you know the story: the analogies across all that writing are compressed to an identity—*the poem*. And the disanalogies are compressed to change for that identity—*revisions of the poem*. It’s like the fences growing, the cars getting three feet longer, the dinosaurs turning into birds. The analogies become a unique element—the poem, and the disanalogies become change for that element, revision. Hugh Kenner, as Tobin quotes, refers to Marianne Moore’s poem “Poetry” as “the one scarred by all those revisions.” In the compressed blend, there is a thing, actions are taken on it, and the actions scar it. Of course, in the input spaces to the blend, there is no *thing* that she worked on. But in the blend, there is.

This kind of compression, this kind of construction of big ideas in an integration network, is also the means we have for conceiving of our own existence, our own identity. What is an individual human being? We have a conception of ourselves as having a personal identity that runs over time. Despite the fact that the baby who came out of my mother is so different for me, Mark Turner at age 55—just look at all those differences—there are still analogies over all those mental spaces over time. And a certain array of those analogy relations is compressed into an identity in the blend. The identity is the individual identity—Mark Turner; look, there I am. Culture tries to tell me that this analogy is very important, and it gives me a proper name so as to encourage all of us to compress the analogies to a personal identity. Culture calls me “Mark Turner”, as if I am one thing. I was given this name when I was born. Some people get their names when they are baptized. There are many cultural moments of naming around the world. The proper name I was given is an encouragement to us all to compress the analogies to a unique element. The disanalogies are compressed to change: we say that Mark Turner changed. Mark Turner learned something. He received certification.

I know what you are thinking. You are wondering, what can this lecturer possibly be saying? Is he saying that individuals don’t “really” exist? That’s not the way to think of it. Having a notion of self seems to be very useful for us. We have notions of self that run over time and space. But look at it this way: there is no evidence that any member of any other species has any conception like this about their personal identity and identity that runs over years, decades. They just don’t work that

way. We cannot interview animals, but to the extent that we can see their thinking in their behavior, there is no evidence for anything at all like the overarching, big-idea conception of self that a human being possesses.

“When I was a child, I spoke as a child, I understood as a child, I thought as a child: but when I became a man, I put away childish things. For now we see through a glass, darkly; but then face to face.” These are the kinds of conceptions we have. They are highly creative, and distinctively human.

We can hold on to the notion personal identity, because of blending. A sense of personal identity is one of those things we pack into our mental rolling bag to take with us, to plug into the world and so engage with the world. This construction of personal identity makes it possible for culture to have the idea of certification. For example, I can now do certain kinds of things in universities because I “received” the PhD degree in 1983. That was 26 years ago. What on earth does that have to do with my teaching now? Well, we are able to establish this idea that there was a change for this individual. This individual now has a feature. This person is “educated”. This person is “qualified”. Various kinds of social ontologies for personal identity become possible because of this big-time compression for personal identity that we perform through conceptual integration.

You remember things from the past. I want to challenge you here, because when you remember something from the past, it seems as if it comes with its own emotion. So you remember something embarrassing from your past; you did something years ago that was embarrassing. This is a problem. Just think about it. No theory of cognition imagines that memories come with their own separate emotional system, that I have emotional system right here active in my brain right now that’s running this scene in this lecture hall, and when I have the memory, the memory brings in an entirely separate emotional system. No. What I must do to feel emotion about that memory is activate certain things in my brain right here, right now. I must use my *present* emotional system in order to have emotions accompanying that memory. To repeat: I must use my *present* emotional and affective system, the one I’ve got right now, and blend it with that past stuff so that my *current* emotions which I feel can provide the blend with emotional content. I can in fact feel the embarrassment right now right here in remembering something from the past. So this now makes it possible for us to have a conception of a past self by blending some things that you take as having reference in the past with some of your current human scale emotional abilities. Even though I’m not at all embarrassed right here right now, I can feel the embarrassment right now in the blend for the action in the past.

This blending is not deterministic. I can do the blending in a number of different ways. Suppose, for example, that the memory comes up, and you remember you were embarrassed by the event at the time. But suppose that in the here and now, you think the only reason you were embarrassed was peer pressure, and now you feel *proud* of what you did in the past, not embarrassed. So in the blend, you have the past event, and the knowledge that it was judged to be embarrassing, but you have the

emotion of pride in the blend, not embarrassment. Alternatively, you can activate more and more knowledge, and somebody can prompt you to remember certain features, and you can bring to bear your emotional system to try to recreate the feeling now that you believe, intellectually, that you had back then.

What I'm saying is that we have a conception of self running over time, and it is possible to have this conception only because we have the ability to do really creative, advanced blending of self and memory to create in the blend a sense of past self connected to ourselves. This is a big deal.

Look at human memory. I'm going to say something now that's going to be pretty pushy. We invent time machines like the Tardis and so on in imagination. And one of the original time machines was conceived by H. G. Wells and presented in a work of fiction, a novella, titled *The Time Machine*. What was Wells' time machine? It was a handle. If the narrator pushed it this way, he went into the past. But notice he had to go through time into the past successively: he had to go back through the day, then back through the night before, and so on through day and night and day and night and . . . This took a lot of "time". If he pushed the handle over hard, the days and the nights changed really fast.

He could also move into the future, by moving the handle in the other direction. Here, of course, we have the usual blend of change of time with movement through a dimension. It is a very standard blend, of time and space. We'll talk about that blend this afternoon when we talk about metaphor.

But notice that human memory does not actually work like H. G. Wells's Time Machine. In human memory you don't have to rewind the tape in order to remember yesterday. What did you do yesterday in the morning? Boom! There it is. You just drop in. Now this is really amazing. Why do we have minds that do what George Eliot writes here:

"Our moods are apt to bring with them images which succeed each other like the magic lantern pictures of a doze; and in certain states of dull forlornness Dorothea all her life continued to see the vastness of St. Peter's, the huge bronze canopy, the excited intention in the attitudes and garments of the prophets and evangelists in the mosaics above, and the red drapery which was being hung for Christmas spreading itself everywhere like a disease of the retina."

Why does memory work like this? The hypothesis I put forward with Gilles Fauconnier is: human memory works like this because it feeds blending. If the only things you have available as inputs to blending are what are in your current experience, that's very limited. But if memory can work so as to activate things that are very far out of your current experience, suddenly you have more material for your blending mill. You have a lot more input spaces that you can use in constructing blends to think about the present, to think about other kinds of things. Human memory has this goofy feature—you can just drop in on anything; it provides wonderful material for blending.

We have conceptual integration networks for understanding memory as working this way. Macintosh operating systems now provide a utility called Time Machine. Time Machine does incremental backups of everything on your hard disk. How it does this—that is, the way a software specialist understands its operation—is one thing.

The everyday understanding of the Macintosh Time Machine is different—it's a human-scale blend. When you enter the Mac Time Machine, that is, when you call up Time Machine and see its visual presentation on your computer screen, you are looking at a window containing the files of the folder you were looking at before you entered Time Machine. In fact, it is the identical window you were looking at before you entered Time Machine. But once you are in Time Machine, you see a stack of windows behind the present one, all presenting the same folder, but at different moments in the past. You see the list of your files right now, but also the folder and its contents at different times in the past. You can go “back” to a previous “version” of the folder, and you can restore what you had in the past. When you look at the screen, it seems as if these “versions” of the folder go all the way back to the beginning of time, or, anyway, to the beginning of your hard disk. Just as in human memory, you can just go back and drop in and activate the previous versions. This is not the way the software actually works, but it is the way we understand the “backup” and “retrieve” mechanism. You can just drop in on the past and activate it and pull it into the present. Nifty.

It's not just the past that we think of as available for activation in the present. We also activate a conception of our future selves. Our future selves are not even here. How can you have a conception of a future self? How can you think about what you are going to be like when you are 80 years old? Well, you have frames for those kinds of age experiences and you can integrate it with your current emotional affective reasoning abilities, and in the blend now you have a future self. You can daydream, or plan, about your future self. The result is an extended self, a big idea, now a big idea of self, and that big idea of the self has an effect on the present self. For example, if you think of a future self that brings tension to the present, you can try to block it out. A future self deathly ill is easily imagined, but the imagining might effect the present self in a way you don't want, so you might try to block out the blend of the future self. We are not always successful in blocking these blends from coming to mind. You don't want to feel anxious at some moments, and so, at those moments, you might try to block the construction of a blend whose inputs are your present self and future conditions, yielding a future self in those future conditions. If you are competing, you can imagine that the future self is a loser, but that conception might degrade your ability to compete, because it induces present feelings, present anxieties. You might instead try to manage your present self by blending in something that would have quite a different effect on your mental state. You might imagine yourself winning, for example.

We also use this kind of blending to construct conceptions of other minds. The conception of another mind is also a big idea, because it takes us very far out of our own experience. We do not have experience of other people's mental states, but by blending, we can imagine them. When we perceive each other, all we see are appearances, behavior. That is all. But we know about our own mental states and we can blend our conception of those mental states with what we see of other people. There are analogies and disanalogies between you and me. I can make a blend, so that

in the blend, you have a mind like mine, and there can be emergent structure in the blend, based on behavior that I perceive. I can accordingly construct slight differences for your mind from mine. Again this seems like nothing to us, but there's no evidence that any other species can do this at anything close to our level.

There is a brilliant researcher named Michael Tomasello, who talks about theory of mind, about being able to conceive of other people and how they operate. And there are evolutionary psychologists who talk about this amazing ability. At least it is recognized in science that explaining this ability to imagine that there is another mind like ours is a huge scientific problem that needs an answer.

There are different hypotheses that have been put forward. One is that you have Swiss Army Knife brains—you have capacities, each a kind of module, none of them related to the others, and one of them gives you social cognition. This is not so much an explanation as a placeholder: you have the ability, but we don't know what goes into that ability.

The view I am proposing is quite different. Human beings have the capacity for double-scope blending, which works according to constitutive and governing principles, and we use this ability to give us the kinds of conceptions that constitute advanced social cognition. Double-scope blending, used throughout human higher-order cognition, gives us the ability to imagine fully other minds. It is a general mental operation—in the sense that it works across conceptual domains—and it provides us with conceptions of other minds.

We use it to project mind not only to other people, but also to animals, ships, buildings, cartoon characters. You can attribute intentionality to a house. We do it all the time. Our projection of mind is not nearly as restricted as proposals for a modular “theory of mind” would have us believe.

It could be that early on in life, you use your double-scope blending ability to come up with a template of what other mind is, but you don't have to rebuild that from scratch every time. You develop it for the most part just the way you develop something like the caused-motion construction in English, for the most part just the way you develop conceptions of time and space. Once you develop these blends anchoring conceptual integration networks, you do not need to remanufacture them every time from scratch. You can activate them partially formed.

It is also possible that there has been some genetic assimilation of the ability to project mind, so that you enjoy a genetically-supported head start in constructing such networks for other minds. Tomasello, in *The Cultural Origins of Human Cognition*, proposes that this ability to have social cognition is the principal ability that distinguishes us from other primates and that makes it possible for us to have culture. The argument I present here is like his argument, except that where he says “evolution of social cognition,” I say, “evolution of double-scope blending”. In my view, social cognition is crucial, but it is a *product* of double-scope blending. Social cognition is a sub-case of double-scope blending. He is right about the importance of social cognition, but I view it as only one of the important things that we are able to do, all of them consequences of double-scope blending.

Blending is adaptive in my theory in just the way that social cognition is adaptive in Tomasello's theory. It makes it possible for us to have culture, including the power of culture to support and maintain a conception, idea, or behavior distributively across organisms, once it has been invented. I have discussed this with Tomasello, who says in response that as long as blending theory includes the necessity for social cognition, he's not unhappy, because the social cognition part is what is needed for what he regards as an adaptive story. He writes, "My particular claim is that in the cognitive realm the biological inheritance of humans is very much like that of other primates." He's right about that, but I think that's a misleading statement. Because what he's going to say is "there is just one major difference, and that is the fact that human beings 'identify' with conspecifics" — that is, with other members of your species — "more deeply than do other primates". That's certainly something we do, but I assert that it's not the only thing we do and it's just one example of double-scope blending. It's one sub-case. That ability for social cognition is quite inadequate to account for the details of the invention of complex numbers, hyperbolic geometry, money, grammatical constructions, counterfactual thought, and so on.

Human beings do many remarkable things. They are very creative at integrating things, not just in conceiving of other minds. One of the things they use blending for is to conceive about the minds. So Tomasello is right that is crucial to conceive of these other minds. But I move the level of explanation one step back. Social cognition is really crucial but it is not the root difference. The root difference is double-scope blending, which makes it possible to have higher-order social cognition.

I can assert from my own experience that we project mind to rattlesnakes. I know, intellectually, that a rattlesnake is not evil, and I know that it doesn't really want to be mean to me, but I once—I'm from Southern California—I once woke up and there was a rattlesnake, and I was freaked out and I wanted to kill the thing. I wanted to kill the thing not just to preserve myself. I wanted to kill the thing because it was mean and evil and trying to scare me and it certainly did scare me. It scared me so directly and effectively that I leapt in one movement bodily to my feet, which was perhaps the wrong action to take, on reflection. Of course I know the rattlesnake doesn't have the goal of scaring me, much less of being evil and mean to me. But in the blend, I attributed that kind of intentionality to the rattlesnake.

We can project mind to an old shirt we have. We wear a shirt. We liked it a great deal. It's old. We even like it when it is old, maybe all the more so. But finally, we wear it completely out, and we throw it away. But you know that it is possible in a moment like that to feel a little bad for the shirt. Does it feel bad being thrown away?

I have had a little robot in my house for the last three years. It sweeps the floor. It's called Roomba. Roomba goes around at 3 o'clock in the morning every night and sweeps the floor, vacuums the floor. My children love Rumba. They talk to it. They give it a name. They know it's a robot. But they talk to it. They say, "Come on, Roomba". They talk about feeding Roomba when they take dirt and put it in Roomba's path. At

one point, Roomba stopped working. I tried various expedients to restore Roomba's functionality. At one point, I shook Roomba and pushed it around a little as it was failing to do its job, so as to indicate to its internal sensors that things were working and that it should just proceed. This intervention worked splendidly, and Roomb's sensors took just the hint I was trying to send. I was probably able to do this because I have worked on some robots in my life. But as I was doing this, my children said, in mock horror, but still with an element of protest—"You're being mean to Rumba!" They were a little surprised to see their father treating anything this way, especially something small that he obviously valued and that had brought the entire family so much benefit. Of course they know that Roomba is just a machine, but this kind of projection of mind, this blending, is something we do all the time. We do it for non-animate objects. We do it for pencils, musical instruments, cars. There is a poem by Thomas M. Disch that expresses just this kind of projection of mind:

Lives there a man with soul so dead
He's never to his toaster said:
"You are my friend; I see in you
An object sturdy, staunch, and true;
A fellow mettlesome and trim;
A brightness that the years can't dim."
Then let us praise the brave appliance
In which we place this just reliance.
And offer it with each fresh slice
Such words of friendship and advice
As "How are things with you tonight?"
Or "Not too dark but not too light."

We do this projection of mind with each other, but in that case we believe in the blend.

I'll end today's lecture with a long list of extraordinary things that human beings can do as a result of having advanced conceptual integration. I won't elaborate on them now. Because of advanced conceptual blending, we are able to have a highly developed and advanced sense of personal identity, to think with feeling about events that are not here in our present scale, to have a sense of the past and the future and ourselves in them, to conceive of other minds. We are able to manage our future selves.

Think of people, for instance, who possess credit cards but cut some of them up because they do not want their future self to get weak and use them. We often do things now because we think that our future self will be different, that our future self will be excited or undisciplined or drunk or something like that, and that we had better do something now to control that future self. We are able to do this kind of self-management. We are able to think of counterfactual scenes, such as the scene in which the world's fisheries are not fishing as much as they in fact actually are. We are able to put together forms and meanings and the little tools we can use to prompt other people to construct various conceptual networks. We are able to have political ontology

and its laws, social memory, concepts like punishment, redemption, guilt, sin, and money—which are concepts that run over long stretches of time and take elements that are very separated in time and space and compress them in the blend. Let me give one example of such a compression. Suppose someone does something, say 20 years ago. Suppose we do something bad to that person now. We have a way of doing that so that we do not have to say, “Oh, that’s one bad thing back then and another bad thing right now.” We say “No. No. This isn’t a bad thing we are doing. This is *punishment*.” Now in the blend there’s an action and a punishment. That’s emergent structure.

Or suppose someone does something bad in the past. Suppose they failed at something. But now, in the present, they do something good, something that is in various ways like the past thing, except that where they failed before, now they succeed. There is a way of looking at this so that we do not count it just as one bad thing and one good thing. In this way of looking at things, the second act is an act of *redemption*. The redemption overcomes the failure that attached to the personal identity from the failure. There are many stories in the world, in movies, and in legends about redemption. Concepts like *redemption* are extremely, extremely creative. They are possible only because we are able to hold on to these long-range conceptions by having the kinds of compressions that can anchor big-time ideas.

This afternoon I am going to talk very specifically about a particular conceptual metaphor that we have. It goes over time and space. I will go through what is really happening in this conceptual metaphor, to show the kinds of compressions it requires, and also to give an idea of how I think advanced metaphor analysis needs to operate. So, as I said, today you will have “bookend” lectures: you began the morning with Dirk’s brilliant analysis of the role of historical diachronic linguistics in social change and metaphor and how this attention to diachrony needs to be incorporated into metaphor analysis. This, for our last lecture, I am going to give an analysis of the nature of mental operations in metaphor. The metaphor I have picked for this afternoon is time is space, because that is a metaphor that gives us some very big ideas. Thank you very much!

Lecture Six

Working in the Mental Network

Thank you very much for that wonderful introduction. My mother would have loved it. Thank you for welcoming me to the beautiful Beijing Language and Culture University. It is a great honor to be here and a great pleasure. I have been looking forward to it and I have already had a wonderful time talking to the students here at your university.

Some of you are new to the crowd. And I will give just a brief introduction to a couple of things that I put in place earlier. I've been talking about the way in which human beings are exceptionally creative — how they are distinguished from other species in having an advanced form of the ability to integrate mental structures, conceptual structures; and how this ability produces the kind of innovation that we find uniquely in our species; how it makes culture possible and how it makes language possible. This is a good topic therefore for the Beijing Language and Culture University! I will give just a little taste of this operation.

This mental operation is called conceptual integration. Conceptual integration occurs when you have two different mental arrays that you are able to integrate selectively into one mental array with emergent properties. The inputs and the blend form a conceptual integration network. Here in the slideshow, you see a yellow circle. Suppose that it represents a man's thought about the fact that he is participating in a wedding. He is a groomsman and he is helping with the wedding. And so in this mental space, there is the bride, there is the groom, and there is the official, and they are on the grass overlooking the Pacific Ocean, and he is doing his job. He is in this wedding frame. He knows his role.

But at the same time that he is participating in the wedding, quite amazingly, he is also thinking about the fact that three weeks ago he was diving with his girl friend off Cabo San Lucas, down in Mexico. And she's not here now. They were diving for treasure. This blue space represents his mental space about diving with his girlfriend off Cabo San Lucas.

Now it is already an amazing problem for cognitive science that he would be able to think of these two things simultaneously. Shouldn't he get confused? What use is it to him to be thinking about diving with his girlfriend in Cabo San Lucas? But notice that he doesn't get confused. He does not swim down the wedding aisle even though in his memory he's swimming. He does not have difficulty speaking even though in his memory he's got an apparatus in his mouth. He does not mistake the bride for a shark, and he doesn't think his girlfriend is actually here.

This is such a big problem that cognitive psychologists like Arthur Glenberg and cognitive scientists like Terry Deacon have dedicated their research to trying to figure out why we should be able to do this, why we should be able to activate two conflicting mental stories at the same time.

For the moment, let's just take this part of our mental operation for granted. The next thing that the man at the wedding can do is make mental conceptual

connections between these two mental spaces, the mental space of the wedding and the mental space of the diving. For instance, he can construct an analogy connection between his girlfriend and the bride.

But he can do more than just connect these two spaces. He can project selective structure from each of these two mental spaces—and this is the important part—into a third, blended mental space, represented by the green circle in the slideshow. In the blended mental space, he can be running a simulation in which he is marrying his girlfriend right here. He is imagining this.

Notice, marriage to the girlfriend is not in the space with the wedding and it is not in the space with the diving. In the slideshow, it is not in the yellow or the blue spaces. But it is in the blend, the green space. Marrying the girlfriend is *emergent structure* in the blend. It is new. In fact, the groomsman at the wedding may have never even thought of it before, but he can project selectively from these two input spaces, and invent emergent structure in the blend. And once he has, he can think about whether or not marrying the girlfriend is a good idea.

We have talked already quite a bit about the way in which blending makes it possible for language to use a relatively few forms to talk about almost anything. So here is an example that I have not provided in previous lectures. In English, there is a clausal structure called “ditransitive”. You all know all about about the transfer of something. If I pick up something and I hand it to somebody, that’s a very basic mental scene. And she hands it back to me, and I hand it back to her again. You should watch little children when they are playing, because they are willing to hand the object over and hand it back, repeatedly. “Would you like it?” and then they take it, and “May I have it?” and then they give it back to you. This is very human scale scene. *Hand*, for example, is a denominal verb for this action in English. *Give* is also rooted in this basic human scene.

But there is also a clausal structure that is associated with the ditransitive — noun phrase, verb phrase, noun phrase, noun phrase. So I say “*I handed her the eraser*”—noun phrase, verb, noun phrase, noun phrase. *I*, noun phrase; *hand*, verb phrase; *her*, noun phrase; *the eraser*, noun phrase. Now that’s a very basic human scene, very compact. We understand it very well. But I can take all kinds of things that we do not understand very well, that are very diffuse, and integrate them with that scene into a blend that now has ditransitive structure.

Suppose you have an agent who acts. And somewhere, related to this action, is a causal event. Perhaps somewhere in the diffuse structure, there is a recipient. There is a patient who can receive. And so on. Adele Goldberg points out that we can now use the ditransitive frame attached to this clause to say *She granted him his wish*—noun phrase, verb phrase, noun phrase, noun phrase. We can do this even though no actual physical transfer is happening.

She gave him that premise in an argument. She allowed him that privilege. She won him a prize. She bequeathed him a farm. These are all cases where we take a range of diffuse structure and blend it with the ditransitive frame to produce a blend with

ditransitive structure even though no actual physical “handing” occurs.

Consider *She bequeathed him a farm*. Think about the extent of time and space that this conception runs over. Think of how many agents are involved. *Bequeath* means I leave it in a will and you get it after I die. That’s a long time. But down here in the blend, bequeathing is like handing you something right here. In using the ditransitive clause, I am asking you to use the ditransitive conceptual frame as an input to the blend. The ditransitive syntax invites you to consider using the ditransitive conceptual frame as an input to a blend meant to anchor a perhaps diffuse network.

Consider *She gave him a headache*. Notice she can give him a headache even though she doesn’t have one. In the ditransitive scene, if I give her the eraser, it’s because I have one. But now under selective projection I can ask you to use the ditransitive, the idea of giving, selectively. In the blend, you may not include the idea that I possessed the object before I gave it to you. We might project to the blend only the reception.

She gave him a headache. She showed me the view. She told me a story. She denied him the job. Notice that you can use the ditransitive clause with a verb like *denied*, which indicated stoppage. The ditransitive clause involves transfer and reception. But we can place a verb in it like *deny* or *refuse*, which is opposed to the notion of transfer and reception.

But that is no problem, because you can put them together into a scene in which “she” did something that had a causal effect on the transfer of something to “him” and what she did was stop it. In the blend, now, the transfer is not completed. Blending is selective.

We must stop here and make an observation. If we were to take to mental spaces and, as an exercise in combinatorics, compute all the possible combinations of elements from those mental spaces, in all the possible ways, we would find that almost all of those possibilities are ignored by human beings who do blending. That is, very, very few of the formally possible combinations are ever achieved. This is because blending is very tightly constrained, by the constitutive principles and the governing principles. But given our human blindness to what we do not construct, it may be difficult for us to see the sharp limitations on blending at work unless we think hard about it. Consider “She denied him the job.” I have just said that the ditransitive frame includes successful transfer and reception, but “denied” indicates the opposite, and that blending can put them together in a very effective way, namely, there is a scene of transfer and reception, a hypothetical or possible one, and “deny” prompts us to construct a blend in which the possible transfer and reception are not achieved because of an action by an agent. You might say, “I see, blending can combine things that are contradictory, so it can combine anything, so it is unlimited and unconstrained, unfettered, and so it means anything can be anything, and that does not tell us much.” Not so fast. Blending can indeed combine things that are contradictory—this is the essence of double-scope blending—but not just in any way at all! It is very constrained in how it does this. Most of the combinations of contradictories are never achieved by human blending. Blending is tightly limited, tightly constrained. The root of science is that it must preserve phenomena: human beings do indeed combine contraries to good effect,

inventively, and successfully. Our theories of cognition and language must embrace that indisputable fact. A theory that does not have room for the effective blending of contraries is to be dismissed out of hand as unscientific, plainly untrue to the data.

Only some things come down from the input spaces. There is emergent structure in the blend.

One of the great things that blending does for us is create *compressions*. Blending takes things that are very diffuse, which we could not understand without blending, because we do not have minds built to wrap around those kinds of structures, except through blending. We work at human scale. Blending takes a diffuse network and creates an anchor, a blend, that is at human scale. It creates a human-scale compression that we can handle and that lets us hold on to the big conceptual network. We have a human scale compression or platform that conceptually we do understand and we can work from there.

How do we achieve these compressions? One way to get a human-scale blend is to use as an input something that is already compressed, something that is already at human scale. This is called “borrowing a compression.” Notice that “borrowing a compression” often borrows grammatical constructions. This is what you just saw in the case of the ditransitive. The ditransitive of handing something to somebody is already compressed. It’s already at human scale and it comes with a linguistic construction — noun phrase, verb phrase, noun phrase, noun phrase. Down into the blend comes that compression and the construction with it.

Here are some examples. I say “*You are digging your own financial grave*”. In the space of digging a grave, there is one agent with one instrument who uses that instrument in the same way again and again—through repetitive action—to produce a result. The scene is at human scale. Now imagine quite a different conceptual scene. In this scene, a conservative father is saying to his adult son, who is investing in what the father thinks are unwise financial instruments. The father thinks these investments will lead to bankruptcy. The father says, “*Son, you are digging your own financial grave*.” In the mental space of financial investing, there is a huge financial problem.

In the space of digging the grave, the reason that a grave is dug is because somebody dies. But over in the space of the financial difficulties, the reason you go bankrupt is because you invested. In the space of digging the grave, it is not the person who is digging the grave who dies. But in the mental space of financial difficulties, it is the person who is doing the investing who goes bankrupt.

So there is a big clash between these two mental spaces in causality. The causality is inverse in the two scenes.

In the space of digging the grave, people do not dig graves without knowing that they are digging graves. But the entire reason for saying, “*You are digging your own financial grave*” to the son is that the son does not realize that he is digging his own grave. He doesn’t realize that he is going to go bankrupt. He does not understand what he is doing.

In the space of digging the grave, the participant structure is that the person who digs the grave is not the person who dies, but over in the financial scene the person who does the investing goes bankrupt. So you see this is an enormous clash. When we

blend the mental space for digging the grave with the mental space for financial difficulty, we take parts of the frames for each space and project them into frame-level structure for the blend. That makes this blend a double-scope blend.. When I say that there is a clash between the frame structures of the input spaces, I mean that the top organizing conceptual structure of the two separate mental spaces are incompatible, and that parts of each of those frames comes into the blend to build the frame structure of the blend. This clash does not stop us from putting together a coherent human-scale blend. On the contrary, this kind of blending is what we are specialized to do.

Luckily, we have already heard from Professor Dirk Geeraerts about Clinton's difficulties with Monica Lewinsky. When Clinton was first starting to have these difficulties, he just brushed them off. There was no problem. There was no scandal for him. I was living in Washington D.C. at the time and somebody in the Washington D.C. area said something that was printed in the newspaper: "*If Clinton were the Titanic, the iceberg would sink.*" The movie *Titanic* was new at that time. Newspapers often do this—they blend things that are currently, but independently, active. Two things come up and you blend them. You are very good at doing this. Notice you understand this conditional sentence well, if you have the cultural frames.

But here's something crucial: the one thing you know about the Titanic is that it sank. But in the blend it doesn't sink! The one crucial thing you know about an iceberg is that ice is less dense than water and it can't sink. It can be submerged but it can't sink.

But in the blend, we now have new kind of physics: ice can sink and the reason it can sink is because Clinton is so powerful. The causality comes from the target, Clinton and his difficulties, not from the source, which has the Titanic and the iceberg.

You have often heard it said about metaphors that the causal structure and the inferential structure come from the source. That is in principle false. That is not the only way in which metaphors can be put together. Compression is a bigger reason for metaphor than inference. These expressions are clearly metaphoric or felt to be metaphoric. It is quite clear that digging the grave is the source for the metaphor in "You are digging your own financial grave" and that the Titanic and the iceberg are the source for the metaphor in "If Clinton were the Titanic, the iceberg would sink", but it is not the case that the main inferential structure, the causal structure, participant structure, intentionality, and so on necessarily come from the source. No, they often come from the target by selective projection.

The point is that we borrow a compression and it brings with it grammatical constructions. The result is human-scale blends with attached grammar. The blends help us understand the network and the grammar gives us a way to refer to the blend and hence to the network it anchors.

There is another way to create compression in the blend. We have seen this way before in these lectures. In this way of creating compression, it is not the case that the compression is already in one of the spaces. Rather, there are "outer-space" connections between input mental spaces, and these outer-space connections are compressed to human scale into the blend. Also, fueled by this compression, other

vital relations are added to the blend, or strengthened in the blend. We have seen this kind of compression, for example in *Dinosaurs turned into birds*. In the input mental spaces, here is one dinosaur, and another dinosaur, all these dinosaurs, and some of them die off and they are all a little different. You have the analogies and disanalogies across all these inputs. And the analogies get packed to a uniqueness in the blend—the dinosaurs—and the disanalogies get packed into change for that unique element — *dinosaurs turned into birds*.

Similarly, in *A fed bear is a dead bear*, we compress outer-space vital relations between mental spaces. The same happens in *Make this envelope disappear*. Or *my tax bill gets bigger every year*. Or *my electric bill gets bigger every year*. In this case, you have electric bills and there are analogies and disanalogies across them. They get compressed to one tax bill and it changes. It gets *bigger*. Now, no tax bill ever got bigger, but in the blend, you have one and it gets bigger. You are not deluded by this human-scale blend. It anchors a diffuse conceptual network. In the blend, many become one.

I am moving a little fast here, but I want to point out that compression of outer-space compression is the operation behind most negatives. Imagine that there is a table and that there are these five chairs around it. And I say “*Put the green tea in front of the missing chair*”. What is the missing chair? What does that mean? Here is what it means. You may have seen the table before with six chairs or you may have the idea that chairs should be equally spaced around the table. And so you have that idea, you have that mental space, and you notice there is a difference between these two mental spaces—the one with five chairs and the one with six. Since one has five chairs and one has six, there is a disanalogy to go along with the robust analogy. We compress these two different tables to one table, and we bring in the sixth chair from our concept of the table with equally spaced chairs. We bring it in. It is in the blend. It now has a feature and that feature is that it is absent. It is like this table is wood, this chair is metal, and this chair over here is missing. This particular chair has a feature in the blend, but the feature is not *wooden* or *metal*; rather, the feature is *missing*. We are very used to things having features. And you can see, conceptually, the missing chair right there when you look at the table with only five chairs. The sixth chair is right there. There it is. There’s the missing chair. You know where the missing chair is. *Put it in front of the missing chair*. *Missing* is now a feature and you already have grammar for expressing properties of things. You use adjective plus noun. So you don’t need new language to refer to the blend. You do not need new grammatical structure. You already have adjective plus noun. In all these cases, you make a compression by compressing cross-space connections between the inputs. This is not unusual.

In Berkeley, where received all of my college schooling, *a caffeine headache* is the headache you develop because you didn’t have your morning coffee. So here, in the slideshow, is one scene in which you have a headache, and some kind of cause, and over here, in a different mental space, you have coffee. There is a disanalogy between these two analogous spaces, namely, over in one space there is a headache, but over in

the other space, there is coffee. You blend these two spaces, and now in the blend, the disanalogy is compressed as absence of caffeine. The cause is now the absence of caffeine. That's why you have a headache. And of course we can refer to absence of caffeine by the word "caffeine". "Caffeine" here is perfectly fine for *absence of caffeine*. Words don't mean. Words are tools that we use to prompt other people to put together a conceptual integration network that is meaningful. The meaningful blend here is that absence of caffeine has caused my headache and I can prompt you to think of that by saying *caffeine headache*. That is not unusual. *Money problem* is not usually a problem that I have because I have too much money. *Money problem* is usually absence of money. *Nicotine fit* is not a fit because you have so much nicotine from cigarettes. It is a fit because you haven't had a cigarette. *Security problem*: that is not used to indicate, "Boy, I've got a problem. There's too much security." No. *Arousal problem*. An *insulin coma* isn't a coma that comes about because you have insulin. It's a coma that comes about because you don't have insulin. *Insulin death, food emergency, honesty crisis, rice famine*. This is not unusual.

Compression of the disanalogy into an absence in the blend where we can suggest the absence by the noun associated with the element.

Safe beach, what's that? What is *safe*? Well, suppose the child is going out on the beach. And you say, "Oh, I'm nervous about my child" and I say, "it's OK, the child is safe". Or I say "It's OK, the beach is safe." Now I can mean exactly the same thing by those two sentences. If the meaning of a sentence was the composition of the meaning of the words, and words had meanings, and "safe" had a meaning, and the meaning was to predicate a feature of the noun, then *the beach is safe* and *the child is safe* should mean two very different meanings, but they don't, or at least they don't have to. They can mean the same thing.

There is the child going out on the beach. Now, there is a hypothetical space in which the child is harmed. And, when I say, "the child is safe," I am prompting you to make a counterfactual connection between the scene we are in and the imaginary scene in which the child gets harmed and compress those. And now in this scene it is not just the scene of the child going out there. Now in the blend the child goes out there and there is absence of harm and now I want you just to conceive of this space as having absence of harm and I do that by using the word *safe*.

I could say, instead, "*The beach is safe*." I could mean that real estate developers will not come ruin the beach. I could mean many things. But the point is a word like *safe* or *missing* or *gap*—as in *there is a gap in the fence*—, these are specialized words that say "Hi, I'm using this word here, and what I want you to do is construct an integration network in which there is a disanalogy across certain input spaces that gets compressed into an absence down here in the blend".

Look at *There is a gap in the fence*. In your mind's eye, you can see the gap in the fence. It is right there. And what that means is there is a mental space in which there is a fence, and another mental space in which there is a fence, and that there is a disanalogy between them because in one space there is a board in this position in the

fence and in the other space there is no counterpart for that board. And I want you to understand that in this scene you are supposed to construe this space as a board that is not there. So there is a *gap* in the fence.

We talk about *a boat house, a house boat, an angry man, a loud man*. Notice *the man is loud*. What we really mean in saying this is that there is a scene in which there is a man with a voice and the voice is loud. But now the property of the voice becomes a character property of the man.

Consider *sugar-free* and *dolphin-safe fishing*. In the United States, we have tuna cans that are labeled *dolphin-safe tuna*. This phrase is not used to indicate that it is safe for the dolphin to eat the tuna, that the dolphin will not be poisoned by the tuna. Rather, it means that this tuna was caught in such a way that the nets did not damage the dolphins. It's dolphin-safe tuna.

Likely candidate can mean not a candidate but somebody who is likely to be a candidate. In *guilty pleasures*, it's not the pleasures that are guilty, but you have certain feelings about the pleasures, and this structure is compressed so that the pleasures now have a feature—*guilty*. These are all cases of the compression of outer-space relations to inner-space features.

Let me give you an example here that I heard on the radio. In the United States, we often have “three-day” weekends. One of these is because of “Memorial Day.” It's like Dragon Boat Festival, a three-day weekend. I saw an advertisement that read, “*At South Shore Lumber, get no sales tax Friday, Saturday, Sunday, and Monday!*” So that is four days, because they are extending this benefit to the day before the three-day weekend.

Sales tax in the United States works like this: you buy something, and it has a price, but then, in addition to the price, there is another amount you pay which is the tax and that goes to the government. It does not go to the federal government; it goes to the state government.

States are sovereign powers in the United States. That's why it's called the *United States*. Notice what is happening here. What does this phrase mean, “*get no sales tax*”? Here is how this works. There is one scene in which the buyer pays the price and gets the lumber, gets the goods. There is another scene in which the buyer pays the price and also pays an additional amount — the sales tax—and gets the good. There is an analogy between these spaces, of course, but there is also a disanalogy. We compress the disanalogy into an absence in the blend. Now, someone can give you the absence. You get not only the good but you get *no sales tax*. The result is a human-scale scene, in which you receive two things: the goods and *no sales tax*. Now of course what happens is you paid only the price, and you received only the goods, but in the blend, the absence is a feature, and you receive that, too.

We perform this kind of compression with personal identity. I give you the example before of a nine year old boy saying “*If we were all chickens, you, William, would be about Elizabeth's age*”— that's the babysitter— “*you, Peyton,*”—who was seven years old at the time—*would be about Dad's age, and me, Dad, and Mom would all be dead of old age. We are all five alive. We are lucky we are not chickens.*” Now I won't go through all the blending of ages and how it is that Peyton can be Dad's age

but Dad is dead because Dad is Dad's age. You can figure that all out. What I want to point out is the amazing compression involved in the concept *lucky*. We were just sitting at the dinner table. That's what was going on. There was no *lucky* in the scene at the dinner table. But then this blend came up, and in the blend, blend Mom, Dad, and Jack are dead and Peyton and William are old. And the scene of our sitting at dinner is now counterfactual with respect to that blend, and so now we are supposed to understand that in our scene we have a special feature. This feature is a compression of a counterfactual relation. The counterfactual relation is compressed into a feature: we are now *lucky*.

We have many words like "lucky". Consider *accident, dent, mistake*. We say, *This is a mistake*. All of these little words are prompts we use to invite someone to construct an integration network that has a compression of a disanalogy. All of them invite us to understand a scene in a certain way.

Here is another example, uttered by a child, who was looking up at the stars. He had just learned about constellations, stars, and didn't understand that the constellations are already set, that you don't just make up new ones. So the child said "*Oh, yes, I see a constellation. It is Bang Shot.*" Bang is the noise that a cannon makes or a gun makes — bang. "*It's a man*", he says, "*He was a soldier in the Revolutionary War. All of the members of his troop were killed, and the British were starting to come in his direction.*" The American colonies were fighting the British in the Revolutionary War. "*So he hid inside the barrel of a cannon. But then reinforcements arrived, his own people, and fired off the cannons. Now, this cannon was not lowered, so it fired him into the sky, and there he is. But he doesn't know he's dead, so he keeps running across the sky, trying to find a way back into the Revolutionary War.*"

There are the stars. They move. They are in a certain shape. This child has learned that constellations have a story. So he comes up with a story, of a guy, and, bang, there it is, Bang Shot.

Notice, everyone, that you understand this story. The story does not present a possible or plausible world. It's not possible for people to be dead and yet to be operating in a certain way because they don't know they are dead. That's not possible. But we do not have any trouble understanding it because we are putting this together for the purpose of thinking about the sky and other kinds of things.

We have been talking about blends. Many of these blends feel metaphoric. For the rest of this lecture, I would like to look at how blending analyses requires a revision of metaphor theory. Let's look at a particular metaphor that we all know—time is space

We say "*We are getting closer to Monday*" or "*Monday is approaching us*". We have all studied this metaphor previously. Both the ego-moving and the ego-stationary dual versions, which have been recognized by rhetoricians for at least scores of years, are laid out in *More than Cool Reason*.

There are many, many blending networks that we could talk about that involve time, and give us Big Ideas. Let's just take one of these. There was a catamaran—two men on a catamaran, two sailors. They wanted to sail from San Francisco around Cape

Horn up to Boston, because there was a world record for doing that, set by a clipper ship in 1859. And they are on the catamaran in 1998 sailing a catamaran called Great America II, and the sailing magazine, the news magazine called *Latitude 38*, which is just wonderful, writes, “*As we went to press, Rich Wilson and Bill Biewenga*” — those are the two sailors on the catamaran, “*were barely maintaining a 4.5 day lead over the ghost of the clipper Northern Light.*”

What does this mean? In 1859, the clipper is moving along. In 1998, the catamaran is moving along. There is only one boat in each time. They are very different but they get compressed into one blend that now has both boats making this run. And the boat from 1859 — Northern Light is projected down into the blend, but its full reality is not projected into the blend, so we call it a ghost.

Once we have this blend, we can understand this is a race and we can say things like “*it’s 4.5 days ahead of Northern Light*”. In fact you could just say that without *ghost*. “Ghost” marks that there is an absence in the 1998 input.

Now nobody is deluded. The guys in the catamaran, the sailors, do not think if they fall off in the sea, that 4.5 days later the Ghost of Northern Light will come by and pick them up. They do not think they can be saved by hanging on for 4.5 days. They are not confused about their reality, but they are using the blend in order to understand and use these compressions across time.

Here, in the slideshow, you see an advertisement of what it’s like to be in a business seat on British Airways. This is a time compression, compressing the businessman with the baby that the businessman once was, in his mother’s arms.

Here, in the slideshow, you see a picture of the fall of the American stock market. The title is “*Lovely while it lasts.*” The Dow Jones Industrial Average is blended with a kite, high in the sky. You know it is going to fall. So look at the compression. The time plot of the Dow Jones industrial average becomes the string of a kite and the sky becomes the graph paper.

Let us look at the basic conceptual connections between time and space, usually analyzed as a basic metaphor, TIME IS SPACE. Let me list briefly the conclusions I will be leading up to.

Simple mappings are simplifications, compressions, of the actual conceptual networks. Metaphor analysis must embrace the fact that conceptual products are almost never the result of a single mapping. Metaphor analysis usually has as its product a mapping from one domain to another domain. Such mappings are always a gross simplification of what is actually going on. Conceptual integration networks are much more complicated than that. That is true for metaphors like time is money or time is space.

Cobbling and sculpting. Integration networks are never built entirely on the fly nor are they pre-existing conventional structures. Integration networks underlying thought and action are always a mix. There are individual components. There are cultural components. On the one hand, cultures build networks over long periods of time and they get transmitted over generations. But on the other hand, people are capable of innovating in any particular context. So you have novel mappings, novel compressions that use entrenched compressions. You cobble and sculpt with these


networks when you try to communicate.

Compression. Now this is a remarkable conclusion of our recent work over the last decade and it was overlooked by both early metaphor theory, which I worked in, and by early blending theory. The insight is that integration networks achieve systematic compressions, and there are standard ways of doing this, standard patterns of compression that we use all the time to create human-scale blends. There are compression and decompression patterns for conceptual integration networks. In many ways, they make conceptual metaphors possible.

Inference. Inference transfer is not the only or maybe even the main motivation behind metaphor. In fact, it is quite typical for “source-domain” inferences to be kept out of the emergent blended space. This is because topologies in the multiple inputs may clash, so not everything gets projected to the blended space. On the contrary, the projection is selective and much of what is in the blended space *is not in any of the inputs*, because, as you make the blend, you can run it, develop it. Remember the blend in which the groomsman is imagining marrying his girlfriend. Marriage to the girlfriend was not in any of the input spaces for the groomsman doing the imagining.

Emergent structure. The focus on single mapping and inference transfer in early metaphor theory left out many of the powers of these integration networks, in particular the ability to develop emergent structure. So we are going to see the development of emergent structure in analyzing TIME IS SPACE.

Here we go. Let’s do a reanalysis of the old basic metaphor TIME IS SPACE. Here is a sketch of the conceptual integration network for understanding time that we will go through during the rest of the lecture:



For starters, let's think of our experience of time. It's day; the sun is up. I used this *succession of days* notion in the story of the Buddhist monk. Here, in the slideshow, you see a representation of a progression of such days. After the first day, it's night, and then it's another day, and then it's night, and then it's another day, and then it's night and then it's another day. They just go on and on like that, in a sequence.

It is already a major compression to create a blend for the *cyclic day*. Let's label this mental space **C**, for cyclic. This blend is a compression of all these days into a single cyclic day. In the blend, we have one day, one conceptual day, and it's cyclic. Across all these individual non-cyclic days that are the inputs, there are analogical connections. Those analogical connections are compressed to identity: there is one cyclic day. In the input spaces, sometimes the sun is up, and in all these input spaces, there is a noon, that is, when the sun is directly above. Noon, noon, noon, noon, noon, noon, all of those noons get compressed in the cyclic day and now, in the blend, there is one noon. There is one dusk. There is one dawn. This compression is something like the compression we saw in the Mythic Race example, with Hicham el-Guerrouj, in which all six of the races over the last six decades get compressed into one race. This is a compression. And now there is a cyclic day, **C**. It's a conception you have, one day, and this compression makes it possible to say things like — and this is good only for the blend — *we've come around to noon again*. In the input spaces, you never come around to noon again, and even if you move from one space to another so that on the second day it is also noontime, meaning that the sun is directly overhead, it is not the case outside the blend that you have come *around to noon again*, because

outside the blend, there is no *around*, no *noon again*. Rather, there are two separate noons, and you experienced one and then you experienced the second. In the blend, you've come around to noon again.

Now notice that in your experience of days you are never at the *same time again*. There is an imaginative film called *Groundhog Day* in which the main character does indeed wake up every day to the *same day*. But in real life, no time is ever actually repeated. Every moment in time that you are in is always manifestly different from every other time you were once in, but now, down in the compressed conceptual structure of the blend, you come *back* or *around* to *noon again*. You are at the same spot in the cyclic day. Once we have this blend, we can say *this park closes at dusk*. We can say *it's time for your morning coffee*. Why? Because across all these different mornings you have coffee, or at least you have it typically and now it comes down into the blend as one coffee. It's your morning coffee. So *morning* can now be a *feature* of the *coffee* experience. One can even say, "I won't be having my morning coffee today", or "why is my morning coffee not here?" or "where is my morning coffee?" That's amazing: there is a morning coffee even for a day in which you are not drinking coffee. It's just that on that day, it's absent. *Absence* is now a feature in the blend, compressing a counterfactual link between today and the cyclic day with the morning coffee.

Once we have these cyclic time blends, we can say things like *your weekly workout*, *your monthly visit to your mother*, *your annual checkup*. There are many of these cycles that we create.

Let's stick with **C**, the cyclic day, though. Now we want to discuss another mental space, which I label here in the slideshow **A**, for anchor. It's very useful for us to have a concept of a periodic event, something that repeats. We need an anchor in the world. One of these anchors is the sun's going up and coming down. It's a very standard one. Another one is a watch. A watch is a very bizarre thing, when you think about it. It's a bunch of glass and metal and you stick it on your wrist. It can get broken easily, but it's right out there, at the end of your arm, highly vulnerable. You can't eat it. It doesn't keep you warm. It doesn't hold a drink. What is it good for? It has some little metal that moves around. If there is a physical event that we all think of as repeating—that is, a physical event where we can all perceive the onset and the termination easily, and all agree easily on when the onset and the termination occur—then we can use that repeating physical event and blend it with our concept of cyclic time.

It works like this. Here, in the slideshow, you see your cyclic day, and over here in the slideshow, is a representation of the object with the repeating physical event. So we have these two mental spaces, **C** and **A**.

Noon, for example, in the cyclic day can now correspond to something in the period that is repeated, that we think of as repeated. There is then a connection between mental space **C** and mental space **A**. And you can blend them, to make a blended mental space that I will label **M**. It is quite ingenious what we invented. You can have a watch where the hand goes around just once every day. And this, in the

slideshow, would be noon, when the rod comes around to this point. It's noon when the rod reaches this point. It's midnight when it reaches this point. You could make a watch like this. I have seen clocks that work like that.

We have chosen instead, for various reasons that are entirely explicable but that I won't go into, to make it so that the rod on the analog watch goes around twice. In that arrangement, when the rod reaches this point, it could be either noon or midnight. And we can divide up the face of the clock so that the rod's sweeping out a certain arc corresponds to an interval of time, and that equal arcs correspond to equal intervals of time. So now the arcs are time intervals. If you look at your watch, it's divided up in beautiful ways so that you have twelve divisions of the face. This is an analogue watch, by the way, with twelve divisions of the face, but the little rod goes around twice. So the number of divisions in the day is the number of times the little rod goes around in one day times the number of divisions on the face. If the little rod goes around twice and there are 12 divisions, then there are 24 equal divisions in a day.

You can map this conceptual cyclic day onto this periodic event, and blend them this way, so that in the blend, you have universal events like *hours*, measured by physical events. We have this agreed-upon blend. In this blend, **M**, you do not have time yet, but you do have agreed-upon universal events, like the *hour*. An hour corresponds to a certain division of the cyclic day, and, in the new blend, to the rod's sweeping out a certain arc on the face of the clock. That interval is 1/24 of the day. And all the hours have equal arcs. We all agree; we can all agree. Now this is amazing. Where is the hour? Can you touch it? Can you see it? Suddenly there is a universal event in the world, an abstract universal event. But what does this abstract universal event—the hour, or the day, or the minute—have to do with an actual event. How do we dance for an hour, run for an hour? Where is the *hour*? We'll get back to that. For now, focus on this blend, which gives you universal temporal events. Notice that the actual universal temporal event, like the hour, isn't actually anywhere as opposed to anywhere else. We will come back to mental spaces **C**, **A**, and **M** later.

There is a great range of linguistic data having to do with *time*. We must account for this data. A theory that sets this data aside is not scientific; a fundamental basis of science is to preserve the phenomena, to respect the data. You will find what I am about to say in an article titled "Rethinking Metaphor", by Fauconnier and Turner.

Let us look at some of this data. We can say *three hours went by, and then he had dinner*. We don't say, most of us, **Three feet went by, and he was at the door*. You can say *minutes are quick but hours are slow*. You don't say, though, when you are moving **Inches go faster than feet*. You can say

- *Those three hours went by slowly for me, but the same three hours went by quickly for him.*
- *For me the hours were minutes but for her the minutes were hours.*

When I say these things, you know what I mean.

- *At the end of the three hours, you will have solved the problem, but at the end of the same three hours, he will have solved it and five more.*
- *Time came to a halt.*

- *Sure, it's Friday afternoon, but Monday morning is already staring us in the face.*
- *Next week was an eternity away.*
- *For me, the three hours were forever, but for her, they did not exist.*
- *It'll go by faster if you stop thinking about it.*
- *Our wedding was just yesterday.*
- *Where have all those years disappeared?*
- *Next week was an eternity away.*
- *I didn't see those years go by.*

You can say all of these things and be understood. That's amazing. Science must account for this ability to talk about time in this way. The explanation is not provided by mapping time onto space the way metaphor analysis usually does. Metaphor analysis must do much more than that to account for these data.

Notice some goofy things. In the domain of space, units of measurement are not moving objects. So if I have a foot, a ruler, you know it doesn't move on its own. I do not look at the world and see feet and yards and meters flying by. That is, I don't see the units of measurement flying by. I may see an actual ruler if someone throws it, and if I am racing down the road, I may "see" parts of my scene flying by, and one of those parts may look to me a meter long. But I do not see some universal measurements of space flying by. Units of measurement are not moving objects. But in the data we just discussed, universal temporal events are moving objects! They move. Hours *go by*; days *go by*.

In the domain of space, observers are not at the same location and they are not looking in the same direction. So I am right here in space, and I am looking over here. I am looking at you but you are looking at me. We are not in the same place and we are not looking at the same things.

But in the data having to do with time, we are all at the same temporal location and we are all looking in the same direction. We are also seeing the same universal events.

When we talk about universal temporal events, we assume that everyone is experiencing the same ones: we see the same *hours*; we see the same *days*. What we actually see in our hours or days can be radically different, but if I see an hour go by, then during that interval, you see the same hour go by, even if I am in California and you are in Beijing. We are in the space spot of time, and we are looking at the same events.

In the domain of space, not all moving objects are on the same path. But in our understanding of time, all of these moving objects, the hours, Monday that is coming at us, Tuesday that is coming at us, the Monday we are approaching—all of these universal events are on the same path. All of the events are on the same path.

We have just reviewed some enormous conceptual clashes between time and space. Here are more. In the domain of space, observers in the same location looking at the same direction would see not only the same moving objects but also the same speeds. But in the linguistic data for talking about time, that's not what we see.

For her, Monday is coming rapidly because on Monday she takes the test. But for me, it's coming very slowly because I'm going on vacation Monday and I want it to get here but it won't. In the domain of space, all objects moving along a path exist, and the closer ones are perceived as closer. But when we talk about time, the times that are farther away can seem closer. In the domain of space, you cannot speed up or slow down the speed of moving objects by the quality of your attention. But the linguistic data for talking about time indicate that you can change the speed of the events by the kind of attention you pay them.

How does all of this happen? How can we explain all of this data?

Here is our analysis. I will walk through it on the slideshow.

There is a big conceptual domain of events. We experience them. For example, this lecture is an event. We are here. This is happening. We are experiencing this event. Our understanding of the domain of events is remarkable, a conceptual achievement, but for the most part we take it for granted. We reify it, and feel that it is true to the world. The wind is blowing: that's an event in our understanding.

Human beings are experts at parsing the world into events and objects. That's what we do. We look at the world. We see events and objects. This is pretty bizarre because your visual field is not parsed into events and objects. You don't see events and objects in separate places. All events have objects and objects are involved in events. And moreover, it's not even clear that the objects are parsed. You have to do the parsing into this object versus that object. The world doesn't come with labels. But we are experts at partitioning and labeling. Neurobiologically, that's what we do. We are mammals. So, in these lectures, we will take it as given that we can do this.

Our expertise in this big conceptual domain of events, which I am labeling **E** here in the big conceptual integration network for time, includes understanding event shapes. We understand such aspects of events as their order, and features such as onset, repetition, completion, and so on. Mental spaces can include events, and our subjective experience of the event can be included. So we can have a mental space for an event and an element in the mental space that our experience was pleasant. Or painful. Or boring. Or frightening.

For example, a lecture might be fun for me and boring for you and challenging for the technician who must run the camera and so on. We can all have different subjective experiences of the events.

Now there is another space, **X**. This space has another kind of knowledge. This is a very important kind of event for human beings, namely, motion through physical space from a point A to a point B. We are human beings. We are mammals. We are built to understand this. An object moves through space from a point A to a point B, with corresponding objective and subjective experiences as you are doing that. **X** is the mental space of experienced motion along a path through space. **X** is a subset of **E**, the mental space of events in general. **X** is a particular kind of event. So the input of experienced motion through physical space—**X**—has many immediate connections to

E.

Let's talk about some of the structure of **X**. If we travel from A to B and then B to C, we know that the event of traveling from A to B is over before the event of traveling from A to C is over. This is part of our experience of moving in space. This comes from our ability to order events. So all else being equal, the relative distances correspond to the ordering of the events: AB is shorter than AC; the event AB is over before the event AC. For **X**, we have a notion of fast and slow, and it's not the one we use in physics. So we can say that AB is faster, that going from A to B is faster than going from A to C even if we move at the same speed on those two paths. I can say, "*Oh, it is faster to walk across the street than to walk to the bakery*", even though you may know that you would go at one mile an hour in both cases. The *fast* in this domain can correspond to the distance of the path of the event or the time interval of the event. That's just a way we use the word *fast*. So in **X**, the event of traversing the path is connected with the path.

We can blend **X** and **E**. The blending of the action of traversing a path and experiencing that motion with any other event was discussed in metaphor theory first under the label events are actions and later as the "event-structure" metaphor. Using this blend, we can understand any event by blending it with experienced motion along a path. I label this blend **E/X**. In **X**, I can go through the park. Now, in **E/X**, I can go through the lecture. I can go through the lecture because in the blend, I am moving along a path. Giving the lecture is moving along the path. I can *go through* the lecture because it's an event and it is an event that in **E** isn't moving along the path, but in the **E/X** blend it is experienced motion along a path.

Once we have **E/X**, any event can be experienced motion along a path. This is quite familiar to everybody. Once we have **E/X**, we can talk about going through a lecture, going through the book, going through college. We can go through events. We can come to the end of the lecture. We can come to the end of the party. We can go through them just the way we can come to the end of the park.

Now recall our wonderful network for the cyclic blended day, **C**, and the periodic physical event that is used as an anchor, **A**, and their blend, **M**. **M** has universal events, like hour, that are thought of as parts of the cyclic day and indicated by the position on the physical periodic event, but we do not yet have *time* anywhere in this network. We use **C** for the cyclic day, **A** for the anchor, the thing in the world that has periodic motion, and **M** for the blend that gives you hours.

Now, of course, we can blend again. We can blend **M** with **E/X** to get what I will label **E/X/M**. And this is the blend that finally gives us a conception of time. This is the first spot in the network in which you get *time*, by which I mean: once you have **E/X/M**, then every event, every local event, every particular event involves a universal event. I will explain that.

Here in **M**, there are hours. Here in **E/X**, there are events that are blended with the particular event of movement from A to B along a path. Now when we blend **M** with **E/X**, it is very easy to do because of course we have all of these wonderful event

connections across all these spaces. Suppose I am delivering a lecture. In **E**, I am delivering the lecture. In **E/X**, I am *going through* the lecture or *coming to the end* of the lecture. But in **E/X/M**, I am also going through universal events, like minutes and hours, which are projected from **M**. In **E/X/M**, to go through the particular event, namely this particular lecture, is also to go through the universal event, like the hour. In **E/X/M**, as I go through the lecture, I am going through the minutes. Where are these minutes? They are the universal events from **M**. And in **E/X/M**, every particular event is inside a universal event, or several universal events.

In fact, this hour that I am going through when I go through the lecture is the same hour that you are going through when you are going through the lecture. It's the same hour that the person across the street from the lecture hall is going through, even though they are going through a very different event. Amazing. This fact is the result of our having constructed **E/X/M**.

Once we have **E/X/M**, we have an understanding of our world, in all its full particularity, according to which there are universal units of measurement for time just the way there are universal units of measurement (foot, meter, etc.) for space. So, for instance, it is not possible once we have universal units of measurement for space for some object to have length but no measure. It must take up so many meters, because there are universal units of length, and any particular length must lie inside some universal measurements of distance. Any particular distance, anywhere in the universe, must go through so many millimeters or centimeters. These are universal units of measurement. You don't have a particular length in the world without it also occupying a universal measure of length expressed in terms of universal units of spatial length.

Now, similarly for time, once we have **E/X/M**, you can't have an event that is a special event that is not also inside the universal event. Our particular event of this lecture lies within a universal event of an hour and a half. We can give it a local label—3:30 to 5pm China Standard Time. But that local label indicates a universal temporal event: all the events in the universe are fitting inside this hour and a half, this universal event.

Once we have **E/X/M**, every event over here in **E**, where there is subjective experience, also lies inside a universal event, which we get from **M**. We all go through the day. We all go through the hour. We all go through the minute. All events—such as the falling of the leaves from the tree—go through the same hour. It can take an hour for the water to drain out of the bucket. We have universal temporal events in **E/X/M**. So you can say things like *I went through the first hour much more quickly than the second hour*.

Notice that we can bring down from the **M** space the sense that we all “objectively” agree about the universal time events. We can all agree, within tolerance, if we check, when the hour starts and ends. It's the “same” hour for all of us, in the sense that an hour is an hour no matter what your experience is during the hour. But over here in **E**, we have subjective experience. The “same” particular event can make each of us feel differently.

In the local event, the first hour can feel very different to you and to me. Suppose I lecture for two hours. In **M**, we all agree that it is two hours. But the first

hour could feel to me very zippy, and the second half very plodding. I can then say, *I went through the first half of the lecture much more quickly than I went through the second half of the lecture.* Now I could mean by this that the two halves of the lecture are taking different amounts of time but I can also mean that, while they take equal amounts of objective time, from **M**, my experience of the particular events is very different.

There is “objective” topology in **M** and “subjective” topology in **E**, and both of them are available to **E/X/M**. We can even use both of them in the same expression. When we refer to the *eight-hour week*, that’s “objective” topology from **M**; that’s the universal measurement. There is a week and there is an eight-hour day. *It's amazing how the eight-hour work day — so that’s the objective topology — is longer on Monday than it is on Friday.* Now, the fact that it is longer on Monday than it is on Friday is subjective topology, from **E**. Monday is when you go back to work after the weekend in the United States. And Friday is when you end work and then have the weekend. So people in the United States constantly say things like, *Monday is dragging, Monday is longer than Friday*, and so on.

Because of **M**, we all agree that the 8-hour workday is the same on Monday and Friday. Both last 8 hours. That’s from **M**.

But the subjective experience from **E** is now that one of them is longer than the other. We can use them both at the same time.

The Onion, a humor magazine, ran a hilarious piece on October 16, 2007. It was titled, “It Only Tuesday”. If you look at the language in this piece, which is about time, you will see that almost none of it makes any sense according to the usual analyses of TIME IS SPACE. But it all makes sense if understood through the elaborate conceptual integration network I am analyzing here in this talk. Notice, for example, the interweaving of objective time from **M** and subjective feeling from **E** as the article works in the conceptual integration network for time I am presenting in this lecture:

WASHINGTON, DC—After running a thousand errands, working hours of overtime, and being stuck in seemingly endless gridlock traffic commuting to and from their jobs, millions of Americans were disheartened to learn that it was, in fact, only Tuesday.

"Tuesday?" San Diego resident Doris Wagner said. "How in the hell is it still Tuesday?"

At this point, the article shows a picture of New Yorkers waiting for a subway train, with the caption, “Already the week is unbearable for these New Yorkers awaiting a subway train, and it’s only fucking Tuesday.” The article continues:

Tuesday's arrival stunned a nation still recovering from the nightmarish slog that was Monday, leaving some to wonder if the week was ever going to end, and others to ask what was taking Saturday so goddamn long.

"Ugh," said Wagner, echoing a national sense of frustration over it not even being Wednesday at the very least.

According to suddenly depressed sources, the feeling that this week may in fact last forever was further compounded by the thought of all the work left to be done

tomorrow, the day after tomorrow, and, if Americans make it that far, possibly even Friday, for Christ's sake.

Fears that the week could actually be going backwards were also expressed.

"Not only do Americans have most of Tuesday morning to contend with, but all of Tuesday afternoon and then Tuesday night," National Labor Relations Board spokesman David Prynne said. "If our calculations are correct, there is a chance we are in effect closer to last weekend than the one coming up."

Added Prynne: "Fuck."

At this point, the article shows a picture of cars in a traffic jam in California, with the caption, "Believe it or not, it's not even goddamn lunchtime yet for these commuters in Southern California." The article continues:

Reports that this all has to be some kind of sick joke could not be confirmed as of press time.

Isolated attempts to make the day go faster, such as glancing at watches or clocks every other minute, compulsively checking e-mail, hiding in the office bathroom, fidgeting, or reading a boring magazine while sitting in the waiting room, have also proven unsuccessful, sources report.

The National Institutes of Standards and Technology, which oversees the official time of the United States, is flatly denying that it has slowed or otherwise tampered with Tuesday's progression.

"The current Tuesday is keeping apace with past Tuesdays with no more than one ten-thousandth of a second's variation at the most," NIST spokeswoman Dr. Geraldine Schach said. "However, I sympathize with the common consensus that this week has already been a colossal pain in the neck."

Labor Secretary Elaine Chao released a statement addressing widespread speculation that it might as well be Monday for all anyone cares.

"We understand this day has been tough on many of you, what with meetings mercilessly dragging on and an entire stack of files still left to organize," Chao's statement read in part. "Yet we urge Americans to show patience. The midweek hump is just around the corner, and we have strong reason to believe that Saturday will be here before you know it."

"Go about your lives as best you can," the statement continued. "Do not, we repeat, do not take a sick day, as it'll make the rest of the week that much harder to endure."

In the meantime, citizens are doing their best to cope with the interminable week, though Tuesday is still hours away from ending.

"The more I try to speed it along, the longer it almost seems to take," said Dale Bouchard, a Chicago-based broker who has been waiting for today to be over since it first began earlier this morning. "Honestly, today could not have come at a worse time this week."

In the meantime, the latest wristwatch consultations indicate that it is somehow still Tuesday, if that makes any sense at all.

Sometimes, we use the subjective topology from **E**; sometimes, we use the

objective topology from **M**.

Notice that we must do a great deal of impressive blending, with emergent structure, and with compressions, just to achieve some of the understandings of time that people think are absolutely basic. We must do a great deal of impressive blending just to account for how people talk about time in everyday language.

If you ask how long it took me to go through the lecture, and I say, “*I went on too long. It was an hour and five minutes long.*” In that case, I am using the objective topology from **M**. If I say, “*I went on too long. It took me centuries*”, then I am using the subjective topology, from **E**. Consider again “It’s amazing how the eight-hour work day is longer on Monday than it is on Friday”. The person who says this is not crazy. And we don’t say, “Oh, you couldn’t possibly be a human being because you are saying what’s false.” In fact, this statement is contradictory if all of the topology is recruited from **M**. But of course, only the first part uses the **M** topology. The second part uses the **E** topology. We interpret the sentence “It’s amazing how the eight-hour work day is longer on Monday than it is on Friday” by recruiting conceptual structure from the big conceptual integration network for understanding time. Words do not mean; words are prompts to construct meanings by applying mental operations we already possess to things we pretty much already know.

Imagine a married couple that must go to dinner with the wife’s parents. The husband says, “*Remember that an hour with your parents is faster for me than it is for you*”. The husband is not saying that the universal events of hours that took us so much work to construct vary for the husband and wife. No. He is admitting that in **M** the universal event is the same for both of them—one hour. But in **E/X** that hour has different speeds for them.

Now for the next step in the conceptual integration network. As has been understood within rhetoric and philology for centuries, there is a dual understanding of **E/X**, which I will label **E/X’**. In **X**, someone is moving past objects. But in **X’**, objects are moving past a stationary person. So, for example, if I am on a train and I’m actually going through the forest, I can say, “*I am going past the trees. I’m going by the trees.*” But given that motion is relative, you can choose as your reference point either an ego who is stationary or an ego who is moving. So I can say *I went by the trees* (that’s **X**) or *the trees went by* (that’s **X’**).

Recall that **X** is experience motion along a path, from A to B. Given the relative nature of motion, we can construe this using either **X** or **X’**. We are not talking about scientific truth: we know the trees are not moving relative to the ground, for example. But they are moving relative to the person on the train. So I can say “*The old toll house went by. The old toll house already went by. The rough stretch of road went by. The forest went by.*”

Just as in **E/X**, I could give you my subjective experience of my motion, so in **E/X’**, I can give you my subjective experience of the motion of things past the stationary ego: *That stretch of road went by effortlessly. The first five miles went by effortlessly.* These are quite normal things to say.

When I blend my dual of motion, X' , with the much bigger domain of events that we experience, E , we get the blend well-known from metaphor theory, in which I can say *the trees went by quickly*.

When I blend that motion dual, X' , with the events, E , I can say *the lecture went by effortlessly*. Suppose someone who did not attend today's lecture asks you, "Was it very difficult listening to Turner for an hour and a half?" and you say, "*No, no, the lecture went by effortlessly.*"

You can also say, using E/X' , *the party went by pleasantly*. This is not regarded as highly creative or bizarre thing to say. On the contrary, this is normal language. Why? Because cultures have developed elaborate integration networks. You can integrate events with motion by an agent along a path, or with its dual, motion along a path by objects moving past a stationary agent. You can make the cyclic day blend, and blend it with periodic motion events to create universal events like hours. You can blend these two blends: the blend of events, E , with motion along a path, X , or its dual, X' , and the blend of the cyclic day, C , and the periodic event, A . You can then blend again, to create either $E/X/M$ or $E/X'/M$.

In both $E/X/M$ and $E/X'/M$, you have the amazing emergent property that all the particular events in the universe lie inside the same universal events, such as minutes and hours and days.

Suppose we are working in E/X . Then I can go through the lecture. If we are working in E/X' , the lecture goes by. If I now blend in the space with universal events, M , I can say, if I am in $E/X/M$, that I went through the *hour* quickly. Or if I am working in $E/X'/M$, I can say that the hour went by quickly. In $E/X'/M$, the universal event, the hour, can move.

Remember that I said at the beginning that in the domain of time, you think that universal events like hours can move, but that you do not have that sense for the domain of SPACE, in which feet do not move. Remember that we discussed all those enormous clashes between your experience of space and your experience of time. So I can say *Friday always goes by faster than Monday*. In that case I'm using the subjective typology from the event domain and I'm using the dual of motion. *The hours sped by for him but dragged by for me. It took centuries for the hour to pass*. These expressions are not nonsense. You know how to interpret them. *Those three hours went by slowly for me, but the same three hours went by quickly for him*. You know what I mean because you know this conceptual integration network for time.

If I use the typology of M , which has the objective topology, with X' , the dual of motion along a path, in which events are going by me instead of my going through them, I can say *minutes go by faster than hours*. I can say *The same hour will go by whether you are suffering or having fun*. In $E/X'/M$, the universal events—minutes, hours—move. They move past you. That's the dual and their topology is coming from M .

But you can have it both ways. There was a man who was awaiting word about his friend Michael who was missing in the bomb detonations in London in July 2005. I

watched the interview on CNN. The interviewer asked him how it was going. He was waiting to hear whether the rescue teams had located his friend, who was in the building that was blown up. He was waiting, and he said, “*Well, time goes by really slowly. At the same time, it goes by really fast*”. It makes perfect sense. How? The answer is that he has the subjective experience of not receiving any news about his friend Michael. He wants to receive news. It is very painful and slow, and so he says that the time goes by slowly. On the other hand, time is going by, and the more the hours go by, the less likely it is that he will receive good news about his friend. So he does not want the time to go by, yet it is going by, and he feels that it is going really fast. He has these conflicting experiences, and he expresses them, using the conceptual integration network we have been analyzing.

There are yet more subnetworks involved in this great conceptual integration network for time. We independently we have the notion of retrieval of memory—call it mental space **R**—as blended with distance and space—call it mental space **S**. Label their blend **R/S**. So, using the blend **R/S**, we can talk about *unearthing* our memory, or *bringing our memory a long way up back to the surface* or *having the reach far to get that memory* and so on. **R/S** is an independent blend.

But we can of course blend **R/S** and **E/X/M** or **E/X'/M**, and, then, in that hyper-blend, we can include memory as part of the network of our experience of time. We will look at that hyper-blend in a moment. But before we do, let us notice that there are other additional subnetworks we could analyze. For example, one of the things we know from the space of **M** is a constraint: we all must go through the same hour. We can personify that constraint, that is, blend it with a human agent. This blend makes an agent out of the constraint. And we can give that agent a name, like “Time”. Then we can say,

- *Time drives us forward through the events.*
- *Time marches on.*
- *Time waits for no man.*
- *Never fear: time will carry them along.*

There was a time in the English language when “the hour” could be a name for time. So one could say, “*Come what come may, time and the hour runs through the roughest day*”. You can personify this constraint. That is an additional blend.

Let’s go back to **R/S**, the blend in which a memory can be close or distant, hard to access, so you can talk about *calling up things from the depths of your memory* or *bringing a forgotten event to the surface*. If you blend **E/X/M** with **R/S**, then the subjective feeling in **R/S** that the wedding is very accessible, very close, is mapped onto the subjective feeling about the events of yesterday. So the blend endows **R/S** with a metric using the notion of time. Memory, or recall, now has a time metric. Accordingly, in the **E/X/M/R/S** blend, the word “yesterday” provides an adequate indication of distance in memory. So in **E/X/M** itself, our wedding was not yesterday (assuming it was 18 years ago). But the memory of the wedding as experienced in **R** projects to yesterday in **E/X/M/R/S**, where the wedding of 18 years ago can now be

"yesterday." And *yesterday* is close to the person moving on the path in **E/X/M**. In this case, the ordering topology of **R/S** dominates over the ordering topology in any version of **E/X/M**.

E/X/M/R/S lets us say, *Our wedding was just yesterday. Where have all those years disappeared? Or, the years have really gone by fast.* You need to account for the clash between **E/X/M**, where the wedding was eighteen years ago, and **R/S**, where the subjective feeling of the memory of the wedding is that it is very close. How can something eighteen years ago be very close? There are two ways, as follows. Consider moving in space. Suppose you are on a train and you are going from A to B, and you get to B, and the trip has taken much less time than you were expecting. You could conclude that B is a lot closer to A than you thought or you could conclude that the train is going a lot faster than you thought. If you choose the resolution according to which you were mistaken about the distance between A and B, namely, A and B really are very close to each other, then, in the **E/X/M/R/S** blend, you have the inference that the years disappeared. Today and your wedding are close because the years are not really there; they disappeared. Alternatively, the journey from A to B can take much less time than you expected because you were moving fast. In **E/X/M/R/S**, you can accordingly conclude that today and your wedding seem so close because the years went by really fast: *My, how those years have flown.*

Notice that all these examples I have been given are really basic, everyday expressions: *the years went by really fast* is a standard way to talk. You don't have any trouble understanding these expressions, or using them. When someone says such a sentence to you, you don't have to stop and think. If your eyes are glazing over as you look at the sketch in the slideshow of conceptual integration networks for time, it may seem that following the analysis of this network is like solving partial differential equations in seventeen dimensions. But you understand the expressions themselves immediately. Nearly all conceptualization, thought, and inference happen in backstage cognition; they take a great deal of work, and you are doing that work all the time. The only thing that pops into consciousness is little reactions, little products. If following the analysis of this conceptual integration network is difficult, it is only because consciousness is so weak. You have not the least difficulty using these networks in backstage cognition. Trying to follow them in consciousness is something quite different: that's hard.

You don't see in consciousness how you walk. You don't see in consciousness how you see. You can't call such things into consciousness, except perhaps a little, when you are really focusing. Most of what you do in language cannot be called into consciousness. We are trying to use human mental powers not meant to do this kind of analysis to do this kind of analysis, to bring to the surface operations that you conduct all the time but that you do not even know you conduct, and that you are not designed to see.

I said at the beginning that conceptual constructions are never a matter of simple metaphoric mappings from one domain to another. This is the big news for metaphor analysis. Metaphor analysis during the last fifty years revealed many important things.

It showed us some mappings, but now we see that those mappings are abbreviations, compressions of much more elaborate blending work. For basic metaphors like time is space or time is money, there are always many mappings going on, many more than people thought.

Now I am going to present to you an example of innovation. This example was noted by Nathaniel Smith. It is a passage from a memoir “Emily's Diary, November 5, 1913,” in *King of Morning, Queen of Day*. The narrator is making his scrapbook up in the attic. That’s the experience that is being narrated. The narrator is up in the attic, where all the old things are stored.

Remarkable—when I am sitting on a cushion on the floor, busy with scissors and glue pot, the time just vanishes. Before I know it the latticed rectangle of pale autumn sunlight has moved from the left wall across the floor to the other wall and Mrs. O'Carolan is calling me for supper. Perhaps time is flowing faster up there in the attic. Perhaps the accumulated mass of the past gathered there is pulling time out of the future faster, like a weight on a line. Or perhaps, more mundanely, it is only that I am getting older every year and that it is the accumulated weight of time behind me that is unreeling the years with ever-increasing speed. What a horrible thing it must be to grow older and find that ever-decreasing number of years hurrying you faster, faster toward your grave, as if time were impatient to be rid of you.

This passage looks highly creative. But notice that you don’t have the slightest difficulty understanding any of it. The reason you have no difficulty understanding it is that every piece of it is using structure that is already in the network I showed you for your everyday expressions about time. There is a little more exploitation of these connections in the network. So, for instance, for physics that we know a bigger weight pulls harder; when you are up there in the attic and the accumulated mass of memories gets blended with the mass of physical stuff, it can pull time faster. There is a blend, a personification of time here. Time hurries you, only now it is hurrying you faster. The movement of the shadow of the latticed rectangle across the room is a periodic event, as we saw in the mental space A. That shadow moves across the room every day. So the shadow in the room is a timepiece. You can map it to your watch.

Today, I have given a reanalysis from the blending perspective of the basic metaphor TIME IS SPACE. You can find the details of this reanalysis in an article titled “Rethinking Metaphor”, which Gilles Fauconnier and I wrote. But I want to emphasize that there is nothing special about TIME IS SPACE. All basic metaphors are simplifications, compressions, for much more elaborate conceptual integration networks. TIME IS SPACE merely provides us with the example for today.

Today, I have tried to make various general points about conceptual metaphor. Let me summarize them.

Integration networks. First, conceptual products are never the result of a single mapping. What we’ve come to call conceptual metaphors turn out to be mental constructions involving many spaces and many mappings in elaborate networks constructed by means of overarching general principles that come with conceptual integration. And these conceptual integration networks are far richer than the bundles

of pairwise mappings that have been considered in most theories of metaphor.

Cobbling and sculpting. Integration networks are never built entirely on the fly, nor are they entirely pre-existing conventional structures. Integration networks underlying thought and action are always a mix. On the one hand, cultures build networks over generations. Over long periods of time they get transmitted down generations. Techniques for building particular networks also get transmitted culturally. But people are capable of innovating in any context because they are using the same mental operations to innovate that were used to create the conceptual metaphors to begin with. Innovating in the network and creating the network depend on the same mental operations. That's why we can change or modify these networks in the moment—we are just doing what has always been done. Integration networks seamlessly combine conventional parts, conventionally-structured parts, and novel mappings and compressions, because all these parts arise through the same mental operations.

Compression. There is a crucial and remarkable conclusion of recent blending theory that was overlooked by early metaphor theory and by early blending theory. It was overlooked by Fauconnier and Turner in the early years, when we were the only blenders in cognitive linguistics. The conclusion is this: Integration networks achieve systematic compressions that make it possible for us to manage the network. The ability to use standard techniques and patterns of compression and decompression enables us to work at once over elaborate integration networks.

Overarching goals other than projection of inference. Metaphor has overarching goals other than the projection of inference. In many cases, crucial inferences that are projected are largely driven by the target rather than the source. Even more important, often the overarching goal is compression rather than projection of inference. Inference seems really important to us, because it has to do with truth. That is why presentations of blending often start with inference—to show that it is important even for what we tend to think are the most important aspects of thought. That's by the way why I start with the Buddhist monk. I say there is a riddle and there is a solution to this riddle and it looks kind of mathematical and you can't get it without blending. This leads people to think blending must be important. This is understandable. But in fact, compression is often more important than inference

Emergent structure. The focus on single mapping and inference transfer in early metaphor theory left out many of the powers of integration networks, in particular the ability to develop emergent structure. time is space is nothing if not a network that creates emergent structure and emergent concepts that were not available in the original input spaces.

Thank you!

Lecture Seven

Frame Blends

Thank you! You are so generous. Thank you for that splendid introduction. Thank you for inviting me to your beautiful university, with its lovely trees and its impressive buildings and its intelligent students.

We have been talking about the way in which a basic and very important mental operation — conceptual integration—makes possible human linguistic ability. Today, what I'm going to talk about is the way in which this operation interacts with two other very well-known human abilities. The first is the ability to frame and the second is the ability to conceive of our experience in terms of stories—that is, suites of actions involving agents who interact.

A frame is an organizing structure that we use to guide thought and action. For example, suppose you go to a library. Most of us have a frame for a library. In the library frame, there are books and there are librarians and there are rooms and you know how to operate in a library because you have a frame for a library. You have a frame for a restaurant. You have a frame for a taxi ride. This makes it possible for me to come to some city in China I have never visited and still succeed in getting a taxi ride, even though I may need to adjust my frame just a little on the basis of my environment. As we discussed in a previous lecture, I have a packed little mental bag of stuff I unpack to plug into the world as I go around it: my frames for library and taxi ride are part of this portable, compressed, human-scale equipment I carry with me to unpack and plug into the world.

It happens that we often blend frames. What do we mean by that? Consider our previous example, in which I asked you to consider a groomsman at a wedding. He is doing his job. He has a conceptual frame for a wedding, and he uses it to guide his action. He unpacks it to plug into the world so he can do his job. The wedding has the bride and the groom and the wedding. And you know what to do at a wedding because you have a frame for wedding. Languages provide keys to frames we know: all we need to do is mention a particular word, or maybe two, or a clause, or a phrase, and the whole frame is activated in the minds of our audience.

We often blend frames. Suppose that this man is at the wedding, but while he is at the wedding, he is thinking that three weeks ago he was diving with his girlfriend in search of treasure off Cabo San Lucas. His girlfriend is not here at this wedding, but he can think of both of these things at the same time, and as we discussed it's a remarkable problem in cognitive science that he can think about both the wedding and about the diving. They conflict with each other. How does he avoid becoming confused, and making mistakes?

As we discussed, he can make a connection between the two different mental spaces. He can connect the bride to his girlfriend. There is an analogy between them. He can project from the two mental spaces to a blend. He can connect himself diving

to himself at the wedding—that would be an identity connection. But he can also connect himself as the boyfriend to the groom. The most important such connections between mental spaces are called vital relations. They are connections such as change, cause-effect, time, space, identity, analogy, similarity, representation, part-whole, and so on.

We, and the groomsman, can project selectively from the wedding and diving input spaces into a third space that we will call the blend. In the blend, the groomsman is imagining that he is marrying his girlfriend right here, right now. The wedding comes from one mental space, the girlfriend comes from the other mental space, and they are blended. As we have discussed, marrying the girlfriend is not in either of the inputs. Just by putting them together, you have already created emergent meaning in the blend. The groomsman can run the simulation and see how he feels about it. Maybe he never thought of it before he participated in this wedding. We are eminently capable of running these mental simulations.

In a moment, in the slideshow, I am going to show you a joke. One of the nice things about frame blends is that on the one hand some of them are quite universal, on the other some are highly culturally specific. Our frames, and accordingly our frame blends, are often quite culturally specific. Consequently, a particular frame blend can be impressive, powerful, or hilarious in one culture and quite unintelligible in another, because different cultures have different frames. Here is the joke frame blend. In the United States, we have universities, and in universities, at least in research universities, the professors can get tenure. That means that they not only have a job but that they get to keep their job unless something really exceptional happens, such as plagiarizing somebody else's research. In order to get tenure, the faculty member must go through a tenure review, which is an elaborate assessment. I have tenure, for example. People who review you for tenure ask themselves questions like, "Well, does this person have good scientific practice? Has this person published a lot at good university presses as sole, first, or last author? Has this person contributed to the community? Has this person brought in enough grant money to pay for their laboratory? And so on."

So, there is one frame: achieving tenure in a university. At many places in the world, this frame is not available. But it is in the United States, and perhaps you are acquiring this "tenure review" frame if you don't already have it, because I am explaining it, and you are blending my explanation with various inputs you already have because you are members of the Chinese university system. In the tenure frame, your job performance undergoes a particularly thorough scrutiny at a set time, and if you do very well, then you get tenure, and so get to keep your job. If you have questions about the "tenure review" frame, ask me now. It can seem quite bizarre. Indeed, it often seems very bizarre to me, even though I have tenure and have participated by now in hundreds of tenure reviews for other people. It's important that you have at least the rudiments of this "tenure review" frame. How familiar is this to you? Are we Ok? Oh, yes, good, somebody is helping me out. Thank you! Xiexie.

Now consider another frame: the God of Christianity. This is a Christian frame, the one that I'm going to be using here. In this frame, there is God and Jesus Christ.

You are familiar with that, I see. Here is the joke. It comes from the internet.

Why God won't get tenure

1. Published only one book.
2. It was in Hebrew.
3. It had no references.
4. He did not publish it in referenced journals.
5. Some doubt He even wrote it Himself.
6. He is not known for His cooperative work.
7. Sure, He created the world, but what has He done lately?
8. He did not get permission from any review board to work with human subjects.
9. When one experiment went awry, He tried to cover it up by drowning all the subjects.
10. When sample subjects do not behave as predicted, He deletes the whole sample.
11. He rarely comes to class—just tells His students to read the Book.
12. It is rumored that He sometimes lets His Son teach the class.
13. Although He only has 10 requirements, His students often fail His tests.
14. He expelled His first two students for learning.
15. His office hours were infrequent and usually held on a mountaintop.

This joke is “Why God won't get tenure?” Now, of course, God is not a professor and a professor is not God, although many of them think so. This joke asks you to take two very different frames: one for the Christian divinity, the other for the tenure review process. I predict that much of this joke will make sense to you, but that some of it won't. The parts that will be hard to understand arise where you are lacking the full details of one frame or the other, as understood by the person who made the joke. This joke asks you to blend these two frames. Why will God not get tenure? Well, he published only one book. You might say that about an assistant professor in the United States coming up for tenure review at a research university.

There is an ingenious accidental connection between the two frames. Over in the Christian God frame, we have the Bible, often referred to as “the Book” or “the Good Book.” *Bible*, etymologically, means book. God “published” only one book. In the blend, “publishing” takes on a blended and extended nature. This is a joke because the Bible is by general agreement far and away the most influential book, ever. The joke is that all that influence is not enough, because it is only one book. This makes the reviewers look silly: they are just counting: 1, 2, 3. So 2 stupid books that were read by only 2 people each—the author and the author's mother—are more important than the most influential book in the West ever. That looks silly. In general, this joke makes the tenure review process in the United States look mindless. You see the silliness because the ridiculousness arises in the blend.

The lack of wisdom in the blend can be projected back to the standard “tenure review” frame, and get us to reconsider what we think of that standard frame. So this joke can be not just a joke, but also rather a satire, a criticism of tenure committees. On

the basis of the blend, we can develop conceptual structure, judgment, and cultural views that we can then impose on the things that we do take seriously, like universities.

The joke says that the Book was in Hebrew. Of course, it wasn't actually in Hebrew; only part of it was in Hebrew. But to discount the book because it was written in a language preferred by the U. S. research community is the kind of thing that somebody in a tenure review might do, at least by implication. The joke says that the book had no references. Now this is a very funny thing to say in the United States. Not citing enough people is taken in the United States as not being an insider in the scientific community.

The joke says that God did not publish his Book in refereed journals. Refereed journals have nothing to do with the frame of the Christian God, but a lot to do with the frame of a tenure review.

The joke says that some doubt that He even wrote the Book himself. This is of course because every book of the Bible actually has an author or multiple authors or a community of authors over time. Perhaps the author is "Moses" or "Mark" or "Luke" or "John" or "Paul."

The joke says that He is not known for His cooperative work. Again, this is from the "tenure review" frame, and in fact it fits well with the God frame, but the criticism is turned back on itself because God has no need to collaborate.

When one applies all these "tenure review" criteria to God, then God is found wanting, and won't get tenure. So the very smartest thing in the universe won't get tenure, which suggests that something is wrong with the tenure review process. This recalls the flood.

My favorite entry in the joke is that God "expelled his first two students for learning". That's Adam and Eve in the Garden. Eve plucked fruit from the tree of knowledge, so God expelled them from the garden. A student is expelled from the university for not learning or for doing the wrong kind of thing.

This example is what I mean by a frame blend. Two quite different frames are blended, with emergent structure that is not in either of the inputs, such as that the tenure review process is misconceived. The frame blend makes a rhetorical point.

Let us look at a frame blend from Dante. Actually, we will see that underlying it is a general template for frame blending in Dante, one often used in Dante's *Inferno*.

Dante's *Inferno* is a book about judgment, cosmic judgment. Dante enters Hell, guided by Virgil, the author of the Aeneid, and he encounters people, or rather souls, who are in Hell, and he asks them how they are doing and how they got there. One of them is Bertran de Born. Bertran de Born was a famous troubadour, a poet and singer. Bertran de Born was also politically active, and, at least as Dante understood it, he contributed to the war between the English king, Henry II Plantagenet, and the son of the English king. In other words, he got father and son to fight with each other. For this, Dante places Bertran de Born in the eighth circle of Hell, for the sowers of schism. Metaphorically, we would say that he "divided" the son from the father. He split them apart. We say such things routinely, as part of everyday language.

Bertran de Born's punishment as portrayed by Dante in the *Inferno* is the result of a frame blend. This general procedure for blending frames in the *Inferno* goes like

this: if there is a metaphoric understanding of the sin, call up both the frame for the sin and the frame we use to understand the sin. In the case of Bertran de Born, you would call up the frame for one agent's causing an opposition between father and son, between King and Prince. You would also call up the frame for dividing a physical object in two, for cutting the physical object in half. Metaphorically, Bertran de Born set father and son apart, separated a unity that should not have been separated, broke a connection that should not have been broken. The punishment is: *impose on the sinner the frame used metaphorically for understanding the sin*. That is, we impose on Bertran de Born the frame for separation of an object. His punishment is that his head is cut off his body, and in Hell he carries his head around in his hand.

Notice that in the blended understanding of the sin, what gets separated is not the sinner, but rather whatever it is that the sower of schism separates. So Bertran de Born separates the King and his son. That is a metaphoric blend. In that blend, Bertran de Born is not separated. Instead, the King and son are separated. That is one blend. But now, Dante makes quite a different blend for inventing the punishment for this sin. In this "cosmic justice" blend, the *sinner* gets separated. Bertran de Born, from the political mental space, is blended with the object that is separated. Now, in the "cosmic justice" blend, the *sinner* gets separated. When Bertran de Born talks to Dante in Hell, it is the detached head that is talking. Bertran de Born says, "*Because I parted people so joined*"—that means, because he separated the father and the son, "*I carry my brain, alas, separated from its root, which is in this trunk. Thus is to be seen in me the retribution.*" In other words, you see the punishment in the sinner. And this is the pattern of cosmic justice. The frame you use for the metaphoric blend used to understand the sin itself is now imposed not metaphorically but literally on the sinner. This is a template for cosmic justice, cosmic justice through a particular scheme of blending. This template for frame blending is used often throughout Dante's *Divina Commedia*.

Here is a frame blend we looked at the other day. It was alleged that Bill Clinton, the American president, had had an affair with Monica Lewinsky, a White House page. When this allegation first broke, Bill Clinton seemed not to suffer any problem at all. Someone said, "*If Clinton were the Titanic, the iceberg would sink.*" People who hear this expression universally feel that it is metaphoric.

Notice that the one thing you know about the Titanic is that it sank. You also know that icebergs cannot sink. They can be submerged, but they cannot sink, because ice is less dense than liquid water. To create this Clinton-Titanic blend, you project some things from the "source" domain of the metaphor, such as the Titanic and the iceberg. They come from the frame of a boat's crashing and sinking. From the political frame of political scandal and potential impeachment, you project down the politician and the difficulty. In the blend, the politician is the boat and the obstruction to the boat is the political difficulty. That's the basis of the frame blend.

From the political mental space, you take the fact that the politician—Bill Clinton in this case—is not having any trouble, despite the fact that the various political commentators and other politicians are trying to give him trouble. They are like the iceberg. But it is no problem for him.

Accordingly, in the blend, the boat/politician is not bothered by the obstruction/difficulty. He does not sink. When we specify the boat/journey/obstruction frame with the particular values for the Titanic and the iceberg, this makes it possible for the blend to have a Clinton-Titanic that strikes the iceberg but does not sink. Using the Titanic and the iceberg creates a great hyperbole: the Titanic was supposed to be unsinkable, but the iceberg was such an obstruction that it sank even the Titanic. In the blend, the Clinton-Titanic is so strong that not even the iceberg can make a dent. It is important to see that this structure comes from the “target” of politics rather than the “source” of the Titanic and the iceberg.

Notice that the two frames here—political difficulty versus obstruction of a boat—are very different. There is selective projection from each of them. We bring elements from each of those frames down to make a new frame for the blend, one with emergent structure. In the blend, but in neither of the input spaces, the Titanic is indeed unsinkable. In the blend, but in neither of the input spaces, an iceberg can sink. This is new, emergent physics.

Here is another frame blend: *You are digging your own financial grave*. Some financially conservative father says to a son who is investing in the stock market: “*You are digging your own financial grave*.” People feel that this is a metaphoric expression. We have two frames here: one for investment, another for digging a grave. There is selective projection and emergent structure in the blend. For example, in the blend, the person who is digging the grave is the one who dies. Moreover, he does not know that he is digging a grave. This structure is not available from the “digging a grave” frame. It is new in the blend.

Now let us look at a more extended frame blend. There is a popular book in the United States, whose title is *Seabiscuit*. Seabiscuit was a horse, a race horse. Seabiscuit raced in 1938. This was in the run-up to World War II, and just at the end of the American depression. Morale in the United States was not very good. Seabiscuit was a kind of lazy horse and Seabiscuit was going to have to run against a very well-trained horse named War Admiral. We have in this understanding a frame for a race between horses. This is a human-scale frame: we have a few animals, and we have a race, right here in front of us.

But we are going to need another frame here, the one for war between nations. This frame is specified as a war between the Allies and Germany, or between the United States and Germany. As we will see, the horserace frame and the war-between-nations frame are blended. See what one reviewer of *Seabiscuit* says: “*Seabiscuit: An American Legend. Seabiscuit tells how an unimpressive older horse with crooked legs and a short tail*” — in other words something that shouldn't win, something that has everything going against it — “*stole the hearts and minds of the American people during the Depression. In 1938, the No. 1 newsmaker was not Franklin Delano Roosevelt or Hitler; it was a horse that defined the word 'underdog.'*” That's a little witty: an underdog is the dog in a dogfight that is not going to win. So here we have a term for dogs applied, through blending, to a horse. “*It was the indefinable quality of 'being game'*” — that means sporty, lively, up for the challenge—“*that captured*

Americans. As one observer put it, Seabiscuit would rather die than be beaten in a race. Yet, unlike many champion thoroughbreds, his off-track personality was low-key, appealing and, frankly, lazy. He was a glutton for food and enjoyed the friendship of a horse named Pumpkin. In short, he seemed the American Everyhorse”.

In this frame blend, the vast character of the people of one of the nations at war is being blended with the temperament of a single horse. This so-called “character of the people” involves millions of people in a time of war and depression. It is already a strong compression, and it is blended with the personality of this horse: “*the American Everyhorse, the equine version of how we see ourselves*”. The race between Seabiscuit and War Admiral was taken as a kind of symbol of the United States as it was going to have to enter or might have to enter the war with a very finely-tooled German army.

We have seen some other frame blends before in these lectures. We saw double-scope blends in which the frames clashed. I won’t go again through the blend involved in *Harold and the Purple Crayon*. As you recall, Harold has a purple crayon, and what he draws is real. One frame is the frame of drawing. Another frame is the frame of real-world objects, like the moon, and paths, and windows. These frames are blended to create a new reality: in this blended frame, what one draws is real and what is real has been drawn. The moon comes into existence by being drawn and then it gives light. This is a very popular book for young children. In this frame blend, there is a new physics: one can travel through the world by drawing. This is locomotion by drawing, something not available in the input frames. Frame blending has much to do with how we are distinguished from other people.

Let us begin to talk about the way in which frame blends give us new language or give us grammatical constructions that can be used in new ways. This is Catullus’s *Elegy 101*:

multas per gentes et multa per aequora vectus
advenio has miseras, frater, ad inferias,
ut te postremo donarem munere mortis
et mutam nequiquam alloquerer cinerem.
quandoquidem fortuna mihi tete abstulit ipsum.
heu miser indigne frater adempte mihi,
nunc tamen interea haec, prisco quae more parentum
tradita sunt tristi munere ad inferias,
accipe fraterno multum manantia fletu,
atque in perpetuum, frater, ave atque vale!

Through many nations and over many seas, the poet has come to carry out the funeral rites for his brother. He is visiting the tomb of his brother, who died in a distant land. But the voice of the poem is the poet speaking to his brother! Addresses his brother directly, using the language of direct address. He says that he has come to speak in vain to speak to the mute ashes.

It’s important to understand ashes are the result of cremating the brother. The speech of the poet to the brother is in vain because the brother cannot answer. So, in

one frame, we have two human beings who can talk to each other. In the other frame, we have a person and the ashes of cremation of another person. These are very different frames, of course. In the blend, the brother and the ashes are fused. Catullus is not deluded. He is not having a mental breakdown. On the contrary. Notice what happens. The projection is selective. In the blend, one can talk to the brother—that is from the frame of two people talking. But the brother cannot answer—that is from the frame of the person visiting the ashes of cremation.

What does Catullus call the ashes? He addresses them as “brother”, *frater*. In the blend, “brother” has become a word that can refer to ashes. That’s not because if you look in a Latin dictionary, you will find that “frater” has as one of its meanings “ashes”. No, it’s because in one of the inputs there are words, there is language, there are linguistic constructions that can attach to the element for the brother. That language can be projected down into the blend, and now we have language for referring to the blend. The ashes in the blend are fused with the brother, so now that element can be referred to as “brother”. In the blend, *brother* is grammatical for speaking in a scene in which everyone knows there is really no brother, but in the blend there is. Catullus uses the Latin word for “you”: *te*. Catullus uses not just individual words, but a large conversational structure of address and discourse in language. These now become available for the blend.

How can we account for this phrase “mutum cinerem”—which means “mute ashes”? That is a very interesting phrase, because outside the blend, we do not speak of mute ashes. If I have a fire in my fireplace, using wood, and after the fire is out, I can ask one of my sons to “clean up the ashes.” I can’t ask them to “clean up the mute ashes”, or if I do, I am telling thereby to make a blend, because the word “mute” is for something who can’t speak even though it used to be able to, or that can speak but for some reason has stopped, or is a living human being someone deprived of the ability to speak. “Mute” is only for something that in some way is expected to be able to speak. So you can’t refer to mute ashes in the fireplace. But to say of the blend, “In vain address your mute ashes” is regarded as particularly poignant. Why? Because in the blend, the ashes are the brother—a selective projection from the brother. In the blend, *brother*, *you* and *mute* all acquire a perfectly natural use. Remember that words do not mean. Words are prompts to get us to construct meaning. It would be a mistake to say “Gee, what is about the word *brother* in the dictionary that makes it suitable for referring to ashes?” That’s not the way to look at it. We use linguistic constructions to prompt for various kinds of conceptual networks.

Those of you who attended the previous lectures will remember Racine’s *Phèdre*, in which Phèdre speaks to Hippolytus, describing an imaginary story in which it is Hippolytus who saves the Cretans from the Minotaur, and it is Phèdre who helps him. In the original story, it was Theseus who killed the Minotaur, and he was helped by Ariadne. Theseus is the father of Hippolytus, whose mother was an Amazon. Phèdre is now married to Theseus, so she Hippolytus is her step-son. You won’t understand a word of this if you were not in the other lectures. In the original story, Ariadne provided Theseus with a thread that he could spool out behind him to help him

find his way out of the labyrinth that held the Minotaur. Theseus and Ariadne were lovers. Phèdre, drawing an analogy between how Hippolytus looks and how Theseus, his father, looked in the days when Phèdre fell in love with him, indicates that in this imaginary story, she, not Ariadne, would have helped Hippolytus escape the labyrinth. These compressions have a lot of warrant: Hippolytus looks like his father, and Phèdre is Ariadne's sister and was in Crete and fell for the hero who killed the Minotaur. So if Ariadne would have helped, why not her sister Phèdre. And if Phèdre fell for Theseus, why not for Hippolytus—in the blend, of course. In her description of the scene to Hippolytus, Phèdre says “but no, a thread would not have satisfied your lover”—*votre amante*. This language, “your lover,” now applies in the blend to Phèdre: in the blend, Phèdre and Hippolytus are lovers. Language that cannot be applied to them outside the blend can now be used of them. “Your lover” picks something out in this blend just as “mute ashes” picked something out in the blend in *Catullus 101*.

We don't have to have new words to refer to these blends. We just need language to pick out what is in the blend. We will see in future lectures that this kind of blending can indeed lead to new language, new grammatical structure.

Wallace Stevens wrote a poem called, “The Snow Man”. It is an amazing frame blend. One frame it uses is the frame for human being. And another frame is for snow, which that can take on various forms. These two frames can be blended to make an everyday snowman. But Wallace Stevens uses more than the usual selective projections to the blend. Wallace Stevens's poem asks us to project human intentionality and cognition into the blend. The result is not just an everyday snowman but also rather a snowman who has some abilities to think. The point of the poem is not to have a little story about a snowman who came to life. Rather, it is to prompt us to think about who we are. The poem is often interpreted as pointing out to us that we are always in the grip of our desires and intentions, and that we impose upon what we see our biases. But the snow man in the poem is nothing. He thinks; he perceives; but he is nothing. And so the snowman is not blinded or misled by imposing his own preconceptions on what he is looking at. This special snowman is a better perceiver than we are. He is a

. . . listener, who listens in the snow,
And, nothing himself, beholds
Nothing that is not there and the nothing that is.

Because he is a snowman and not a real man, he is not full of biases the way we are full of biases, he doesn't project what is inside him onto what he perceives. He sees nothing that is not there. We, by contrast, do. One purpose of this blend is to get you to think about what it means to be a human being. You can't be a snowman but you can put together the blend and imagine it. The emergent structure in the blend can lead you to project back to the frame of human beings, and perhaps give you different ideas about how you might operate, just as the joke about God and tenure prompted you to reconsider the tenure review process.

We find frame blending in every domain of human communication. We are linguists and must be aware of all of it. If we were botanists, we would need to look at weeds as much as flowers. As linguists, we must look at expressions that might be offensive. Here is a memorable example. “*Jail Bait*” is a nominal compound that I heard in my youth. I am happy to say that I have now seen a movie in which it is applied to a 16-year-old boy. In my youth, I heard it applied only to some girls aged 15, 16, or 17. *Jail Bait* was a phrase that was used to refer to an under-age girl that an of-age man found sexually desirable. So what does it mean to be under-age? Well, in the United States, there are various laws according to which people younger than a certain age are viewed as not being competent to give consent to sexual conduct, so sex with them is forbidden, regardless of whether they consent. People under this certain age are referred to as “minors.” These laws can become complicated, I suppose, but “jail bait” was a term for referring to an anatomically developed and attractive teenage girl who was a minor. For an adult male to engage in consensual sex with such a girl could result in his prosecution and incarceration, no matter what the girl claimed to prefer. Now *jail* comes from the domain of human criminality and *bait* comes from fishing. These are two very different frames. *Jail* is where you put the criminal. *Bait* is what you put on the hook to fish with. These are two clashing frames.

Notice that in this blend, the jail is not the counterpart of the bait. In these nominal compounds, it is usually the case that one noun comes from one frame and the other noun comes from the other, and this is a way to prompt one’s audience to create a frame blend.

In this blend, the girl is the bait. The projections are highly selective. The blend provides a great compression. Over in the frame of sexual conduct, there can be a long causal chain, a variety of people, different times, different places, an entire complicated narrative. A long time—months, years—can pass in the narrative before the sexual act occurs.

We have not only the frame of sexual conduct but also the frame of criminality, with arrests and trials and jails. Again, the causal chains, agents, and time scales in this frame can be quite complex.

But in the fishing frame, things are very tightly packed. Out goes the bait, the fish is caught. It’s one fisherman, one fish, stereotypically one bite. The consequences are very clear—if the fish bites with commitment, then the fish is caught. We all know that fishing can be a little more complicated than that in practice. But the frame is at human scale. To achieve this blend, all three frames must be used: criminality, sexual conduct, and fishing. The purpose of the blend is to make the consequences clear to someone who is not paying attention to them. It could be used as an alert: “Wait a minute. You’re starting to get into something here, and you might not see where it leads. This girl is *jail bait*.” In the blend, the adult is the fish. Nobody wants to be caught on the hook.

In this lecture, I have discussed poetry and idioms, but I want to emphasize that we use frame blending to make major decisions about finance, politics, national identities, and the self. We use frame blending in all the basic and important parts of

life. For example, there was an argument made in political science along these lines: “*If Churchill had been prime minister in 1938 instead of Neville Chamberlain, Hitler would have been deposed and World War II averted.*” Remember that Neville Chamberlain was Prime Minister of Great Britain at the time and Churchill was not. This expression prompts you to construct a blend. In one mental space, it is 1938, and there is Churchill, who at the time was regarded as having bad judgment and an extreme position of the sort that would make it impossible for him to become Prime Minister. He was a back bencher. In another mental space, it is 1938, Neville Chamberlain is prime minister, and Chamberlain participates in the appeasement of Hitler. There is yet another mental space, in which it is some years later, and Winston Churchill is the celebrated Prime Minister, the wartime leader of England. You can project to the blend Churchill from 1938, the role of Prime Minister from 1938, and the character of Churchill from the mental space in which he is the wartime leader of England. In the blend, it is 1938, and Churchill is prime minister and has the kind of character we know from his later life. In the blend, instead of appeasement, Churchill confronts Hitler directly, threatens. He has great spine. There is wonderful emergent structure in the blend: Hitler backs down. And there is more emergent structure: no holocaust.

We can see the standard patterns of blending in this example.

Blending exploits counterpart connections between the mental spaces. There are identity connections connecting up all the Churchills in these mental spaces. They are all compressed to one Churchill in the blend, although the projection is selective. The role in 1938 *prime minister* is projected down to the blend, but not the value of that role in 1938, Neville Chamberlain. Instead, Churchill is projected into that role in the blend.

We bring Hitler into the blend and he is the same Hitler in both spaces despite the fact he is a very, very different Hitler in the blend. In the blend, we have a Hitler who backs down.

These blends can be extended. Somebody could respond, “*That's only because Hitler was irrational: a more rational Hitler would have seen that his strategic chances were still excellent, and would not have backed down.*” In this blend, now we have Hitler blended with a more rational frame for a politician. Now we get a rational Hitler in the blend; this Hitler continues to push for domination, but this Hitler uses quite different reasoning than the historical Hitler did.

These blends can be reused, cited again, turned into input spaces in their own right. Margaret Thatcher, prime minister of the UK, talking about aggressors in the former Yugoslavia, called up the blend in which Hitler backed down, and said, in so many words, *if Churchill had been prime minister, Hitler would back down and we need to be just like that in Yugoslavia.* Her point was that we should not appease the aggressors; we should be strong in order to prevent atrocities of various kinds. Margaret Thatcher was prime minister, holding the same political office as Neville Chamberlain and Winston Churchill. Her blend was part of reasoning for sending a nation to war. We see in this case not after-the-fact cosmic justice or an explanation of the popularity of a racehorse like Seabiscuit, but rather reasoning toward national policy.

Here is a crucial point about frame blending. The purpose of frame blending is

sometimes to build a model that we can follow in the world. But sometimes the purpose is not to build a model that is possible for the world but instead to give us insight into the input spaces. Here is a gruesome example that illustrates this point. A woman in a coma and was in a nursing home. She was placed in the nursing home when she was 20 years old because of the coma. She was raped nineteen years later and became pregnant. The court had to decide what to do. This was a real human question of life and choice and future and judgment and justice. One side said that they knew that when the woman was nineteen, she was pro-life, meaning that she opposed abortion. A different view was taken by a law professor, who wrote in *The Los Angeles Times*, “Even if everyone agrees that she was pro-life at the age of 19, she is now 29 and has lived in persistent vegetative state”—that’s a coma—“for 10 years. Do we ask: ‘Was she pro-life?’ Or do we more appropriately ask: ‘Would she be a pro-life as a rape victim in a persistent vegetative state at 29 years of life?’”

Think of that blend. In the blend, there is a woman who is pregnant because she was raped while was in a coma. She is still in a coma but she’s now thinking about the fact that she is a person in a coma. In logic, this is what we call “P and not-P”—it means that two contradictory statements are both true. Blending is not restricted to making possible worlds: when the purpose of the blend is to give insight into the inputs rather than to make a model that we can follow, there is no reason that the blend cannot contain contradictory matter. We see this many times. A *reductio ad absurdum* in mathematics, for example, involves blending the accepted mathematical system with a particular proposition, for the purpose of constructing a space that one shows does indeed contain a contradiction. That is indeed the point, to show that the result is a self-contradictory system. Many such mathematical spaces have been constructed in practice, to good effect. The self-contradictory blend shows us something about the inputs, namely that the proposition is not to be included as part of accepted mathematics.

In the blend with the woman in the coma, she is thinking and aware of the fact that she is not thinking and not aware. This is P and not-P. No one dismissed this blend as pointless and silly. No one dismissed it at all. The purpose of this blend is to get us to think about the inputs, and it succeeds at that job. It leads us to think about, or at least it seems to have been intended to lead us to think about, what it means for a woman to have “choice” when she is in a coma. We don’t have to agree with any one side or the other in this case to see the usefulness of frame blending in trying to help us make decisions about quite crucial matters.

Seana Coulson found a wonderful example in 1992. It was written by Lee Ezell. To understand this, you need to understand that a “Planned Parenthood Clinic” has many goals—including helping those who want to prevent pregnancy, but also helping those who have an unwelcome pregnancy. Some who had unwanted pregnancies have gone to Planned Parenthood Clinics to arrange for an abortion. Here is the example:

I say thanks that no Planned Parenthood Clinic was available to me in 1963, when, as a virgin teenager, I was raped and became pregnant. The state of California would have been taking advantage of me in my crisis state by offering me this seemingly easy out. As an unwanted child myself, I decided abortion was too permanent a solution

to my temporary problem.

There are very many frame blends here. Notice that, if you know the right cultural frames, then you don't have any difficulty understanding this at all. No one thought this example was unintelligible or strange. No, putting together these kinds of elaborate mental integration networks with blend upon blend upon blend—this is just what we're built for. This is the easiest thing for us to do.

Remember that the mother who writes this letter is called Lee. Let's give her daughter a name—Julie. In order to understand this passage, you must construct the counterfactual blend in which Lee's mother did not have Lee. You need that space in order to understand that the woman who gave birth to Lee did not want the pregnancy or the child. Remember that Lee wrote that she was an "unwanted child." What Lee's mother wanted was for the mental space in which she was not pregnant to continue to be true to reality; instead, it became counterfactual to reality. In the space where the woman who became Lee's mother is not pregnant and does not become pregnant, there is no unwanted pregnancy, no unwanted child, no Lee. You also must construct the counterfactual blend in which Lee in 1963 was not raped and therefore did not have Julie. This counterfactual blend is the normal state, as opposed to what Lee calls the "crisis state." You also must construct the space in which it is 1992 and there are Planned Parenthood Clinics providing abortions to teenage rape victims. You need very many spaces; I won't describe them all. My point is that to follow Lee's argument, you must put all these different frame blends together and use them as inputs to yet other blends. Otherwise you can't get the right kinds of inferences. You don't realize the great complexity of blending that you are doing as you understand this passage, but so what? You don't see in consciousness the incredible complexity that goes into vision. Perhaps 50% of neocortex is implicated in vision—all that power, all that computation, all the time, necessary even for seeing a tree. You do fabulous amounts of work in backstage cognition. That is mostly the way it is with language and meaning. We are not built with our mental operations to be able to analyze language or meaning in a direct way. We are built to engage and deploy language, to construct meaning, not to analyze how we do it. Cognitive linguists are people who try to take mental capacities that we happen to have and trick them into looking at language and meaning.

Here is another frame blend about important and expensive matters. Justice William O. Douglas of the Supreme Court wrote the opinion finding that the Federal Urban Renewal Program was constitutional. This was a program that permitted the federal government of the United States of America to go into a neighborhood and knock down privately-owned buildings that were themselves just fine, and rip up roads that were themselves just fine. In the United States, this is something truly amazing for the government to be permitted to do. It is almost unheard-of. Governments can take over properties for the sake of eminent domain, and also to require substandard properties to be brought up to legal standards. But the Federal Urban Renewal program permitted the federal government to take over properties when it was not claiming eminent domain and even when one or more of the individual properties did

meet legal standards. Was this constitutional? Justice Douglas wrote:

The experts concluded that if the community were to be healthy, if it were not to revert again to a blighted or slum area, as though possessed of a congenital disease, the area must be planned as a whole. It was not enough, they believed, to remove existing buildings that were unsanitary or unsightly. It was important to redesign the whole area so as to eliminate the conditions that cause slums—the overcrowding of dwellings, the lack of parks, the lack of adequate streets and alleys, the absence of recreational areas, the lack of light and air, the presence of outmoded street patterns. It was believed that the piecemeal approach, the removal of individual structures that were offensive, would be only a palliative. The entire area needed redesigning so that a balanced, integrated plan could be developed for the region including not only new homes but also schools, churches, parks, streets, and shopping centers. In this way it was hoped that the cycle of decay of the area could be controlled and the birth of future slums prevented.

Justice Douglas is asking us to blend the frame of urban life—with buildings and roads and people—with the frame for crops and blight. When you have blight on a crop, what do you do? You must cut down the entire crop, wait for the infecting little organisms to die, and then replant the crop again. The *blight* frame is completely unlike the frame for the urban environment. If you were to connect the two by straight analogy, it would go like this: you would say, “Oh, I see. Over here in blight we cut down the whole crop so that all of the residents on the crop die and don’t come back; therefore, by analogy, over here in the city what we are going to do is knock down all the buildings so all the citizens die and don’t come back.” That is the straightforward analogy. But that is not what we want at all, not how we understand this passage at all.

In the blend, what you get is that you knock down the crop so that something similar but different—something healthy and differently designed and arranged—comes back, but not that you mean to eliminate the residents.

In the blight frame, you replant the identical crop again, in exactly the same design! But that is opposed to the principal meaning we get in the blend. We most certainly do not want the same design to be restored to the urban scene. The urban renewal will replace what is eliminated with something following quite a different design. The renewers will not rebuild everything exactly as it was. It is important to recognize here that this frame blend was the basis for a decision that allowed the federal government of the United States of America to do billions of dollars worth of alteration to America’s major cities, which it would not otherwise have been allowed to do. Frame blends are not merely fanciful. Blending is at the heart of what humanity does

We have talked previously about polysemy. Here I am going to sum up some of what we have said, in the context of frame blending. Expressions that are applied to an input can be projected to apply to counterparts in the blend. In this way, the blend harnesses existing words in order to express the new meanings that arise in the blend.

Polysemy Principle 1: Through selective projection, expressions applied to an

input can be projected to apply to counterparts in the blend. In this way, blends harness existing words in order to express the new meanings that arise in the blend.

So for instance, if you make a frame blend between a virus that attacks the body and something that does bad things to your computer, and in the blend you have this bad thing, you can now call it a “virus” because the word is projected down from the input. Now, you can say, “*I got a virus from your email message.*” This is perfectly fine. This kind of projection is one mechanism of polysemy. The word now can be used to refer to this kind of thing and it can become entrenched.

Consider “same-sex marriage”. Across the United States of America, there are movements to allow marriage between man and man or woman and woman. So “marriage,” which was something between a man and a woman, now gets projected down into the blend for a relationship that is not between a man and a woman. Try not to have an ideological reaction—we are talking only about the mechanisms of polysemy, of how there can be a change in the way the word “marriage” is used. Blending is what makes it possible for the word to come to be used to prompt for conceptual structure for which it once did not prompt.

Consider the word “number”. You begin with counting numbers—one, two, three, . . .” You can blend those counting numbers with points on a line, and with containers that have contents. So, in the blend, three has structure from the third in the sequence of counting, and a particular spot on a line of progress, and to three objects in a basket, a container. Interestingly, in the blend, zero becomes a possibility, because you can have a container with nothing in it. Notice the counting numbers themselves don’t have a zero. The steps of progress did not have a zero. But now, in the blend, zero counts as a number, although previously it was not a number.

Then polysemy goes to work. “Number” previously did not refer to nothing. But now, in the blend, the word “number” comes down from one of the inputs, and since zero has the right status in the blend, zero can be referred to as a “number”. Similarly, when we develop the blend for rational numbers, as outlined in *The Way We Think*, now these objects become “numbers”, through the same procedure of polysemy. And on through the list of real numbers, transcendental numbers, negative numbers, and so on.

When you combine real numbers with two-dimensional space, as we analyzed in *The Way We Think*, we can get complex numbers. The important point of this is that the internal structure of what a *number* is changes as you build blend upon blend upon blend. But you keep the word “number”. Complex numbers are very wild things, but now they can be called “numbers”, just like one and two. You don’t need new words to refer to an enormously new creative meaning. “Number” continues to be projected down from blend to hyper-blend to hyper-hyper-blend, and so on.

We saw polysemy with words like *father*. In *Paul is the father of Sally*, you have the kinship frame blended with the mental space for Paul and Sally. This is a simplex blend. In *Zeus was the father of Sarpedon*, we have an immortal father in the blend and

still call him “father”. This is an example of polysemy. In *Zeus is the father of Athena*, there is no mother! And Athena is born fully clad in armor from the brow of her father! This is extraordinary emergent structure, in dramatic clash with the kinship frame. But “*father*” in the blend can refer to that element, because it is projected down from the kinship frame. Here are some other examples of how blending provides polysemy:

- *Joseph was the father of Jesus.*
- *The Pope is the father of all Catholics.*
- *George Washington is the father of our country.*
- *Fear is the father of cruelty.*
- *The Child is Father of the Man.*

The same word, “*father*”, can be used to pick out quite different elements in all these blends, because it is in the input space, the kinship frame, and can be projected to its counterpart in the blend and then apply to that meaning that has been constructed in the blend.

Polysemy Principle 2: Combinations of expressions from the inputs may be appropriate for picking out structure in the blend even though those combinations are inappropriate grammatically for the inputs.

Consider again “number”. Before the invention of complex numbers or imaginary numbers, there were expressions like “square root” and “negative one”, but “the square root of negative one” had no meaning. It had no application. It wasn’t grammatical, because you could take square roots of only nonnegative numbers. Once we have the blend, *the square root of negative one* picks out something in the blend. We abbreviate the square root of negative one now by *i*. We see here an example of how combinations of expressions from the inputs can be used to pick out elements in the blend that had no counterparts in the inputs. If same-sex marriage is established widely, I will be able to say, without failing expectations attached to the *wedding* frame for “bride”, that *the brides married each other* at noon. That expression would be perfectly unremarkable. “Brides” could unremarkably, for everyone, regardless of ideology, refer to two agents who are marrying each other.

Polysemy Principle 3: We often have terms for emergent structure in the blend and can use them even though those terms cannot be applied to the inputs themselves.

You remember the Debate With Kant integration network, and the expression, “Kant has no answer for me.” That’s a perfectly grammatical expression for human conversation but it cannot apply to the input with the modern philosopher or the input with Kant centuries ago. But it can apply to the blend, because in the blend Kant and the philosopher are engaged in the debate.

Polysemy Principle 4: Blending routinely and inevitably extends the uses of words. In this case, the word itself can become a prompt to create an entire kind of integration network.

We saw this in examples like *safe*. “Safe” does not so much pick out a single new element in the blend as it prompts for the creation of an entire kind of integration network. “Safe” prompts us to take one scene and then use a frame of harm, blend them to create a scene of harm, a specific scene, understand that specific scene of harm

is counterfactual with respect to the original specific scene, compress that counterfactual link into a particular condition in our new, blended scene, namely *absence of harm*. “Safe” is a prompt to make a complicated kind of frame blend. We have analyzed this before.

Fictivity

Fictivity is a major component of human languages studied by a number of people, such as Len Talmy. Fictivity is a result of frame blending. Let’s look at Talmy’s fictive motion blends. Talmy points out that we say things like *the mountain range goes all the way from Mexico to Canada*. Now of course mountain ranges don’t go. They are static. And you can’t explain this as metonymy, by which we mean this: You can understand that there is a road down to the beach and that cars and people do indeed go down the road to the beach, so, the explanation would go, “the road goes all the way down the beach” is a metonymy, using “the road” to mean “things that go down the road.” But there is no such road along the mountain range and no organism or vehicle that goes along the peaks of the mountain range from Mexico to Canada, so the metonymy is not available. But you can imagine that there is a frame in which something travels from origin to destination, along a path. We use this frame conceptually very often.

Then there can be another frame, of, say, a mountain range. This is a frame for a remarkable geological feature. But the mountain range is static.

You can blend these two frames so that, in the blend, there is conceptually (but not actually or perhaps even possibly) something that moves along the mountain range with an orientation, from one end to the other. This gives us fictive motion. In this case, Mexico is mapped to the beginning point, the origin of the path. Canada is mapped to the endpoint of the path. You do not have to say “across the United States.” That is implied.

Talmy gives a taxonomy of types of fictive motion. Let’s look at a few others.

Access Paths. Consider *The bakery is across the street from the bank*. Now, *across* is a word from the frame of motion along a path. It can be projected to the blend, even though there is no motion in the static scene of the bakery, the street, and the bank. The path now corresponds to the street, but going perpendicular to the flow of automobile traffic on the street. It might be that it is possible to walk across the street from one to the other, but that does not matter: even if it is in reality impossible, because of a wall or a ditch or a fence, one can still imagine it, by projection from the frame of motion along a path. “Street” is a word we bring down from the static scene. It can be used to pick out the path. Notice that, interestingly, *the bakery is across the street from the bank* and *the bank is across the street from the bakery* should mean exactly the same kind of thing. In a way they do, but they ask you for very different construals of where you are standing and what you’re seeing in the blend.

Advent Paths. Consider *The palm trees clustered together around the oasis*. Nobody thinks, unless this is a scene from *Macbeth*, that trees can walk or that they can

move or that they can cluster. In cases like this, a static configuration is blended with a motion frame, in this case of coming to a place, advent. The result is that in the blend, the configuration of the trees is the end state of the motion of advent. Another example is, “*As I painted the ceiling, paint spots slowly progressed across the floor*”. In this case, there is actual motion: the paint drops fall. But that motion is not projected to the blend. Rather, a motion along a path on the floor is projected to the blend, and the configuration of the paint drops is understood in the blend as motion across the floor. Let’s review the details: You are moving your hand as you paint and the paint drops are falling. But the paint drops themselves don’t move across the floor. The fictive motion according to which they move across the floor is like the situation when you see a series of light bulbs go on, one after another, in a sequence, and you can’t help but see this as the light moving along a path. You can say “Wow, the light swept out a path” even though each of the light bulbs is static.

Shadow Paths. Consider *The tree threw its shadow down the valley*. This is a case where the static situation is understood by blending with the caused-motion frame, which has a path along which the object moves. Here, the shadow is blended with the object.

Fictive Interaction

Just as there can be fictive motion in blends, so there can be fictive interaction. The Debate with Kant is an example of fictive interaction, as is the Mythic Race, with Hicham el-Guerrouj. In none of these cases do the people who interact in the blend interact in the input spaces. But the blend creates a fictive interaction that helps us understand the network that is anchored by the blend. Catullus 101 is also a fictive interaction: there is no actual interaction between Catullus and his brother when Catullus’s brother is dead. But in the blend, there is a fictive interaction.

We have considered very serious situations—the U.S. federal government redesigning entire urban areas, a court trying to decide what to do when a woman in a coma is pregnant, Planned Parenthood Centers, political debates, and so on. Let’s go back to poetry and literature.

You recall that I mentioned *The Runaway Bunny*, a picture book for two-year-olds. In *The Runaway Bunny*, there is a little bunny, and the bunny’s mother. They talk to each other. Talking animals are the mainstay of every nursery, and are of course a frame blend. A human being talks; an animal doesn’t talk. A stuffed object is also a frame, and it can be added into the blending network, with the result that we get a stuffed animal doll that we pretend can talk. Children play with the stuffed bunny in the nursery. *The Runaway Bunny* does not represent the bunny as a stuffed toy overtly, but there are scenes in it that suggest a stuffed animal. Anyway, in *The Runaway Bunny*, the little bunny says to his mother, “*I’m going to run away*”. And the mother says: “*Well, if you run away, I will come after you, for you are my little bunny.*” And he says: “*If you come after me, I will turn into a fish and swim away.*” At this point, the book has a picture of the bunny jumping into the water. And the mother says: “*If you become a fish, I will become a fisherman, and I will catch you.*” And he says: “*If you*

become a fisherman and you catch me then I will become a flower high in the mountain.” And she says: “*If you become a flower high in the mountain, I will become a mountain climber and I will find you, for you are my little bunny.*” This is one frame blend after another, and the entire book proceeds in this fashion.

In fact, there is a passage of literature that works exactly like this, and that, for all I know, provided the model for *The Runaway Bunny*. Frederic Mistral, who was awarded the Nobel Prize for literature, published a work in 1858 called *Mireille*. He embedded in that work a Provençal song titled *O, Magali*. In *O, Magali*, there is a suitor who is singing to Magali up in her balcony. The suitor says, “*You will be mine.*” He is playing the guitar. She says “*No. No. No.*” He says, “*I will pursue you.*” She says: “*If you pursue me, I will become a fish in the sea and swim away from you.*” He says: “*If you become a fish, I will become a fisherman*”. She says, “*Then I will become a bird and fly away.*” He says: “*Then I will become a hunter and hunt you*”. She says, “*Then I will become a flowering herb in the wild.*” He says: “*Then I will become water and sprinkle you*”. She says, “*Then I will become cloud and float away to America.*” I love that! So he says, “*Then I will become the seabreeze and carry you.*” This pattern continues for a long time. Finally, she says, “*You’ll never catch me. I will go into the convent and become a nun.*” The suitor says, “*If you do that I will become a priest and I’ll be your confessor and hear you.*” And she says—this is the big one, “*If you pass through the portal of the convent, you will find all the nuns walking in a circle around me, because you will see me laid out under a shroud.*” But he does not miss a beat. He says, “*If you become the poor dead girl, I will therefore become the earth. And then I shall have you.*” Because if she is dead, then she is going to be buried in the earth. She then says: “*Now I begin to believe that you are not merely engaging in pleasantries with me. Here is my little glass ring for remembrance, handsome young man.*”

Let me make one last point. The frame blends that you see in *O, Magali* are very impressive, but in a way easy to see. Here is the fish, here is the girl, and they are blended. Here is the cloud, here is the girl, and they are blended. Again and again. Fine. But stay attentive and watch for blends because there is a different kind of blend that is happening right before our eyes in both *Runaway Bunny* and *O, Magali*—a blend for the purpose of rhetorical persuasion. What do I mean?

When you are a two-year-old child thinking about the future, there is nobody who can actually guarantee you a secure future. You have to live the future in order to know how it turns out. Things happen as you live; and it takes a long time; and it involves many agents. And when you are dating somebody, when you are thinking about getting married, naturally one is going to be concerned about whether or not it will all work out. Will the marriage be a good one? The only way actually to know is by running that future, and finding out. You can’t say: “Oh, I see. I’ll marry this man for 30 years and if it doesn’t work and I just go back to the way I was, when I was 19 or 26. There are some things you can’t undo, but you would like to know how it is going to turn out before you do it.

There is a pattern that we all use in trying to assess the future. We take the structure of a discourse interaction and blend it with the structure of the future. The

conversation can happen right now, between two or a few people. It can be relatively brief. It is already compressed to human scale. The conversation can happen at a café, or over dinner. The discourse interaction can take place in a little story that a mother reads to a child. We can take the structure and the dynamism and the psychology from that little frame of discourse and blend it with vast and expansive life events that take a long time to play out. This is what happens in both *The Runaway Bunny* and *O, Magali*. The little bunny in *The Runaway Bunny* and the beloved in *O, Magali* keep saying they are going to get away. Rhetorically, what they do is this: they say, I will become such and such. In this discourse, it is their right to throw themselves, by blending, into a new frame. So the bunny and the beloved both throw themselves, as it were, into the frame of a fish getting away.

This presents a challenge to the mother bunny and to the suitor: they must find a way to throw themselves, blend themselves, into the new frame, under the constraint that they must achieve a connection. That's pretty to do when the bunny or young woman decides to become a fish: If you become a fish, I'll become a fisherman. But notice the selective projection: the fisherman will not hurt, kill, cook, and eat the fish. If you become a cloud and float away to America, how will the inventive rhetorical challenge be met: I'll become the breeze.

Notice that rhetorically, inventively, the mother and the suitor stay with the child and the beloved. Every time they rhetorically escape to a new blend, the mother and suitor must find a way to insert themselves into that blend. For example, the suitor is not permitted to say to this rhetorical challenge, "Oh, don't be ridiculous. You never become a cloud." If he does that, what happens? Well, that's the end of his suit—the beloved has escaped. No, to meet the rhetorical challenge and show his ability and his dedication, he must stay with it. Both the mother and the suitor must show discursively and imaginatively and linguistically that they will stay with the person they are pursuing. And that tenacity and loyalty, if you will, gets projected into the blend of life—they are showing that they can deal with whatever is thrown at them and still keep the connection. Their rhetorical performance is blended with how they will behave in life. The little bunny throws a challenge at mother; mother will find a way. The beloved throws a challenge at the suitor; the suitor will find a way. That becomes persuasive or is taken as a signal of persuasion, because in the blend the structure of the discourse is the structure of the life. So the little bunny says in the end: "Shucks! I might as well stay here and be your little bunny." And the mother says: "Have a carrot." At the end of *O, Magali*, the girl says "I see now you are not just toying with me. Here is my little ring for remembrance". The ring is a symbol all in itself.

The frame blends that are in fact often the most persuasive and important are those we do not notice, like these. But they have enormous effect. They can indeed determine the course of a life, for both the little bunny and the beloved.

Thank you very much.

Lecture Eight Blended Stories

Thank you so much for inviting me to the lovely Beijing Foreign Studies University. It's been wonderful to be here. Thank you so much for coming and sitting in the bitterness of late afternoon in the heat at 4:30 when you could all be somewhere else dancing.


I have been talking about conceptual integration. For some of you, this lecture is the first you have attended. This morning, I talked about frame blending. Frame blending occurs when we integrate two frames.

Let's do a quick review, but with a different example, one involving language. Fillmore and Atkins have a classic analysis of the verb *risk*, and the syntax of the verb *risk*, and its meaning. Charles Fillmore, as you know, is the director of FrameNet and the linguist known for establishing the field of frame semantics.


Fillmore and Atkins give an analysis of the frame for *risk*. They don't quite put it this way, but what they say is consistent with saying that the frame for *risk* is a blend of the frames for *chance* and *harm*.

It is important to see that the frames for *chance* and *harm* are independent. If I say "Well, there is a chance that it will be 30°C tomorrow and a chance that it will be 31°C, but we don't care which", that's just means that there is a possibility of one or the other, but there is no question of harm. Similarly, if there is harm coming and the occurrence of this harm is 100% certain, then chance and possibility are not an issue. So harm and chance are quite independent as frames.

But when you integrate the frames for *chance* and *harm*, you get one of the basic frames for *risk*, in particular, *running a risk*. Fillmore and Atkins use the diagram I recreate here:



In this diagram, a circle means *chance*. There is a chance of harm. There is also a chance of something else. This is the structure of the basic frame of running a risk. But Fillmore and Atkins point out that there is yet another blend, one that adds in an additional frame, namely the frame of *choice*. You can *choose* to place yourself into a position where there is a *chance of harm*. This can be because you are betting on a horse, for example, or because you like the thrill of driving fast or something like that. Here is the diagram:



In this diagram, a circle means chance, and a square means choice.

In effect, Fillmore and Atkins have analyzed the lexical meaning of *risk* as a set of frame blends. They don't put it that way, but I think that is the best way to put it.

Importantly, blending is selective: you do not take everything from the frame of *chance* and everything from the frame of *harm* and everything from the frame of *choice* and put them altogether. For example, the *harm* frame automatically brings into an evaluator. If there is harm, there has to be harm to somebody who evaluates it that way. Think, for example, of a diamond and an owner. If the owner wants it cut, then the cutting does not count as harm. But if the owner did not want it cut, then the same cutting counts as harm. Of course, you might say that a connoisseur of fine diamonds, not the owner, might be caused aesthetic pain by the incompetent cutting of the diamond, even if the ignorant owner did not mind. But that means that then the connoisseur is the evaluator who sees the harm. But in no case is the harm actually to the diamond. It is to the evaluator.

Fillmore and Atkins talk about what they call “derivative syntax”. I would call it “blended syntax”. In these cases, the syntax follows the blending of the conceptual frames. Fillmore and Atkins consider, for example, the verb *smear*: when you smear something on a surface in such a way that the surface is covered by what you have smeared, then the verb *smear* acquires the syntax of *cover*, as in *I smeared the wall with mud*. In that case, the verb *smear* can be placed where *cover* would go. Similarly, when loading hay onto a truck results in filling the truck, then *load* can take on the syntax of *fill*, as in *I loaded the truck with hay*. You can always say that you *filled the truck with hay*, but when the loading results in filling, you can then say *I loaded the truck with hay*. The point is that when you blend these frames, you get selective projection to the blend, and this includes projection of the linguistic elements that attach to those frames. Accordingly, blended syntax follows frame blending.

Let me point this out in the case of *risk*. Fillmore and Atkins point out that when *risk* is *expose*, then *risk* can take on the syntax of *expose*. It can acquire the syntax for *expose*, as in something like *we'd have to reinforce the boat before risking it to the waves*. “Risk” can now occur in this syntactic position because *risk* has been blended with *expose* and “expose” can be used to say “exposing it to the waves”. “Risk” can acquire through blending the syntax for *investing in* something, as in *Roosevelt risked more than \$50,000 of his patrimony in ranch lands in Dakota Territory*. That's because you can blend *risk* with *invest* and then “risk” can occur where “invest” is possible.

These are all examples where emergent syntax arises because it is part of the emergent structure that arises from blending frames.

Human beings are highly adept at frame-blending. Part of double-scope blending is the blending of grammar associated with the frames being blended.

This morning, we were in a library, and I commented that we have a frame for *library*. A library has books. There are librarians. We can check the books out and so on. These frames are frequently culturally inflected, sometimes entirely culturally built. When we do frame semantics, frame analysis, and accordingly syntax, we must look at the cultural role and the historical role in building those frames.

I promised that in this afternoon's talk, I would discuss the blending of stories. *Framing* is crucial to higher-order human cognition, and we blend frames. *Story* is also crucial. *Story* may be the most basic mental operation for understanding. *Story* involves a small group of agents, not always intentional agents, interacting with each other in a human-scale scene. Sometimes we personify those agents; sometimes we don't. When the rain rains on us and we get wet, that's a small spatial story. Bigger stories are things like the legends about Dragon Boat Festival day.

Just as we blend frames, so we blend stories. Today, we are going to talk about blending stories. You've already seen one example of a blended story, the story of Phèdre and her stepson Hippolytus. Hippolytus was a beautiful man, who resembled his father Theseus when young, but his father Theseus is now traveling. Phèdre, who is much younger than her husband, Theseus, develops a passionate love for her stepson, Hippolytus. She is very restrained and tries to overcome her passion. But finally she decides to avow her love to Hippolytus. And she starts off by recounting the story in which Theseus went to Crète and saved the people by delivering them from the Minotaur. There was a monster inside a labyrinth, a monster inside a maze. And her older sister Ariadne helped Theseus. Theseus the hero came down, went into the maze, conquered the monster, and managed to get out, because there was a thread that Ariadne gave him to help him get out.

At one point, Hippolytus asks Phèdre: "Don't you miss your husband?" And Phèdre replies, in effect, "Sure I miss him. I miss him not the way he is now, unfaithful, running around the world, but the way he was when I met him." And she describes him in a wonderful term and says that "he was the way we paint our gods". That's already a blend. And then she says "or the way you look."

This is an analogy. It's not hard to imagine the aptness of the analogy, because of course Hippolytus is the adult son of Theseus. So of course Hippolytus looks something like Theseus when he was young. Phèdre remarks that Hippolytus has the courage, bearing, speech of Theseus when young. Then she says, "Why couldn't you have come down to Crète?" Of course that was impossible—Hippolytus wasn't even born yet—but she is imagining a story. In the story, Hippolytus is now blended with his father Theseus. That is, Hippolytus here in this scene in front of Phèdre is now blended with Theseus who went down and defeated the monster, defeated the Minotaur.

Phèdre imagines a blend in which Hippolytus does what his father did. She says, "You could have been the one to delivery us from the monster. You could have done it. You could have been the one to whom my sister gave the thread." We are doing

OK so far in this blended story, but now she takes a remarkable step. She says: “Oh no, I would have gone with you”. She means that, in this new blended story, she is to be blended with her sister Ariadne, so that in the blend, what Ariadne did, Phèdre does instead. She says to Hippolytus, “I would have helped you out. I would have given you the thread.” Then she says: “But that would have not been enough for your lover. I would have gone into the maze with you, to succeed or die in the maze with you. My love would have ...” Hippolytus is surprised at her words, because in the blend, Phèdre and Hippolytus are now lovers. Hippolytus asks, “Madame, have you forgotten that Theseus is my father?” And she says: “What makes you think I’ve forgotten him?” Her response is basically, what do you mean? this is just an analogy! And he says: “Oh I see. I am sorry.” And she says: “Oh, cruel one. You have understood me all too well.”

You’ve seen this before. This is an example of blending stories. There are two quite distinct stories here, with different characters. One of them is the story of Theseus and Ariadne and the Minotaur. That story has a place for the very young Phèdre, but she certainly did not go into the maze with Theseus. The other story is the story of Phèdre, Hippolytus, and Theseus. There is amazing projection to the blend. Only some things from the story of Theseus and the Minotaur are projected down into the blend.

There is also emergent structure in this blended story. In the emergent structure, Phèdre and Hippolytus are lovers. They are not lovers in either of the input stories, but they become lovers in the blended story. Of course, that’s exactly what Phèdre wants to communicate to him. He understands that the blended story can be a comment on the scene that he and Phèdre inhabit.

There are often things people cannot easily say directly to each other. So they hint. One of the ways of hinting is by constructing a blend that has structure in the blend that allows them to make utterances of the blend that they would prefer to make, if it were not unsuitable or difficult, of one of the inputs to the blend. Then they can utter the expressions without its being claimed that they have uttered those expressions about a particular scene or content. In films, for example, there might be characters that are a high school boy and girl, where the characters are actors in a high school play, and the boy is playing Romeo and the girl is playing Juliet. They the boy gets to express with his body and his words his undying love. He can pretend that the expressions are in the blend in which we have the boy/Romeo and the girl/Juliet, but the boy means the expressions to apply to the input space where there is a boy and a girl but no Romeo and Juliet.

In the blending network whose blend is the story of Phèdre and Hippolytus in the labyrinth, there are input spaces that Phèdre connects with standard vital relations: identity, analogy, cause-effect, part-whole, time, space, all of the usual connections, which by now we have seen many times.

There is selective projection to the blend: only parts of each of the input spaces come into the blend. There is emergent structure. In the blend, it is Hippolytus who conquers the Minotaur. In the blend, Hippolytus and Phèdre are now lovers. And as we mentioned many times, there is now language that can apply to the blend that is not

appropriate to the two input stories. So she can say, relative to the blend, “*your lover*”, “*votre amante*”, of herself, because in the blend this term applies, even though it does not apply to either of the input spaces.

There is a poem by William Butler Yeats, a lyric poem in the voice of a man who is about sixty years old and is visiting a school. He is a public official and he is coming to visit the school and say nice things to the children. That’s one story: here’s a sixty-year-old public official visiting a school. And in another story there is the woman that he loved passionately when he was young. These are two very different things, two very different stories.

He begins talking in the voice of the man who is in the school: *I walk through the long schoolroom questioning; A kind old nun — the nuns, this is Ireland, the nuns are often the teachers — A kind old nun in a white hood replies; The children learn to cipher— that’s means to do arithmetic — and to sing,*

*To study reading-books and history,
To cut and sew, be neat in everything
In the best modern way — the children's eyes
In momentary wonder stare upon
A sixty-year-old smiling public man.*

So here is the sixty-year-old smiling public man and the children look at him. He is visiting the school.

He is thinking, however, of a very different story: the story of the woman he loved:

I dream of a Ledaean body

That is a reference to the myth of Leda and the swan. Leda was a beautiful woman. Zeus visited her in the form of a swan and mated with her.

*I dream of a Ledaean body, bent
Above a sinking fire, a tale that she
Told of a harsh reproof, or trivial event
That changed some childish day to tragedy --*

So he is remembering the woman he was in love with. Now it happens that she, the beloved, is remembering a story from her childhood in which perhaps she was punished or criticized or corrected, and something harsh was said to her. We now have three stories: the story of the sixty-year-old man in the school, the story of him and his beloved, and the story from the beloved’s youth of the harsh reproof.

*That changed some childish day to tragedy --
Told, and it seemed that our two natures blent
Into a sphere from youthful sympathy,
Or else, to alter Plato's parable,
into the yolk and white of the one shell.*

That last bit is quite complicated; we won’t go into that. But here comes the blend. He says,

*And thinking of that fit of grief or rage
I look upon one child or t'other there
And wonder if she stood so at that age—*

At this point, we have a very robust analogy: Could the beloved have stood like that? The man didn't know the beloved when she was a child.

*For even daughters of the swan can share
Something of every paddler's heritage—*

That means even really exceptional women could have been, when young, like normal girls in some ways.

And had that colour upon cheek or hair,

The analogy is here growing tighter between the child who later became his lover and the school girls he is now looking at.

*And thereupon my heart is driven wild:
She stands before me as a living child.*

The blend comes in with full force now: the child who grew into the woman the man loved is blended with one of the living schoolgirls, and now in the blend, the woman, or the child who became the woman, is alive and before him. The great network of his life is now compressed down in some salient ways into a person who is standing in front of him. The beloved is here, in front of him.

Notice the emergent structure. Nobody to whom I have ever shown this poem has interpreted it as indicating that he wants to have sex with a little girl. You see, everyone laughs when I mention it. That's good! But why do we not interpret it that way? Notice that we do not project the sexual passion from the space of the man and the woman down into the blend.

The next crucial thing to notice is that now he can use a word like "She" to refer to the element in the blend that is alive and in front of him. But who is this *she*? *She* applies to the lover up in the story of the man and the woman, or the child who became that woman. It doesn't apply to the schoolgirl in the space with the sixty-year-old man and the schoolchildren. But it does apply to the blended element in the blended story who is also the schoolgirl.

The poet can say, "*She stands*". Notice that this is present tense. What is strange about this? Please answer. What is strange about using the present tense for his lover? That's right—the beloved does not exist simultaneous with the space of the sixty-year-old man, that is, the space of the temporal viewpoint. But he can use the present tense "stands" of his beloved because in the blend, the beloved is also the schoolgirl, and the schoolgirl is contemporaneous with the space of temporal viewpoint.

He also uses the phrase "a living child". That would be an odd expression to use in the story of the sixty-year-old man and the schoolchildren. Suppose the sixty-year-old man said to the schoolchildren, "I am so happy to see you living students." That would be quite odd. It might sound like an expression said by someone whose English is poor. But "living" is quite acceptable here, because up in one of the input stories, the beloved is not living anymore, and "living" emphasizes the disanalogy between the blend and that story. The grammar in "*she stands before me as a living child*" is legitimated by the blend. The grammar picks out things that are in the blend. And you know how to map the blend back to all of the input stories in the conceptual

integration network.

Even in a spectacular, pyrotechnic case like this—which involves stories of the sixty-year-old man and the schoolchildren, of Leda and the Swan, of the man and the woman as lovers, of the girl who received the harsh reproof—there can be many elements of the blend that we put together but that we do not particularly notice. The grammar that becomes available—*she stands before me as a living child*—does not leap out as a blend, but a blend it is. Notice how I can make all the students in this lecture hall laugh by greeting them as “living students”. That expression seems very odd, funny. But when “living child” is used of the blend in the poem by William Butler Yeats, it provokes not at all laughter, but rather something deep. It works because all the backstage conceptual integration to make it work has already been done by you.

Notice how I could make you laugh by calling you *living students*. But it seems perfectly fine to you because in the back stage of our cognition, you’ve already done all of these work. Similarly, we saw how language comes up for the blend in the scene in which Phèdre, speaking to Hippolytus, refer to the element in the blend that corresponds to herself as *your lover, votre amante*. This morning, we discussed how Catullus could refer to *mute ashes*. That is language that is not grammatical for ashes but is fine for the blend—no problem.

In a few minutes, I will go back to talking about the blending of stories. But to move forward in that analysis, we need to look at something very basic that plays a big role in blended stories. We need to look at how we understand other minds, and even how we understand our own past and future minds. Understanding other minds often goes in cognitive science under the term “theory of mind”—but that label is misleading in two ways. First, it is taken as concerning only how we understand the minds of other *human beings*, when in fact, projection of mind is not nearly so restricted. Second, it is taken as having to do with understanding only *other* human beings, but in fact it has to do with how we understand our own past and future minds. So at least we would need to supplement *theory of mind* with *theory of self*, but I will for the most part let these misleading labels go and instead refer to *projection of mind*. Projection of mind is made possible by double-scope blending.

Suppose you look at a seal. There is a seal, in the water. It is a water animal. Fine, no problem. The seal is analogous to you, and you can recognize the analogies. It has eyes and the eyes are open and they have a direction. And they are active. You cannot see that the seal has a mind. You can’t see that. You can’t feel that. But you know that *you* have a mind. When you look at something with your eyes, you know that you are seeing something and you are paying attention to it. Consequently, you can make a blend—without ever recognizing that you are doing anything at all—of you and the seal. You are one input and the seal is another, and they are blended in the blended space, so that, in the blend, the seal has a mind. Accordingly, in the blend, the seal’s eyes are not just *open*, and *round* and *active*, but they are additionally *inspecting* and *attentive*. They have *intention* behind them. They have *perspective*. Notice that you can’t see or perceive *inspection*, *curiosity*, *intention*. You can’t even see that that the seal is *perceiving*. You cannot see that the seal has *goals*. But in the blend there is a seal with a mind. Let’s go slowly over this. In this mental space, with the seal, the seal

has features you can perceive. And the seal in that mental space has analogies with you: two eyes, a head, the turning of the head, the opening of the eyes, and so on. But how does your perception of the seal turn into an understanding of the seal as having a *mind*?

The answer is that the seal is blended with you. Now, in the blend, quality of mind is projected from the space with *you*. And now the seal in the blend has a mind, a rich mind. This is the easiest thing in the world for you to do, but that is only because you are a double-scoping cognitively modern human being. Other species are really terrible at this, even though it seems inevitable to us.

Even when I point out the blending that is needed to understand the seal as having a mind, it seems overly-complicated. Why should we even need any explanation at all? Isn't it just obvious that a seal has a mind? Am I so crazy as to think that a seal does not have a mind?

Of course I think a seal has a mind. There are many blends that we reify, that we regard as true of the world. I do not think that the blend is false in this case, although it is a scientific enterprise to investigate just how much of my own mental ability I should be projecting into the blend for the seal. It is very easy for human beings to overproject not only to animals but even to each other.

The scientific hypothesis of a mind for the seal is a blend. This is the mental procedure according to which scientific hypotheses are created. And we think that this blend is true. In the blend, the seal can perceive, and we think the blend is accurate to this extent.

In fact, this projection of mind is just what we do with each other. The seal has a perspective, for example. I think you have a perspective. I can't see from your perspective but I can imagine it. I can see from my perspective. For example, I want you all right now to imagine that you are me and I want you to see yourselves, in imagination. Go ahead, you can do it right? In the blend, you can be you with your mind and your location but my perspective. You can create this meaning through blending. You can take on my perspective in the imagination. If you couldn't do that, you would be very bad at social cognition. If you could not do that, language could not work the way it does, because much of language depends upon being able to assign perspectives, often ones not our own, often ones that are not even available in our local environment. I can say, "In Paris, from the Institut de France on the Left Bank, you can look across the Pont des Arts and see the Louvre on the Right Bank." But we can do much more.

Now let us look at a different blend of a seal and a human being. We can do selective projection from the seal and from the human being so that in the blend we create not a seal with a mind but a *selkie* with a mind. I didn't know what a selkie was until I was in my thirties. A selkie is a legendary being from the Orkney Islands in the United Kingdom. Selkies are shape-shifters. There's a large category of shape shifters in mythology and in legend. A selkie is like a seal. It looks like a seal. But it can take off its coat. And when it takes off its coat, wow, does it look great—it looks like a human being, and in most legends, a very attractive human being! That's at least how the legend usually goes. The story is told of how selkies dance with each other on the

level shore in the moonlight, and the really attractive, really, really attractive male selkies take off their coats and slip into the villages and have sex with the human females, who are really, really, really happy about it. In the Orkney Islands, there are legends about how if you have a child whose toes look a little webbed, then it probably had a selkie for a father. Indeed, human beings sometimes do have webbed toes.

In the selkie blend, we have a being that is a blend of a seal and a human being, but with different projection and different emergent properties. In the blend, just as human beings can put on and take off clothes, so the selkies can take off their coats. Selkies have inspection, attention, intention, perspective, just as we saw for the seal in the blend. But selkies in the blend have much richer projections from the mental space of the human being. They have some of what the seal has and some of what the human being has, and emergent structure not available to either the seal or the human being. The seal understands the sea and the elements, and knows the weather. So do selkies. Selkies can talk.

We are very familiar with such blended talking animals who are quite unlike either animals or human beings. Consider the different blends of *dog* and *human being* we have for Goofy, Pluto, and Scooby Doo. These are three different blends of *human being* and *dog*. As Gilles Fauconnier has pointed out, Pluto doesn't talk at all—even though he has spectacular abilities to understand speech and to communicate—but Goofy talks in a goofy way. Scooby Doo talks in a kind of human voice with doggish overtones. Donald Duck of course is the standard example of a talking animal. No dog or human being talks like Scooby Doo or Donald Duck, although a talented human being can and did speak that way. This is in fact how the voice Donald Duck was manufactured for the cartoons, by having a human being talk in a strange way.

This selkie has a removable coat which is something like the clothes that human beings have. It's a coat that the selkie puts on. When the coat is on, the selkie looks like a seal and can swim like a seal. But when the coat is off, the selkie looks like a human being and can't swim. The selkie without the coat is not exactly a human being, of course, because an actual human cannot put on a coat and turn into a seal. And the selkie with the coat on is not exactly a seal because seals cannot take off their coats and turn into human beings. The selkies are vulnerable when they take off their coats, because they are stranded on land. They are vulnerable, in something like a way a human being is vulnerable when naked—you can't go walking in freezing weather, for example.

So here comes the myth. The myth is that some Scottish Orkney Island human males would like to capture the coat of a selkie female so that she will marry him. If he has her coats, she can't go anywhere. This is quite a manipulative scenario. There are stories about how a man falls in love with a selkie and tries to get her coat and promises her that if she has children with him and stays with him—selkies live a very long time—, then he will return her coat to her at some point.

In the selkie blend, or what I will refer to as “Selkie World,” there are selkies. Selkie World is a blend of stories—the story of the seal's life and the human being's

life. But there is a great deal of emergent structure in the blend that is in neither of the input lives, and it is common when we establish such a blended world—like the Land of Oz or the Through-The-Looking-Glass world—to come up with an explanation for the origin of this emergent structure.

We do not explain the origin of selkies the way we explain the origin of seals with a mind. The explanation for the origin in the blend of the seal with the mind is to assume that the structure in the blend is in the input. That is, if we ask ourselves how seals with minds originated, we just respond: “What do you mean? Seals have minds already?” We do not notice the emergent structure for the seal that is in the blend, and so we assume that it is already in the input. We think the seal has a mind in the blend because it already had one in the input. This is just what we do with each other. Why does another human being have a mind in the blend? We assume that it is because the human being already has one in the input, although we cannot perceive that mind. So if I ask, “Why do I think you have a mind?”, it is easy to think that the answer is because you do have a mind. Indeed, you do, but the way I understand it is by blending. I have many secondary scientific reasons for thinking you have a mind—your brain looks like mine, for example—but of course the way I can conceive of your having a mind is by blending. Blending is a road to mathematical insight and scientific truth.

But the case is quite different with the selkie blend. You don’t say, if asked why selkies can take off their coats, “Because seals can.” You don’t think seals can take off their coats, so you can’t make the mistake of thinking that the emergent structure in the blend is already in the input with the seal. You can’t just say, “Oh, selkies can take off their coats and become human beings because seals can.” You are not content with that explanation.

Instead, you can say, “I don’t need an explanation at all. It is just a fantasy.” Or you can come up with some other standard kinds of explanations for emergent structure for imaginary species. You can claim that there was an event that was causal for the emergent structure. Consider Spiderman. How did a person become *Spiderman* and acquire all of those strange abilities. The answer in the story is that the person was bitten by a radioactive spider, and this produced changes in the person to transform the person into Spiderman. Similarly, there is a talking dog in a story—*Martha Blah-Blah*. If we were to ask, “Why can Goofy the Disney Dog talk?”, the answer would be, “It’s just a fantasy—the Disney world is one where lots of animals talk.” But the reason Martha Blah-Blah can talk is that she eats alphabet soup, and the letters go not to her stomach but to her brain, so now she can talk. These stories are often full of little explanations like this one for how these kinds of animals come about.

In the case of selkies, there are three explanations for their origin. One explanation is that if a human being dies by drowning, it becomes a selkie. The second explanation is that they are transformed according to cosmic justice. (You saw divine justice in the case of Bertran de Born.) According to the cosmic justice in Selkie World, if you do something bad, you can be damned to become a selkie. A third explanation is that selkies are fallen angels. When the rebel angels were ejected from heaven, they fell. And those that landed on the earth became fairies and those that

landed on sea became selkies. You can invent other explanations in the blend if you like.

Now that we have Selkie World in place, let me tell you about a story that makes use of Selkie World. This is a story for children. It is called “Aunt Charlotte and the NGA Portraits”. In the story, there was a woman named Olga Weathers. She lives in Ocracoke, which is an island in the outer banks of North Carolina in the United States. It is quite remote, with sand dunes and the sea and waving wind. About half way through the story, you figure out that Olga Weather is a selkie. It takes that long. The word *selkie* is never used in this story. And when I read this story, I had never heard of the legend of selkie, so I thought this being was invented by the author. You don’t need to know anything about the selkies to follow the story, because the story introduces you to the Selkie World blend itself.

It so happens that a bad man stole Olga’s coat, to get her to marry him, but she wouldn’t marry him because she knew she would never get her coat back from this particular man. This man, knowing how capable Oglia was, had to work very hard to hide her coat from her. He hid it in a picture. What can that mean? I mean the following. This will take a little while to explain.

In this story, there is not only Selkie World, but also Picture World. Picture World is another world that arises through blending. In a minute, I will discuss Picture World and how it arises, imaginatively, through blending, but here let me say where we are going: the world of “Aunt Charlotte and the NGA Portraits” is a blend of the two blended worlds, Selkie World and Picture World. There is emergent structure in this hyper-blend not available from the two other worlds.

Let us talk about Picture World. It’s very common for us to blend a represented element with its representation. We saw this in *Harold and the Purple Crayon*, where the real moon that gives light is blended with the drawing of the moon. When he draws the moon, the drawn moon gives light. In the blend of Harold’s world, there is now a unique *moon* element that compresses the moon and the drawing of the moon. When Harold draws something, it is real. *Harold and the Purple Crayon* is a popular book for three-year-olds. Children have no difficulty with this imaginative kind of blending, because they are born to blend, according to the constitutive and governing principles of blending.

Such blending of a representation and the represented element is quite standard, and not limited to fantasy at all. If I show you a little flat, two-inch by two-inch object that doesn’t move at all and I say, “this is my son”, you have no problem. The little flat, square object is a picture. You are not deluded. You know that this is a blend, in the sense that you do not think you can talk to the picture and have it talk back to you. The projection to the blend is selective. This is a compression of a representation link and an analogy link, because the picture of the face looks something like his face. The representation and the represented object share some topological features. Notice that I do not have to say, “*This is a picture of the head of my son.*” I can say “*This is my son*”, because in the blend, it refers to the blended element. No problem. And you may think that you are impervious to the feelings that come up here, but that’s not true. I have a friend who had a picture of his son, a young child, on the desk. The picture was of the face of the child in kindergarten, and the photo was less

than a year old. The picture fell over for some reason, fell over face-down. This friend immediately snatched the picture up to see if the glass had broken, did not even wait to do it carefully so that he might not be cut by shattered glass. Now, why? He had many copies of this picture, and if the glass broken, glass is very cheap and easy to replace, so what's the problem? What's the big deal? The big deal is that in the blend, that was the face of his kindergarten child under the glass. Just think about pictures that you have; how you feel when you look at the pictures. People talk to pictures. You are not crazy. Icons evoke similar behavior. Look at icons and other kinds of representations. These representations compress networks to human scale and so make various immediate emotions available. We are not deluded, but we are strongly affected by the blends that compress representation and analogy links in the outer network.

In Olga's world, there are paintings. The paintings in her world are just the ones in our world, because her world is our world, mostly. The paintings in her world have names, and they are just the names that they have in our world. The difference is that in Olga's world, all the paintings are in Picture World—the painting is a blend of the painting in our world and what the painting represents. Not everyone in Olga's world recognizes that the pictures belong to Picture World. In fact, very few see that. In Olga's world, if you look at a painting very closely, and you are also a very intelligent kind of person, sensitive to emotion, sensitive to patterns, then you can start to see the people in the picture moving. You can see the wind blow, and the trees. You can see the waves move.

This is not so strange. Look out a window. Just glance out the window quickly. What you see is a framed glimpse of the world, and it is very analogous to a painting in certain ways. One of the inputs that we incorporate into this Picture World is our knowledge of framed glimpses. Or look out any portal, like a door. Just glimpse. What you see is a framed glimpse. One of the things you know about windows, doors, and apertures is that what you see is just part of the world on the other side of the aperture. You know that the world goes on past the boundaries of the window, the door, the frame, the aperture. When we look out the window, we see a building, but we know that there are more buildings that we cannot see because they are outside the frame. People can walk past the window or the door, and when they disappear from our field of vision, we do not think that they disappear from the world. They are somewhere in the rest of the world.

Another thing we know about windows, doors, apertures, frames is that we can go through them ourselves. I can walk through the doorframe, or go out through the window frame.

Project this knowledge into Picture World. If we blend the picture with the notion of a framed glimpse, then there are things in the Picture World, and they can move, and the world extends beyond the limits of what you can see in the frame, and you can go through the frame, into the Picture World, just as you can go through the doorframe and out into the corridor of the hall.

This structure in Picture World is used often in television, films, and literature. In these stories, elements can come out of the television or the film into your life, into

your living room, or you can chase them back into the television or the film. This is a standard trope in fiction.

Notice that when I show you a photograph, of say, someone standing on part of a bridge, and the photograph shows only what it shows, you assume that the bridge continues beyond the limits of the frame. You don't see it, but you assume that it is there. Notice also that you assume that part of the bridge is *behind* the person in the picture. You don't see it, but you assume that it is there, because that's just how it works in the real world. This tendency to complete the representation is very thorough, even though we do not notice it. Consider that we assume that the person on the bridge, who is facing us, has a back, even though we do not see that person's back.

You already complete to an extensive degree the world for the representation, and this is true even for a painting that you know to be fictional. Suppose somebody paints some green alien creature standing on a bridge on a fictional colony on Mars. You know that none of the represented items, except Mars, exists. But you complete it. You conceive of the bridge as extending *behind* the creature standing on it, even though there is no representation of that part of the bridge. I can say, "What do you think is standing behind this green monster?", and that question will be intelligible.

So, what did I mean when I said that the bad man hid Olga's coat *in a painting*? I meant that in Olga's world, you can actually take the selkie's coat and enter the picture and hide it in the painting. This is what the bad man did. The coat was in Olga's world but not in Picture World. He moved it into Picture World in the sense that it was physically moved from outside Picture World inside Picture World. In fact, he hid it in a painting by Canaletto of Venice.

Olga Whethers did an awful lot of work to find it, including an awful lot of art historical work. The bad man in the meantime hunted in our world, and got a wound from a narwhal, which turned septic. And he died.

Why can't Olga just go into the painting once she has located the coat and bring it out? The answer lies in the fact that Selkie World in "Aunt Charlotte and the NGA Portraits" is not like the standard one in the myths of the Orkney Islands. When she takes off her coat, she is not young and lithe. Instead, she is quite ample, like a large seal. She is stout. The frame of the picture is something she cannot physically get through. Even worse, this painting has the Grand Canal of Venice across the bottom, so if she did go through the picture frame, she would fall into the water. Remember that a selkie without its coat cannot swim. She would drown. Think of that. It is really very inventive to manipulate the selkie with fear of drowning, because swimming is what they do best when they are wearing their coats.

So what is Olga going to do? Selkie World and Picture World are already blends with our world as one input. Olga is in a blend that looks like our world but is a blend of Selkie World and Picture World. In this world, there is a young girl named Charlotte. She never has any friends and she is lonely, and her parents love her but they ignore her, because they are busy. She lives in Washington, D.C. She has to go to stupid tea parties with her mother. She loves her parents, but her situation is not working out ideally for her. Charlotte is on vacation with her parents in Ocracoke, but not during the tourist season. There is very little activity on the island for a young girl.

Olga finds Charlotte on Ocracoke, and discovers that the child is a very inventive. Charlotte is actually quite creative; she is good at jigsaw puzzles and otherwise at perceiving patterns in complexity.

Olga and Charlotte become friends. Olga wants to protect Charlotte. But finally, Charlotte can see that the painting is moving—the Canaletto painting, which Olga has acquired and hung in her house in Ocracoke. Charlotte has learned from Olga how to see even better than she usually sees. Olga at last allows Charlotte to go into the painting to try to get the coat back.

Charlotte does just that. She puts her foot upon the frame, which feels solid as a rock banister — and the reason it feels as solid as a rock banister is that it is a rock banister in the blend; it is the edge of the bridge. Charlotte throws herself in, lands in the water of the Grand Canal, and swims over to the side of the canal.

The trouble is that she can't get out, because the edge of the canal is stone and rather higher than the level of the water.

But there are other interesting things in this world. For example, one of the things you learn about this world is that although it has all of the physics of our world, the people in the painting can't see or hear or otherwise perceive her. She is not there *for them*. Why? Because she was not painted. Some people help her up the side of the canal. But how can they see and hear her? How can they help her up, if nobody can see her? This will take a second. There are analogies across different paintings. All the different paintings can be blended to form a kind of picture world, one with some boundaries, but not absolute boundaries. It is possible in this world for people from one painting to move into a different painting. So Olga has located some children in other paintings and persuaded them to go to the painting that Charlotte has entered. These children have names: Celeste, Antonio, Caroline, and Rannuccio. They come from other paintings of the Renaissance. Celeste, Antonio, Caroline, and Rannuccio pull Charlotte up out of the canal.

Charlotte finally figures out that they come from other paintings. The children in the story discuss what a picture world is like. So for instance, if you eat a chicken and you go back, the chicken is still there. If you penetrate into the buildings too far, the world turns into a great mist, because as Rannuccio says "it is just a picture". It is not a full world.

So these children in the Canaletto painting go looking for the coat. They find a door with a doorbell and a number; it has a number: 5478-B. Charlotte, who is a great puzzle master, understands that this is not from the original painting. This is not part of Canaletto's Venus. She starts to enter, but the other children find they cannot follow her, because this place is not from the world of paintings, and they cannot leave that world. "This is not the painting," they say.

Charlotte goes through the door. Inside, she finds a room full of animal trophies: trophies of the hunt. There is Olga's coat. She picks it up and is about to leave.

Then a harlequin, in a chair in a corner, whom she had not previously noticed, wakes up, and says, "What would you want with that old thing?" This is an amazing blend. The harlequin is a blend of the Venutians with a being that can read her mind. Perhaps the harlequin is in a way a projection of her own mind. He says, in effect,

“Listen, listen, you don’t want to go back home. Back home you have no friends. You have to go to those boring tea parties. And you know that woman Olga. She doesn’t really love you. She is just using you. Why don’t you just stay here? You could have friends.” He is extremely persuasive. But he makes a mistake. He says, “You could live forever.” And suddenly, Charlotte realizes that she doesn’t want to live forever; she wants to grow up; she wants to go to college. This is a shocking thing for her to think, because the period in which this story is taking place on Ocracoke is somewhere in the relatively early twentieth century. Maybe Charlotte had implicitly thought she wanted to live forever until the thought was actually expressed by the harlequin. She rejects it. She scurries out of the room. She takes the coat to Olga. Olga turns back into a selkie and goes into the sea.

What’s the point of all this blending? Well, some of the delight is fantasy. But there are serious inferences from this blending of stories. This story is being told by Charlotte when she is somebody’s great aunt and lives in a townhouse in Washington, D.C. The story is titled “Aunt Charlotte and the NGA Portraits”. The narrator is a girl named Marguerite, who has gone to visit her great-aunt Charlotte. Within the story, Great-aunt Charlotte tells Marguerite the story of what happened to her when she was a child on Ocracoke. Charlotte in the story of Ocracoke is about the age Marguerite is when she visits her great-aunt Charlotte. Marguerite, you see, is not getting along very well with her parents, and doesn’t have many friends, and so on. She is kind of lonely. There are many analogies between Marguerite and the young Charlotte. Marguerite visits her great-aunt Charlotte, who seems rather remote. But when they are taking a taxicab ride down to the National Gallery of Art, the one on the mall in Washington, D.C., Charlotte tells her story to Marguerite. “NGA” stands for “National Gallery of Art.”

When the story of Ocracoke is over, Marguerite arrives with her great-aunt Charlotte at the NGA. By this point, it is clear that the elderly great-aunt Charlotte is quite wealthy. She had married a man who did very well, and he loved her, and he bought various paintings for the National Gallery of Art. Indeed, he bought the ones that contain the children Charlotte, in the story of Ocracoke, meets inside Canaletto’s painting of the grand canal of Venice, the one that, in the story of Ocracoke, contained Olga’s coat. In fact, the paintings of the children mentioned in the story of Ocracoke do exist in our literal world, and most of them do belong to the National Gallery of Art.

So, you see, there is another essential blend that is constructed here: the blend of Marguerite, the child, and Charlotte as a child, because great-aunt Charlotte is using the story to say to Marguerite: “Listen, you can have a life. You can grow up. You can imagine.”

And naturally, as is almost always the case, there is another overarching conceptual integration network, one in which Marguerite and Charlotte as a child is blended with the child reading the book. It is a children’s book. In the blend, the reader of the book can explore what it means to choose a life.

It is important to realize that this kind of blending upon blending of stories is what we in the United States would call “duck soup.” I will bet that “duck soup” does

not mean the same thing in China! “Duck soup” in the United States means “something that is very easy to accomplish.” Human beings create these blends upon blends of stories without hesitating, without difficulty. Just read children’s books or watch cartoons—blends upon blends of stories, and “duck soup” for the child.

It’s blend upon blend upon blend upon blend, selective projection, emergent structure, language that only makes sense because it refers to the blend, blends that are prompted by certain kinds of linguistic constructions. Tomorrow, I will be talking about technical aspects of language. The kind of multiple blending we see here is indispensable for language. Double-scope blending is indeed what makes language possible. The blending I will discuss tomorrow is mostly the kind no one even notices. Here, when we see these fabulous blended stories, it is easier to see that the blending is occurring. Blending is not a different or special kind of thought. It is the basic kind of thought that accounts for human higher-order cognition. It is our main method of creativity. It is why we can have cultural invention. It is why we are able to have language in the sense that without it we wouldn’t be able to have language. It is why we have art, music, mathematic insight, scientific discovery, fashion of dress, advanced social cognition and these kinds of things.

In this lecture, we began with a discussion of the *risk* frame, and its blending of *harm*, *chance*, and *choice*. This was not very fanciful. I mean, you use the verb *risk* routinely, without ever wondering why you can say *risk* without a particle versus *risk in* versus *risk to* versus *risk on*. We were talking about mental events below the level of anything that human beings are disposed to notice.

We notice blending only when something goes wrong, or more often, when an artist has done considerable work to put blending onstage. It is difficult to show blending to human beings. It is difficult to get them to notice blending, because for them it is designed to occur almost entirely in backstage cognition. Many of these mental operations cannot be called into consciousness any more than you can call vision into consciousness. Blending is very complicated, too complicated for consciousness to follow. Still, following stereotypical thinking, you might be disposed to say, “Well, it’s just a kid’s story. Kids are, you know, very imaginative, and this story of Charlotte and Olga and Marguerite is just a kid’s story.” I would reply that I think that comment is fundamentally in error. Children can perform this pyrotechnic blending very well, in children’s stories and in learning how to use verbs like *risk*. The fact that they can do it so well so early should tell us something central about what makes us human.

Just to make certain we are on the same page, let us look at a passage not from a children’s story but rather from the high canon of literature. This passage is from *King Henry the Sixth, Part One*, act 4, scene 7, by Shakespeare. In this play, Lord Talbot is in France fighting against the French. Lord Talbot is English. His son, John, has come over to France from England against the wishes of his father. Lord Talbot and his son have been apart for seven years. They end up in a very fierce battle with French soldiers. The French are dominating the field. The father, Lord Talbot, tries to persuade his son, John, to flee. John will not. He fights against the French. He is

mortally wounded, and dies later in the arms of his father, his old father. Lord Talbot speaks to Death right before John dies in his arms.

Triumphant death, smear'd with captivity,
Young Talbot's valour makes me smile at thee.
When he perceiv'd me shrink and on my knee,
His bloody sword he brandish'd over me,
And like a hungry lion did commence
Rough deeds of rage and stern impatience;
But when my angry guardant stood alone,
Tend'ring my ruin and assail'd of none,
Dizzy-ey'd fury and great rage of heart
Suddenly made him from my side to start
Into the clust'ring battle of the French;
And in that sea of blood my boy did drench
His overmounting spirit; and there died,
My Icarus, my blossom, in his pride.

This may be hard for you to follow, but I will go through some of the important parts. It tells the story of the battle. It personifies death as “*Thou antic death, which laugh'st us here to scorn*”. and now he is personifying death. Death is laughing at them. He calls Death a tyrant, but indicates that he and his son will escape the tyranny of death by escaping to heaven. The father and son will, after death, be coupled in bonds of perpetuity.

Lord Talbot then addresses his son:

O, thou, whose wounds become hard-favour'd death,
Speak to thy father ere thou yield thy breath!

In other words, he is saying to his son, “Before you die, speak to me”.

Lord Talbot goes on:

Brave death by speaking, whether he will or no;

Imagine him a Frenchman and thy foe.

Poor boy! he smiles, methinks, as who should say,

Had death been French, then death had died to-day.

Death is being personified, and in this situation, Lord Talbot is trying to get his son to speak to him. For John, the son, to be able to speak to the father would be equivalent to “braving” death, which means “to face death courageously.” In this situation, several forces are pulling the boy in the same direction, to speak. Those forces are: loyalty to the father, support for the English, and courage in battle against death. John does not speak, but he smiles, which Lord Talbot interprets as being equivalent to speech, speech that would be “If Death had been French, I would have killed him. The only reason I couldn’t defeat Death is that Death is not French.”

Let us look at some of the blends in this passage. Here, in the slideshow, I will go through them one at a time. In the final blend, down here on the bottom, we have Death, the French warrior foe, who gets slain in this imaginary blend. That is a blend of Death the Warrior Foe and a French soldier, so that in the blend, Death is a French Soldier. John Talbot is in both input stories, and he is projected to the blend, but now

in the blend, he defeats not only French soldiers, but also Death, because Death is a French warrior foe. Notice the emergent structure. Above this blend, John Talbot dies. But now in this counterfactual imaginary blended story, it is Death who dies. John defeats him.

We do this kind of blending often, with actual people. We might say, “*If John were David, he would be married by now*”. Now notice, that doesn’t mean that John or David has to be married. Maybe David is rich, and John is not, and David doesn’t want to get married, but John does. And if John were rich, he could get married. So, we can have sentences like

If John were David, he would be married by now.

. . . he would be rich by now.

. . . he would be a stockbroker.

. . . he would use his mobile phone.

If John were French, he would have attended the École des Hautes Études en Sciences Sociales.

. . . he would prefer to lecture in English.

. . . he would live in the Lyonnais.

We blend different people in order to tell us something about the inputs. So if Death were French, had Death been French, then, Death had died today. This looks spectacular, but it is just something we do all the time prompted by these kinds of linguistic constructions.

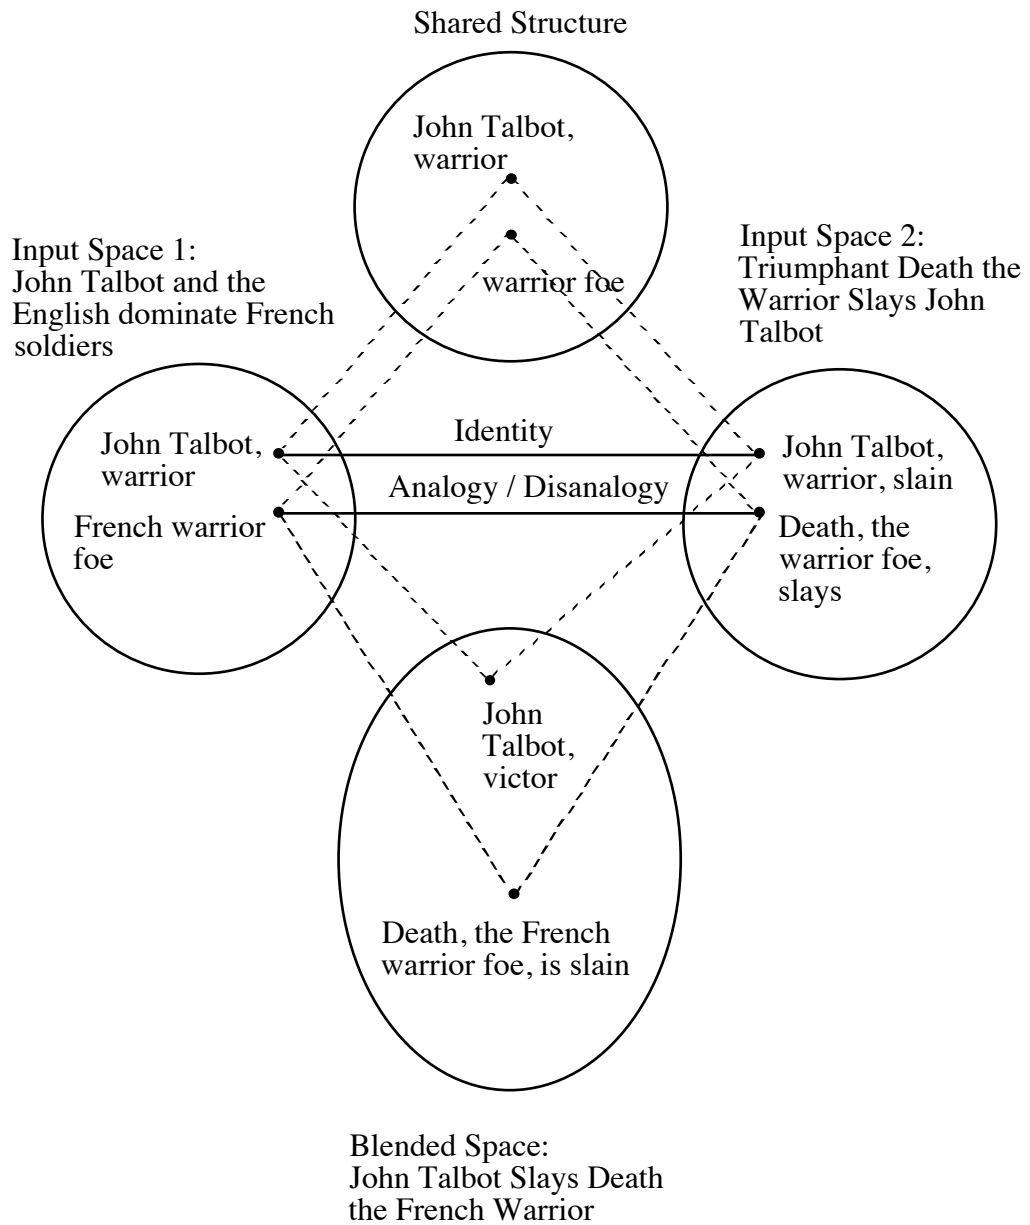


Figure 1: Integration Network — John Talbot Slays Death the French Warrior

But wait a minute, there is Death, the warrior foe, slaying John Talbot, in one of the input stories. How do you get that structure? That is a prior blend.

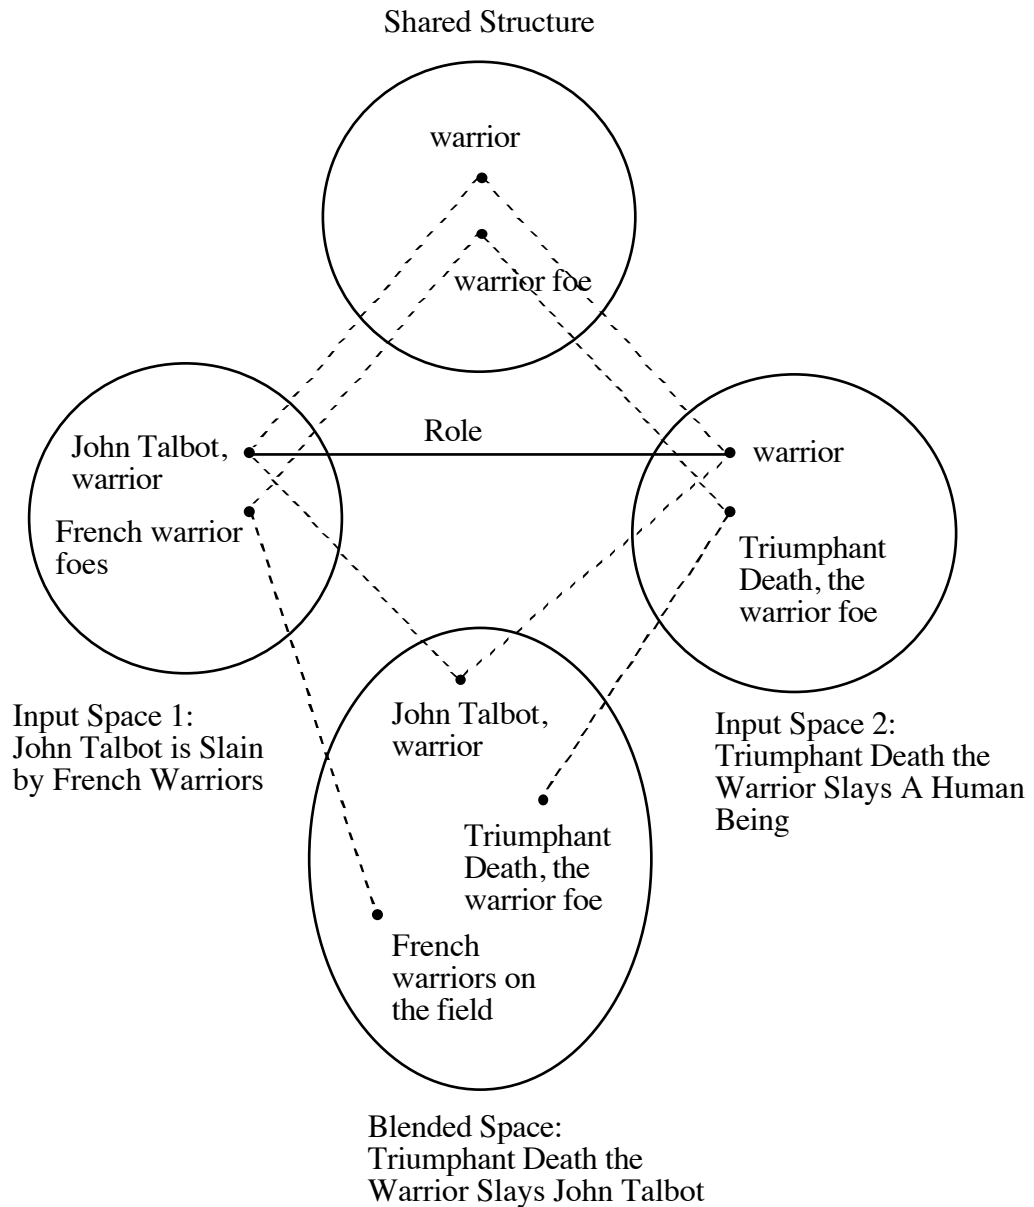


Figure 2: Integration Network — Triumphant Death the Warrior Slays John Talbot

This blend has two inputs: *Triumphant Death the Warrior Slays A Human Being*, and *John Talbot is Slain by French Warriors*.

Well, that is a blend. The role for *human being* is now integrated with the value *John Talbot*. In the blend, Death the Warrior Foe kills John Talbot.

But wait: One of these inputs has Triumphant Death, the Warrior Foe, slays a human being. Where does that come from? It is a prior blend.

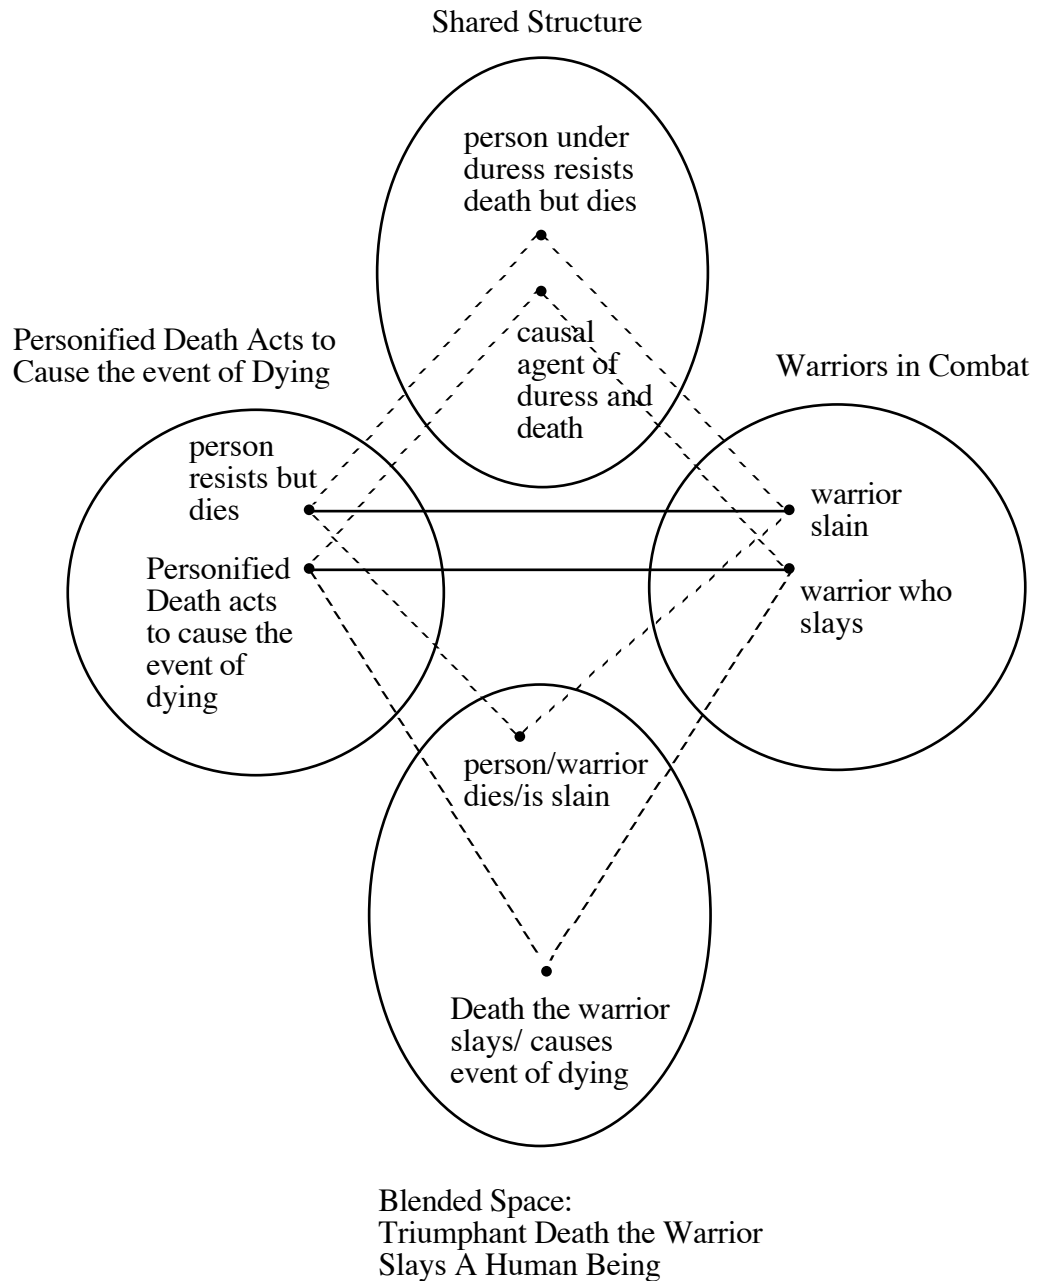


Figure 3: Integration Network — Triumphant Death the Warrior Slays A Human Being

In one of the inputs to this blend, Personified Death causes a person, who resists death, to die. In the other, one warrior defeats another. In the blend, Personified Death is the Triumphant Warrior.

But wait: in one of these input spaces, we have Personified Death. Where does that come from? That is a prior blend.




Figure 4: Integration Network — Personified Death Causes An Event Of Dying

In one of the inputs to this blend, a person dies, and Death-in-General, a general cause in the world but not personified, causes the death. In the other input, one person causes another person’s death. In the blend, Death-in-General is now a person who causes the death of the person who dies.

By Death-in-General, I mean what we mean when we say, “Death comes to everyone” or “Everyone must face death,” and so on, when we mean that it is a general cause, not an agent.

But wait a minute: here we have Death-in-General as a cause in one of the inputs. Where does that come from? It is a prior blend:

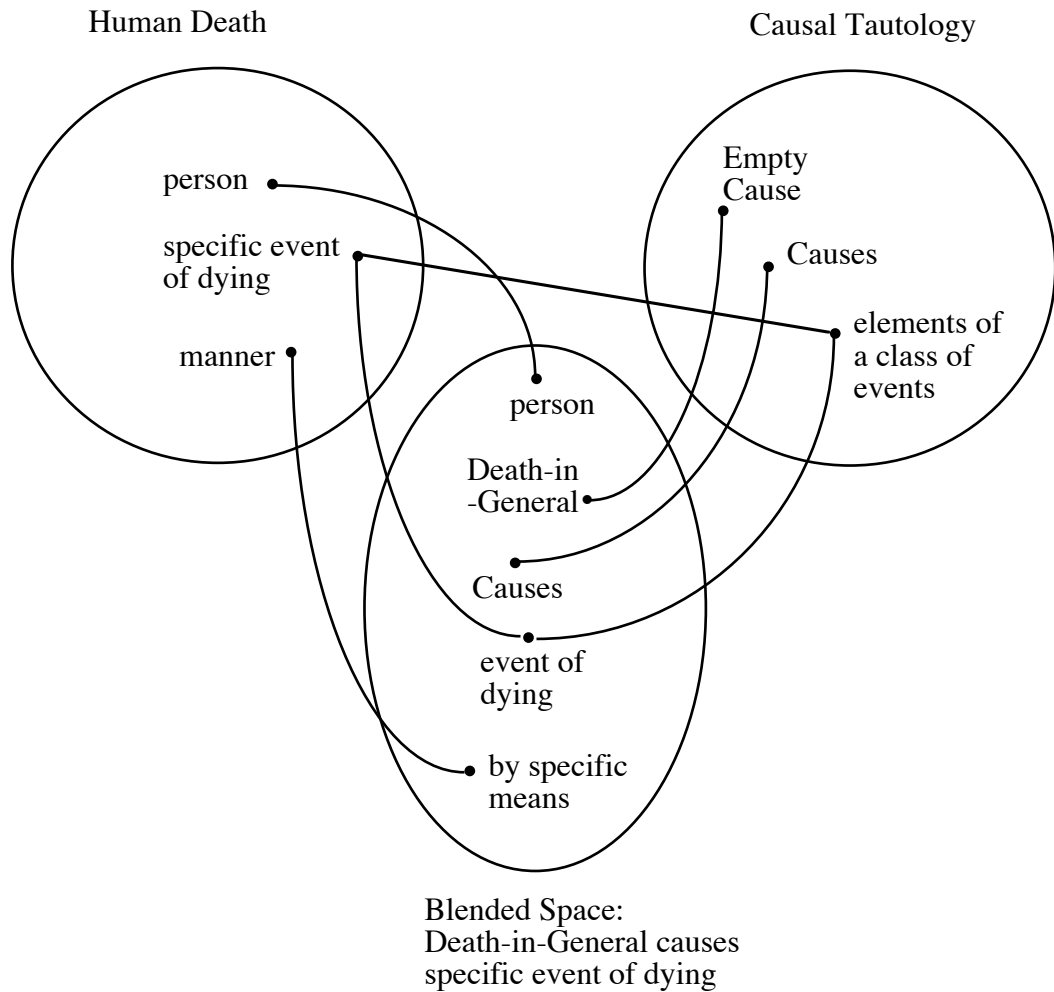


Figure 5: Integration Network — Death-in-General Causes A Specific Event Of Dying

Now that may sound strange. Of course death happens for everyone, sooner or later. But think about it: there are many different kinds of death. People die of old ages, they die of disease, they die of accidents, they die of murder, and they die of many different causes. But above all these, we think there is Death. And everyone is subject to death, as the general cause. This is a kind of causal tautology we construct. What causes death? Death. What causes lust? Lust. And individual act of lust is caused by Lust. Well, what causes hunger? Hunger. Hunger made me eat all that food. There are general causes that we think of as operating in the world. Sloth causes sloth. We say: watch out for sloth! Laziness causes laziness. We say: Laziness will stop him from doing it!

Now it is time to put together all of these blends. You need all of them, in a cascade, to understand

Brave death by speaking, whether he will or no;
Imagine him a Frenchman and thy foe.
Poor boy! he smiles, methinks, as who should say,
Had death been French, then death had died to-day.

Here is a diagram of the cascade of blending that we, or at least many readers, use in understanding this passage:

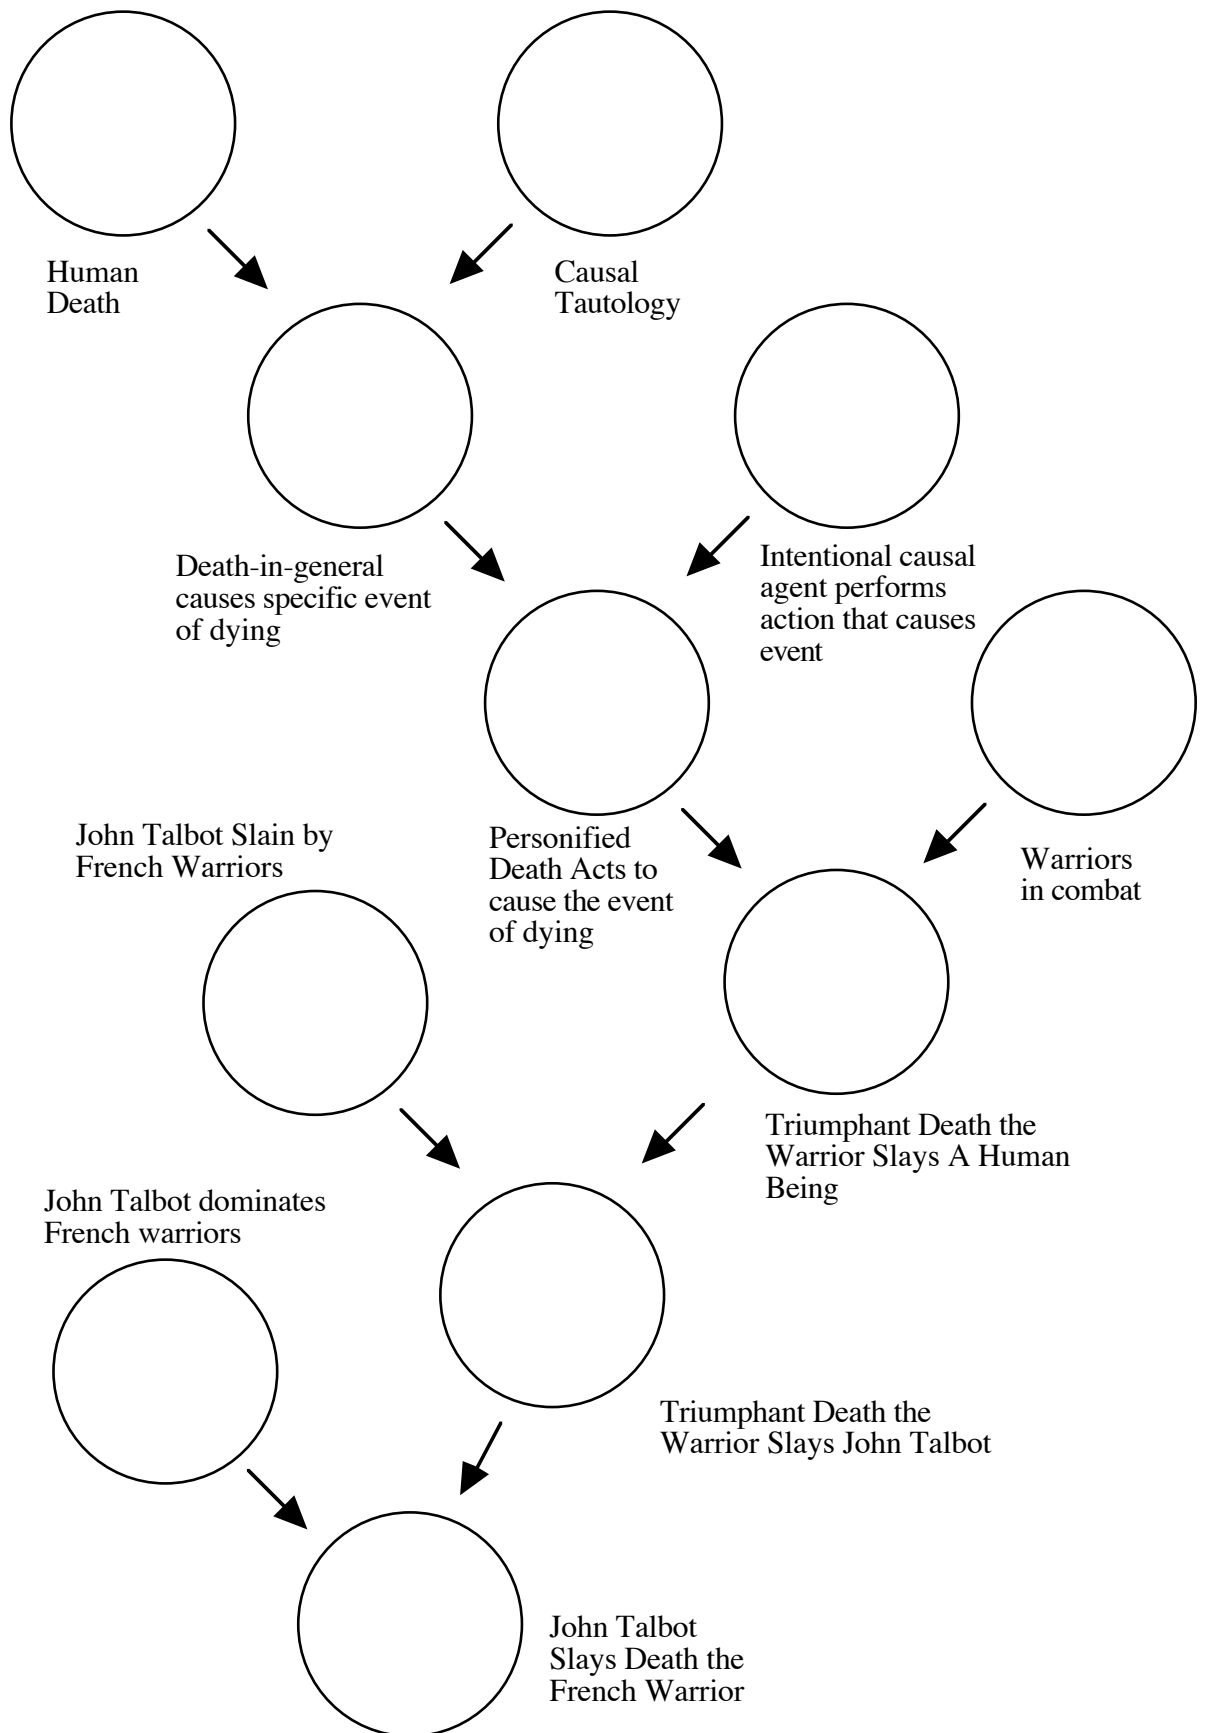


Figure 6: Hyper-Integration Network

Lord Talbot does all this work to try to get his son to speak to him. It is the

defining passage for father and son in the play. This is what we do, often, when we wish to give human meaning to a situation, and persuade people.

Today, we do not have time to go into other, even more elaborate examples, but I refer you here to some other spectacularly inventive blended stories, such as the story in Book Two of *Paradise Lost*, containing the allegory of Satan, Sin, and Death. Satan in the allegory conceives of Sin, and so Sin, a beautiful young woman, is born from his brow, the way Athena is born from the brow of Zeus. She is attractive, so Satan mates with her, and she conceives a child, who turns out to be Death. Of course, Satan is already a blend—and belief is not at issue. We are talking about how human beings can conceive, mentally. For the study of conception, the study is not whether you believe this or that or something else, whether you feel that the blend refers to something in existence, but rather how you put together the conceptual network.

Satan mates with sin. They have an affair, a sexual affair. In the theological allegory, if you are attracted to Sin, engage with Sin, the consequence is Death, the death of your soul. You thought that engaging in sin was going to be wonderful, but it turns out to be horrible and painful. Indeed, the personified Death in this allegory is a terrible monster. Sin is terribly deformed by his birth, so that she ends up being ugly, made ugly, in the allegory by the fact that sin produces death. The allegory is exceptionally complex. I provide an analysis of this in *Death is the Mother of Beauty* and an improved analysis in “Figure” and “The Literal Versus Figurative Dichotomy,” both available from my author page on the Cognitive Science Network, at <http://ssrn.com/author=1058129>. The *Dream of the Rood* is also an example of pyrotechnic blend upon blend upon blend. I analyze it in

Here is another case of blended stories. We have a standard frame of punishment. Somebody sins, or somebody commits a crime, and then we do something bad to them. But we do not frame this as “first they do something bad, then we do something bad to them”. On the contrary, these two acts, the first and the second, are compressed into a balanced unit of crime and punishment. The action we perform on them is in its features just the same, but in the blend its status changes because it is part of this blended unit.

We can also have the concept of an *unpunished* crime or an *unpunished* sin. In that case, you have the sin or the crime, but no punishment. And people can be quite distressed about the fact that somebody committed a crime but did not get punished for it. Well, then, there is this guy named Jesus, who is unsinning. You can take all the unpunished sins of all of humanity and compress them down into the blend so that now they belong to Jesus. Jesus takes the sins upon himself. He is punished and that clears, that balance is out. Note that other blends come in this, such as the Lamb of God.

That’s just background. Now let’s go back to *The Dream of the Rood* — “Rood” means Cross. I analyze this pyrotechnic hyper-blend in “Double-Scope Stories,” also available on the Cognitive Science Network, at <http://ssrn.com/author=1058129>. In *The Dream of the Rood*, there is this thane—a thane is an attendant of a lord, something like a samurai, but not really—and this thane has a dream, and in the dream, the Cross, the Rood, appears, and starts talking to him. It is the Cross upon which Jesus was crucified.

The Cross speaks. This may seem strange, but think about it: physical objects seem to speak to us, or communicate with us, or prompt for meaning. For example, here is a mobile phone, which I am using in China. I remember when I had it in Greece. I remember when I had it in Portugal. It is as if the mobile phone is reminding me, communicating with me. We often feel, in the blend, as if physical objects are communicating with us.

In the blend in *The Dream of the Rood*, the emergent structure is more robust. The Cross does not just prompt for meaning, or communicate, but actually speaks to the thane. It says to the dreamer, in effect, “Look, you feel like a sinner. You feel as if you have sinned so much that you cannot be redeemed, but you can. I felt the same way. I was the instrument that caused the death of Christ. I was stained with his blood. The same nails were driven into me. I participated in a bad design, just like you. But I was redeemed. You can be redeemed.”

There’s another blend in this literary work that is very culturally specific. A thane is a retainer, someone who works for his lord, a human being who is a leader. The Cross helped Christ fulfill his destiny. In the blend presented by the Cross, Christ is now not only the divine Lord but also the human lord, the one who is served by the thane. This is a very old story from England. Passages from *The Dream of the Rood* are carved on the Ruthwell Cross.

Most religious paintings rely on blended stories. In the slideshow, I present a painting of the Annunciation. Here is the Virgin Mary. Historical detail is projected to this blend not just from the historical period of Mary but also from the historical period of the painter. In this painting, Mary is reading a book. Of course, there were no books in the historical period of Mary. In fact, in paintings of the Annunciation, she is usually reading the story of the Annunciation, out of the gospels. Nobody finds such a painting chaotic. The painting you are looking at is by Rogier van der Weyden. The modillion above the bed depicts the Resurrection. So Christ is not yet born here in this scene. This representation is a unification of the whole of the Christ’s history, from not yet being born to rising from the dead. All of that history is compressed to a human scale scene in this painting.

Notice that the angel comes in at human scale. You do not see only the foot of an angel who is sixteen meters tall, for example. The angel speaks to Mary in human voice, because of course that is all that Mary, who is human, can understand.

Here is another painting of the Annunciation. This one is by the Master of Flemalle. In this one, again, we have a lectionary, tracery windows of the sort one sees in a church, and so on. Up here, in the upper left, you see a tiny homunculus. That is Christ. He is tolerating his own Cross, as he did in the Stations of the Cross. The breath of God, or the beams of light, are coming through the window, aimed at Mary’s womb, where Christ will be implanted. God is out there, beyond the window. God is sending Christ right into Mary’s womb. So in the Annunciation, at the moment of the Annunciation, the future is clear. Christ is going to be born. These depictions of the Annunciation ask you to blend many stories—the story of Christ’s crucifixion, of his carrying the Cross, the story of Mary and her learning that she has been chosen by God. All these stories are blended into one scene in which, as it were, eternity is

presented in one moment.

If you would like to see more developed examples of the kind I gave earlier, of the *risk* frame and the kind of linguistic constructions that go with it, then show up tomorrow and the day after, when we will be discussing the role of double-scope blending in the formation of linguistic constructions. This kind of blending looks spectacular when we see these amazing literary examples, but this operation is the basis of the cognitively modern human mind. The principles that make these literary examples possible are the same ones that make language, syntax, and frame semantics possible. See you tomorrow! Thank you very much.

Lecture Nine The Nature of Language

Thank you for that wonderful introduction.

Welcome back to the infinite cognitive science seminar, the infinite cognitive linguistics series. Thank you for inviting me to such a beautiful, historic, and famous university. I have always wanted to see Peking University, and now here I am. Thank you for your patience and your determination and your persistence in following these lectures.

Today, I will talk about the nature and origins of language. I'm often surprised when I see discussions about the origins of higher-order human cognition. These discussions often begin with an assumption that human beings were somehow in a "state of nature," without institutions, rituals, social contracts, and so on, but equipped with full language, and that then they developed other abilities aside from language. So, imagine: it is as if human beings could converse fully, able to use conditional and counterfactual constructions; constructions for focus, viewpoint, and perspective; causal constructions; and so on, but had not used any of that mental firepower to develop any art, mathematics, science, social institutions, and so on. I find this assumption unaccountable. It does not seem to me possible that human beings could have had all the mental abilities that language requires and not use them to develop substantially in the other aspects of human higher-order cognition.

So I ask myself how someone could imagine that first came full language and then other aspects of higher-order cognition. How could theorists think like this? How, for instance, could social contract theorists think that we were in a state of nature, with language, and then decided afterward to invent things like social contracts? My guess is that this way of thinking came about as follows. If we look at the history of the last few thousand years, which is a very brief amount of time, we see that human beings have often invented new institutions, rituals, social contracts, and other activities having to do with social ontology and deontology. So it is not difficult to imagine a time when some of the particular cultural products of our mental creativity were not here, even though we had full language ability. I suppose that we can imagine holding constant our full language ability as we go backward in time while decreasing these other kinds of behaviors and products—having to do with social ontology and deontology—down to zero. That would give us a condition of full language but no social ontology or deontology.

But in fact that condition makes no sense. While it is imaginable, it must be false. The mental abilities required for language include the mental abilities for social ontology and deontology, and full language already has elaborate social cognition built into its core. I do not think that first language came fully into existence and then subsequently other behaviors came up. I think instead that double-scope blending made it possible for all of these kinds of behavior to come up simultaneously, each supporting all the others.

In what follows, I will present a different view, according to which language is

viewed as one of a group of performances made possible by the development of double-scope blending.

What is language and how did it originate? Much of what I will say here is to be found in a book Gilles Fauconnier and I published: *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. You can find the references for this talk in that book.

There are many stories about where language came from. You are familiar with them. There are nativist stories according to which something dramatic happened, some dramatic genetic change, suddenly producing language. There are modular stories according to which the appearance of language is unrelated to other capacities. Nativism and modularity do not have to go together, but sometimes they do.

Terry Deacon has put forward a co-evolutionary theory according to which the brain and language gradually developed together over the last 2.5 million years or so. There are associative and learning theories according to which human beings have developed wonderful abilities for statistically extraction that have made language possible.

I am going to present here a different view. Before we begin, let me sketch out two common fallacies into which human reasoning often falls, especially when thinking about human cognition.

One of these fallacies is *cause-effect* isomorphism. Sometimes, we see a big effect in the world, and we assume that the cause must be equally big and dramatic. We think that because the effect is sudden and noticeable, the cause must have been sudden and noticeable. But that's not necessary at all. A large and dramatic effect can be the result of causal changes that are gradual.

Those of you who have been following these lectures will recall that I have discussed our need to compress things to human scale, so we can hold onto them. We commonly compress cause and effect so as to hold them together. This gives us a simple principle for how the world works. Consider, for example, that when you look at something in the world—a horse, for example—you see it as one thing. Why not? Suppose someone asked, why do you see it as one thing? Your answer would be, *because it is one thing*. That is a cause-effect compression. The effect is that it seems to be one thing. In the blend, the cause is that it is one thing.

In fact, that's not all the reason that you see something in the world as one thing. All the different aspects of something you take to be one thing are processed variously in the brain and not brought together anatomically. Yet it seems to us as if the oneness of the thing goes straight from the external world to our internal mind, right into the brain. Why do I see it as one thing? Because it is one thing! And so the effect is the same as the cause. They are compressed in the blend.

Although this is quite an effective way for us to think, it is important in science not to conflate the cause and the effect.

So if language is suddenly here and human beings have it, that doesn't mean that there was some sudden cause that made language appear. That way of thinking is a cause-effect compression. It is the fallacy of cause-effect isomorphism.

A second fallacy is the function-organ isomorphism.

People are somewhat more aware that this is a fallacy. In this fallacy, you think that because people have tongues and use them to speak, the tongue is an organ for speaking. The tongue has a function. So in this fallacy, it becomes an organ for that function. You are aware that primates and mammals had wonderfully articulated tongues long before any organism arose evolutionarily who could speak. So the tongue is an organ, but just because it is used in speech does not mean that it is an *organ for speaking*.

An opossum is an animal that hangs by its tail from a tree. When you look at an opossum, it is easy to think that the tail evolved in order to make it possible for the opossum to hang from a tree. The tail is then viewed as an organ for hanging. But, no, there were tails long before opossums would hang from trees. It is important to be careful to avoid the function-organ fallacy.

I will today in this lecture propose a theory that escapes both of these fallacies, or escapes both of these problems. I am going to argue that language is a product of a gradual change in a certain capacity. It is a by-product of the origin of the evolution of modern cognition.

Hyman beings have been anatomically modern for about 150,000 years. By that I mean that human beings over that time interval resembled us anatomically. But it seems as if we have been cognitively modern for only fifty, sixty, seventy thousand years. We became cognitively modern because evolution equipped us with a more advanced form of conceptual integration, namely double-scope integration. Evolution did not so much make us human as give us the mental capacity that we use to make ourselves human.

Rudimentary blending seems to be in evidence as far back as early mammals at least. It has been evolving gradually since that time. Cognitively modern human beings took one extra step in blending. They were endowed with the highest form of blending, which we call double-scope blending.

Once they had double-scope blending, then they had a capacity to manage large conceptual integration networks that other species are not able to manage. This ability made it possible for us to have art, music, advanced social cognition, advanced tool use, fashion, mathematical insights, scientific discovery, and language. These things came up together, supporting each other. The suite of higher-order cognitive abilities arose not because one of them came up first and made the others possible, but because double-scope blending evolved through one more gradual step, and that made all these various higher-order cognitive abilities possible.

According to the theory I am putting forward today—this is a preview of coming attractions—there is a long gradual evolution of a mental ability, namely conceptual integration, a gradual evolution of the ability from early mammals up to us. And then, when you get to the final step, the one where double-scope integration becomes available, then language becomes possible. So the effect—language—is fairly sudden in the history of the species. It is a big deal, a singularity. But the causes are neither sudden nor singular. The causal development has been going on for a very, very long

time.

What really separates us so dramatically from other species is not that we have the ability to perform conceptual integration, but that we have the ability to perform the most advanced form of conceptual integration. Other species can do rudimentary blending, but we can do the most advanced form. Other species have the ability, as I said, to do rudimentary blending; it has been around a long time.

The evolution of double-scope blending made a variety of human singularities possible: language, art, science, advanced tool use, and so on. All of these performances are products of double-scope blending, and they have been invented by cognitively modern human beings in cultural time”, not in evolutionary time.

“Cultural time” does not mean “overnight,” but cultural invention moves much, much faster than biological invention, as I discussed in the first lecture. Biological evolution takes a long time. You get differences in different species. Cultural evolution is swamped out by cultural evolution once culture becomes possible. Once you get to the point of double-scope blending, you can develop new meanings without the need of new biology.

To review the elements of double-scope blending: in double-scope blending, you have at least two mental spaces as inputs. In the slideshow, I represent these mental spaces as yellow and blue. For example, let this yellow circle in the slideshow represent the mental space that someone is thinking about, in which he is boarding a plane and thinking about boarding a plane. In another mental space, represented here in the slideshow by the blue circle, the person is thinking about a time he was in Beijing and consumed roast duck or duck soup or “zongzi” on Dragon Boat Festival day or something like that. In this mental space, he is having a wonderful time, and great conversations. It’s the kind of mental space I have been having here during my visit to Beijing.

Now, someone can make a blend, between how the plane changes the professor’s location, and how the wonderful Chinese students of cognitive linguistics change the professor’s thinking while he is in Beijing.

You have seen simple blends and complicated blends. We have seen how blending networks bring in previous blending networks that culture has provided. In this particular blend of the plane trip and the conversation, conversation with brilliant Chinese students can be a ride on a magic jet plane, a jet plane that takes you places you never thought you would go. The result is a new kind of conception.

Sometimes, it is hard to see that any blending is going on, because we are equipped by evolution to do the blending but not to see that we are doing the blending. So let me start with a couple of pyrotechnic blend where it is impossible to miss the blending

The first is a poem by Craig Raine. The poem is titled “A Martian Sends A Postcard Home”. This Martian comes to earth and then sends a postcard home about what’s on earth. The poem begins,

Caxtons are mechanical birds with many wings

and some are treasured for their markings—

they cause the eyes to melt
or the body to shriek without pain

The Martian uses the word “Caxtons” for the book—Caxton was an inventor of the book. The Martian is looking at the book and wonders what it could be. He sees the pages and imagines that it is a bird, except with many wings rather than two. He conceives of the object as a blend of the concept *bird* and the physical features he sees. The illustrations in the book are understood as markings. Crying is understood as *the eyes melting*; and laughing is understood as the body shrieking without pain. The eyes melting are crying, and the body shrieking without pain is laughing that it has lots of pages, and he’s trying to figure out what it could possibly be in saying “Oh, that is a bird except that it has many wings instead of two”. And *some are treasured for their markings*, there are illustrations inside, and *they cause the eyes to melt or the body to shrink without pain*. So this is crying, the eyes melt, and laughing is the body shrinking or to shrink without pain. He says,

I have never seen one fly, but

sometimes they perch on the hand.

The Martian, looking at something he has never seen before, someone holding a book in hand and turning pages, blends it with something he does understand—a bird. The result is a bird, but not like any bird he knows. There is emergent structure in this blend.

The Martian’s postcard continues on in this way about many other objects and activities that are familiar to us. What I want to point out is that this Martian is doing double-scope blending! So he is just like a human being, except that he lacks many standard conceptual frames that are quite familiar to all of us. What makes the Martian different is not his mental processes, but that he is from a very different culture. This is an example of the important point that an analysis of how conceptual integration works does not spare you the work of analyzing how different cultures have used it to assemble conceptual networks and linguistic constructions. Blending is a process. Knowing the process does not specify what the products will be.

Let us now turn to the question of what the study of conceptual integration can tell us about the origin of language.

What features should a theory of the origin of language have? We must work indirectly on this question, because we do not have a time machine to go back and look for the origin of language. I usually work on living human beings. I can interview them, gather data in ecologically valid environments, run experiments in the rare cases where the data need to be elicited, and so on. But there is no corpus for language fifty thousand years ago, so I cannot do corpus analysis on the origin of language, which is a shame. I have applied to the United States National Science Foundation for a time machine, so I can go back fifty thousand years and do the science that needs to be done,

but for some reason they have not given me.

I do wish I had that time machine. It would make things so much simpler. But until it arrives, we will not have robust data from the period we are talking about, so we must work indirectly, with long inferential chains.

Language does not survive from that era, and brains do not, either. It is accordingly somewhat odd to talk about a theory of the origin of language. There are many such theories, though, all of which I think are mistaken. I propose a rival hypothesis here, one conceived by me and Gilles Fauconnier.

To begin, a theory of the origin of language needs to recognize the singularity of language, in this sense: if we look back through history, at the descent of language, there is no evidence that I know of to support the notion that there were ever sustained stages of intermediate language. Language does not seem to have been developed by starting very small and simple and growing ever fuller and more robust. There is no evidence, for example, that twenty-five thousand years ago there was a language that lasted for a thousand years that consisted of only three-word combinations and only one clause and only a thousand words, a language without conditional expressions or negative expressions, and so on. No such language has ever been found and, I predict, no such language ever will be found. There is no evidence whatever that there were ever any sustained intermediate stages of language. All of the world's languages we know have immense complexity. That's remarkable and important. All dialects, all languages have immense grammatical complexity.

So let us not dismiss that. It is an important condition of language. Why is it that, for example, five of the world's languages consisted of just simple little grammars? But they don't. We find no evidence for that in the past and we find no evidence for it now. Any one in the audience could prove me wrong right now by pointing out that there is a natural human language that has been around for hundreds of years that has only a simple little grammar. Thousands of natural human languages in the world and in the past, and not a single one has a simple little grammar. It is not just that no one has found a counterexample, but rather, I assert, that no one ever will find a counterexample. Whatever theory of the origin of language we propose must embrace that fact.

I also think we should try to avoid the fallacy of cause-effect isomorphism. We should avoid the automatic assumption that because language is singular and seems to come up suddenly—relative to the long evolutionary descent—there must be some sudden and spectacular cause for it. We should avoid this fallacy because evolution does not typically operate that way. As François Jacob observed, evolution is like a tinker. It takes what is in the shop and adjusts it slightly. So a theorist who proposes a spectacular genetic spontaneous event as a source of language must bear the burden of proof for such an atypical evolutionary event.

A theory of the origin of language should seek a continuous path of evolutionary change over a very long period as the cause of language. That's how evolution usually

works. The path should be a plausibly adaptive story evolutionarily, one according to which each change along the path is adaptive in itself, regardless of where the path ultimately leads. This is a standard point: Evolution does not get to say (I am personifying evolution now): “It would be really useful if I could get to such-and-such a point ten steps down the line, so just leave me alone for a while until I get there.” No, in a good evolutionary story, each step would be adaptive. This is not to say that all biological change we see in the record had to be immediately adaptive—there are many mechanisms of evolution, such as genetic drift, for example, that produce changes. And pressures of selection vary over time. But since we do not have a time machine and are trying to come up with the best inferential theory, we will be in a stronger position if we can have a story where each step is adaptive. So we are looking for a continuous evolutionary path that can nonetheless produce singularity. In that case, the cause is continuous and is evolutionarily adaptive all the way, but produces singularities at various points. We would like to have a model of what mental operations developed along that path, and in what order. It would be very good if we had an explicit account of what continuous changes produced what singularities and how they did it.

We would also like to see evidence from many different areas that human beings now actually perform the mental operations that we are hypothesizing for that evolutionary path of mental operations. We look therefore for intermediate stages in cognitive ability, but not in the function of language itself.

That is just the kind of theory of the origin of language that I am outlining here: there was a run-up throughout the mammalian history of conceptual integration abilities; they became more advanced, gradually.

We would of course like to see evidence in the anatomy or behavior of today’s human beings pointing to the history of these steps, just as anatomical evidence in today’s human beings points to our once not having been bipedal.

And other things being equal, it’s not bad to have a parsimonious way of explaining the emergence of many related human singularities as products that arose along the same continuous evolutionary path.

In the slideshow, I am presenting a diagram that we use in talking about conceptual integration networks, just as something to point to. Blending is a mental operation. It operates throughout conception. Many patterns come up routinely in networks. These routine patterns are just reference points, a few key spots in the great variety of networks that can come up: simplex networks, mirror networks, frame-compatible networks, single-scope networks, and double-scope networks. Above all, when you are analyzing a conceptual integration network, do not ask which five of these patterns it must fit. These are just a few key patterns. Blending runs over all conception, and these are just a few special points. In numbers, for example, 1, 10, and 100 are special points, but there are many other numbers, along gradients of various sorts.

Blending networks compress networks to human scale in the blend. We have seen some of the ways in which that can happen. The network can borrow a compression from one of the inputs, as in *you are digging your own financial grave*.

The digging of the grave is already compressed, and we use that compression in the blend to provide an anchor for the vast and diffuse network concerning an investor's going bankrupt. *If Clinton were the Titanic, the iceberg would sink* also borrows a compression from one of the inputs. The input with the Titanic hitting the iceberg is very compressed: the boat hits the iceberg and sinks. In the blend, we use that compression to provide an anchor for the vast and diffuse network concerning a great range of political activity, with many agents over extended time intervals, and so on. There is emergent structure in this blend: One thing you know about the Titanic is it sank; but in the blend it does not. One thing you know about an iceberg is that it cannot sink, because ice is less dense than water. Ice can be submerged, but not sink. But in the blend, ice can sink. This emergent structure is influenced by the causal patterns in the input space for finances, not the input space for the Titanic.

A second major method of creating a human-scale blend is by compressing outer-space vital relations between the inputs into inner-space structure in the blend. We have seen many examples of this. We saw this in the Mythic Race, in which six of the world-record holders in the one-mile race, over four decades, ranging from 1954 to 1996, are compressed down into a single one-mile race in the blend. They are all on the same track in the blend. In this network, there are outer-space relations of time and analogy connecting the six input spaces. The network compresses them down to tight structure inside the blend, a single mental space. In the blend, we have one race, and it is happening at one time.

We saw counterfactuals, prompted, for example, by a word like "missing", as in "Put the pear juice in front of the missing chair". In one mental space you have a table with chairs around it, and in another mental space you have a table with more chairs around it, either because you are used to their being one more chair or you have an expectation that chairs should be evenly spaced around a table, so now when you look at the table, you can see where the missing chair is. The missing chair in the blend is a compression that results from work on vital relations of analogy and disanalogy between mental spaces. "Missing" is a special word that we take as indicating, "OK, there is going to be some disanalogy between what you see or what you perceive and a normal conception, or maybe a memory, and I want you to compress that disanalogy down into a special feature, and that feature is *missing*." In the blend, the chair is conceptually present, and it has a feature: it is *missing*. It's like *the red chair, the wooden chair, the missing chair*. We have many words that prompt for such compressions, such as *gap, mistake, accident, hole*. All of these words are prompts to invite you to compress an outer-space network into a feature in the blend.

We also saw the pattern of compression according to which outer-space analogy and disanalogy vital relations are compressed into the blend to create a unique element and change for that element. We saw this in cases like "The cars get three feet bigger when you enter Pacific Heights", and other examples:

- *The fences get taller as you move westward across the United States.*
- *The paint gets darker as you move down the wall.*

Of course, outside the blend, no paint is getting darker. Rather, the paint at one end of the wall is darker than the paint at the other end of the wall, and there is a gradient of variation from one to the other end. There are analogy and disanalogy links across all the individual sections of paint. They are compressed into the blend to make a unique element—the paint—and change for that unique element—it gets darker. We do not think that on the wall anything is changing. It is static. We are not deluded by the compression in the blend. But I can say, without objection, “*The paint gets lighter as you move to the window, as it goes to the window.*”

Some of the important findings of conceptual integration theory are listed above in the slideshow.

One finding is that conceptual integration is at work in category extensions, analogies, metaphor, framing, counterfactuals, grammatical constructions, and so on. Phenomena that have been taken to derive from very different kinds of mental operations, with different systems, are, it turns out, not driven by different mental operations. These are all products of conceptual integration. They feel different, and we must account for the fact that they feel different. But they do not differ in the underlying basic mental operations.

A second finding is that we are often able to take diffuse conceptual network and compress them to human-scale blends.

A third finding is that conceptual integration networks lie on a gradient of complexity. There are some down at the lower end of complexity, such as simplex networks, which we have no reason to think are not within the competence of most mammals. But there are also mirror networks and single-scope networks and double-scope networks. There is a continuum of difficulty in putting together these networks. Cognitively modern human beings do not have any difficulty with any of them. For us, it is all “duck soup.” But no other species, despite their great talents, can handle the blending networks that we handle.

The central problem of language is that you have relatively few constructions. Grammar is very complicated, but it’s nothing as complicated as conceptual structure. If you had to have one word for every different frame you couldn’t handle that many words. You need relatively few linguistic constructions that you can use and combine to talk about almost anything. You need words like *food* that can apply very widely, or words like *you* or *I* or *brown*. You must have closed class words like *I* and *you*, in addition to open class words like *food* and *brown*.

Consider something like *my cow is brown*. This seems very simple, but think of all the people who could say that in all the different situations of all the different cows, including things that you might not think of as a cow until somebody used the word “cow” of it, and of how many different colors in how many different places might be suggested by “brown”.

Think of the resultatives we talked about, like *Catherine painted the wall white*, with the syntax noun phrase, verb phrase, noun phrase, adjective. Other resultatives are

- *Roman imperialism made Latin universal.*
- *He hammered it flat.*
- *I boiled the pan dry.*

These are all cases of resultative clausal constructions: noun phrase, verb phrase, noun phrase, adjective. The events described in these examples lie in completely different domains. These resultatives range over elements like Roman politics versus blacksmithing. *He hammered it flat* is from the domain of blacksmithing, but *Roman imperialism made Latin universal* is from the domains of Roman politics and linguistic geography.

These examples have strikingly different time spans. The time span in which Latin becomes universal involves hundreds, may a thousand years or more. But *the earthquake shook the building apart* involves only at most minutes.

These examples also have strikingly different spatial ranges. The spatial range for the Roman Empire and Latin is most of what is now Europe, as well as some other regions. But the spatial range for *I boiled the pan dry* is maybe one square meter: the stovetop.

These examples also have strikingly different ranges of intentionality: Roman imperialism involves hordes of people over centuries, not to mention all those speakers of Latin, while *I boiled the pan dry* involves only one cook and *The earthquake shook the building apart* has an earthquake in the causal position.

These examples involve very different kinds of cause and effect. The cause is a hammer blow in *He hammered it flat*. But *Last night's dinner made me sick* involves eating the meal one day and feeling sick later on through a long chain of biological events.

Nonetheless, the resultative construction can be used to apply across this vast, vast range of different kinds of conceptual structure, because the resultative frame is used to give a human-scale compression to the blend, and that compression comes with the resultative construction

What kind of theory of the origin of language do we need, given these features of language and its use? Double-scope integration is just what we need to make a word like *food* have very wide application. We have food and its common frame in one input, and a different situation in another input, and they are blended, and then whatever *food* is blended with in the blend can be referred to as “food”. We can talk about *bird food*. If we find people eating something that we have never seen before, we can still use the word “food” of it.

I can talk about *bird food* and when I find a dish that some people are eating I have never seen before, it's still *their food*.

Across all these different examples, quite different conceptual structure is being developed. But we do not need a new word. I can use “food” because we all have double-scope blending. The word “food” is projected down from one input space and we can use it of the element in the blend.

Linguistic constructions attaching to the inputs can be projected down to counterparts in the blend, so, for the most part, you do not need new language to express

new conceptual structure. You do not need new grammar to express new meaning. On the contrary, the language that you have attached inputs provides you with the language you need to express quite different things in the blend. In this way, as a general principle, double-scope blending solves the central problem in the origin of language.

Blending also avoids the fallacy of cause-effect isomorphism and provides a reasonable adaptationist account. The development of the capacity for blending was itself gradual and required a long expanse of evolutionary time. Basic blending goes back deep into the mammalian line, as far as we can tell. Each step in the development of blending capacity was adaptive. From very simple simplex blends like framing all along the steps of the gradient up to very creative double-scope blends, each step of the capacity would have been adaptive, because each step gives increasing cognitive ability to integrate and compress.

There is ample evidence of intermediate stages in development of the capacity for blending. Some species, for example, seem to be able to create simplex networks. We see these species now. Other species seem to be able to do slightly more complicated simplex networks. If you take a chimpanzee, as Michael Tomasello has, and put the chimpanzee into action in strong domains, such as dominant-subordinate relationships, and consumption of food, you can see the chimpanzee working at the limits of the ape brain, perhaps to understand that the other chimpanzee is goal-directed and is unaware of certain facts. But there is no indication that the chimpanzee can understand that the other chimpanzee has a false belief.

Right now, in the present, we see evidence of earlier blending abilities. We see those abilities in other species. But there is also evidence in living human beings: human beings, for example, can do not only double-scope blending but also simplex blending, frame-compatible blending, and so on. We can do all the rudimentary forms of conceptual blending.

A special level of capacity for conceptual integration has to be achieved before a system of expressions with a limited number of combinable forms can cover an open-ended number of situations. A full language has equipotentiality: however limited your grammatical constructions, you can use them to talk about anything. That is possible if you can do double-scope blending. Even small children who have relatively few words and relatively few constructions can use that limited range to talk about what they want to talk about. As we go through life, we develop a larger array of linguistic constructions, form-meaning pairs, developing a relational network of form-meaning pairs. But even when that network of available form-meaning pairs is limited, we use it with equipotentiality because we can do double-scope blending. Children can look at a lion and say, "Look, kitty!" No problem. We can understand *kitty* just as we can understand *food*. The child doesn't know the word for *lion* yet, but it's got *kitty*, points and blends the meaning of *kitty* with this thing and now the word comes down into the blend, and the child can use *kitty* or *kitten* or *cat* to express what it wants to express. That's equipotentiality. Of course, the child's repertoire of linguistic constructions will improve, but equipotentiality is available early.

Kanzi, the genius bonobo, taps out at about 250 pretty specific tokens. Kanzi does not seem to be able to get past that. A human child, who blows past that barrier without even noticing it, isn't doing what Kanzi is doing, mentally. The human child has double-scope blending, and that makes all the difference. The smartest bonobo anybody has ever found, with the best developmental profile and the best human cultural support, cannot enter the vocabulary explosion the way every human child does.

The development of double-scope blending was not itself a cataclysmic event, but rather an achievement along a continual scale of blending capacity. So there is no cause-effect isomorphism in the theory I am presenting here: in the origin of language, the cause was continuous, but the effect was a singularity. Language arose as a singularity. It's a new behavior that emerged naturally once the capacity of blending reached the double-scope level.

Language is like flight. It's an all-or-nothing behavior. Now when you look at the evolution of wings, the evolution of wings from dinosaurs to birds, the story there is gradual. At each step, it was adaptive, even though the organism developing the wings could not fly. Those who had these wings could hop better, stay aloft longer, and so on, but not fly. There comes a moment in which the organism is truly airborne, and that changes everything. It is a tiny difference in causes with a fabulous difference in effects. The moment of being airborne is when the organism can stay in the air regardless of the initial force of launch. Science knows of many such situations, in which a small change in causes produces a large change in effects. A tiny increase in the temperature causes water to turn from solid ice to liquid.

I grew up swimming in the Pacific Ocean, off Del Mar, California, on the North Coast of San Diego. I can tell you from experience that a tiny difference in body density—including body fat but also the exact wet suit you happen to be wearing—makes every difference in whether you float or sink. One more gram of fat, and you can float in the saltwater ocean without paddling. One gram fewer, and you sink, not immediately, but you must expend effort to stay afloat. We are talking here about drowning versus living—this is a huge difference. If you are asleep on the water, on your back, and you have the right density, then you live. If you have a slightly larger density, then you die. Small difference in causes, huge difference in effects. I think small changes in causes produced a huge difference: language became possible.

The reason, I think, that a culture cannot stop at a simpler language—such as just a subject and a verb—is that once you have achieved the capacity for double-scope blending, then all of the ability to build linguistic constructions just comes right up, so you will not find uncomplex, simpler grammars among human beings. You will not find them, I assert, in the history of the species. I do not mean that suddenly a tribe awoke and everybody had full language. Rather, once double-scope blending had some distribution in the community, then it took only cultural time, not evolutionary time, for language to develop.

The story of the origin of language does indeed have room for intermediate stages in the capacity, but not the product.

The hallmark virtue of language is its ability to use grammatical patterns suitable for basic human scenes to capture and convey much less tidy meanings. We have been

talking about this throughout this series of lectures. The pyrotechic examples are like *If Clinton were the Titanic, the iceberg would sink*, but the same is true of an expression like *she granted me my wish*, which is a ditransitive. The basic scene of a ditransitive is one such as that indicated by *he handed me the apple*: one person transfers something to somebody else. That's what the ditransitive construction is built for. But you can say *she granted me my wish*, which now uses the ditransitive frame as an input to a social situation. With the ditransitive as an input to the blend, I can say things like *the country denied me entry*, where I'm using as an input a transfer meaning, but also using as an input a frame for deny, which means blockage. In the blend, we have blockage of a transfer. Now, even though we are trying to manage a vast and diffuse network that does not fit into a little physical scene, the blend can have a ditransitive structure, even though the agent is a country. We are not deluded. We do not think that there is an intentional agent called *the country* and a physical object called *entry*. But the ditransitive frame is an input, and it provides syntax to the blend. That ditransitive frame provides a compression to the blend, and language for referring to elements in the blend.

Language, in the strong sense, is equipotential. To say something new, we do not need to invent new grammar, and a good thing, too! What we must be able to do is to conceive of a blend that let existing grammar come into play, via projection from an input.

There are others who have proposed something like blending as the operation that provided for large advances in human cognition. Stephen Mithen, for example, talked about "cognitive fluidity" as a big breakthrough. His theory is quite different from the theory of conceptual integration in a number of ways. *The Way We Think* outlines the differences. Nonetheless, Mithen did propose that the ability to integrate disparate conceptual structures was important. I find it unaccountable that Mithen would imagine that language was an input to cognitive fluidity rather than an output. I do not follow his reasoning.

Here are more of the facts we need to account for in a theory of the origin of language. Biological evolution happens gradually, but human language appears, in evolutionary terms, very recently. Art, science, religion, and tool use also appear very suddenly, maybe 50 thousand years ago, or 60, or 70, depending on what the archeological evidence turns out to be, and how we interpret scratches on bones and so on. I do not have a time machine and cannot go back to see what was happening 50 thousand years ago.

Anatomically modern human beings have been around from about 150,000 years. That is, organisms that look like us and have the same bone structure and so on have been around for about 150,000 years. But behaviorally modern human beings date from about 50,000 years ago, give or take, depending on the archeological evidence. So we are looking for an advance for modern human cognition, somewhere around 40, 50, 60 or 70 thousand years ago.

There is no evidence of simple languages in any other species. There is no evidence of simpler language in human groups. Children learn complex languages

remarkably easily. But they go through what look like intermediate stages in the grammatical run-up.

None of the proposals we have seen before incorporates all of these facts into a coherent theory of the origin of language. But there are anthropologists Richard Klein, for example, who offers the hypothesis in *The Human Career* that there was some kind of mutation about 50,000 years ago that produced neurological change that made all of these behaviors possible. Klein does not offer a theory of what that operation might have been, but he thinks that such a theory is needed. The proposal I am outlining here has ample room for full linkage across all of these singularities. Social cognition comes up at the same time as language, at the same time as advanced tool use, at the same time as counterfactual thinking. All of these behaviors cooperate and support each other. So conceptual integration theory concurs with Klein that these singularities are linked. No one of them is the cause of the others. Rather, double-scope blending arose and proved to be the underlying mental operation that makes it possible for all of these higher-order behaviors to come into existence.

The Way We Think reviews theories from other fields that bear upon the origin of language. Luigi Luca Cavalli-Sforza proposes that language arose as an invention of behaviorally modern human beings. He puts the date of origin at about 50,000 years ago. In *The Way We Think*, we review proposals that converge to locate a period of rapid cultural invention, producing a coordinated suite of modern human performances, dating from the same epoch, perhaps 50, 60, 70 thousand years ago. We've argued that all of these modern human performances, which appear as singularities in human evolution, are the common consequence of the human mind's reaching a critical level of blending capacity, double-scope conceptual integration. Although we do not have a time machine and cannot go back and verify this, we think from the evidence that exists now of the kinds of languages you can see, of the way in which the children use language, of what we know about the way evolutionary biology works, that this is the best and most defensible hypothesis in the face of all these evidence.

Let me spend a few minutes giving you a closing example, and then I will stop and take questions.

Consider a behavior like writing. Writing is not the same thing as language. Writing depends upon language. Writing is at most eight thousand years ago, more like five, really, and its only been a few hundred years since many people knew how to read and write. I know of no theorist who thinks that there is a mental module that was evolved for writing and reading. It took human culture perhaps 40, maybe 42 thousand years even to invent writing and reading. That is one of the great things about cultural invention. It can take culture quite an extended interval to invent something, but once it does, the invention can be communicated across the culture quite rapidly. Once the invention is established, then the child comes into the culture and because the child has conceptual integration abilities, the invention can be established in the conceptual system of the new organism. Children come into the world able to do double-scope blending, and culture stands ready to feed the child the particular conceptual integration networks it has invented.

You come into the world, and if you are Chinese, you are exposed to Chinese, and

tones, and chopsticks, and Dragon Boat Festival rituals, and because you are equipped to do double-scope conceptual integration, you can pick it up very fast.

The invention of complex numbers took a long time. The greatest mathematical minds suffered with it for hundreds of years. I can teach a 17-year-old complex numbers in a week and they get perfect scores in the SAT test, not because the 17-year-old is a mathematical genius but because you can give them the understanding of the conceptual integration network that underlies the invention of complex numbers very quickly, even though it took a few hundred years to achieve.

Writing is one of those activities that arose in cultural time through the mental operation of conceptual blending. Writing hardly seems like the same kind of thing as a watch or a coin or a cathedral. But when we look at writing, what is it? Well, you see physical marks on a stone, a paper, or a computer screen, and these marks are circulated through the community. Now by themselves these marks are meaningless. We just can't see it that way. When you are a child and you look at writing, you don't see the words, you don't see the language, you just see the marks. You had to work very hard to see those marks as language. That is because the marks are in fact just marks, not language. If we had a sheet with writing on it, and sent it back 50,000 years, it would be useless. I mean, the piece of paper would be a marvel, but the writing would be unintelligible. Human beings fifty thousand years ago would not have known what to do with the marks, because culture had not done the work to invent writing and no one had taught them the relevant conceptual integration networks.

Suppose a woman is reading a letter from her fiancé, who is a soldier at the front. He is in the war. What is she doing? From one perspective, she is looking at distinguished marks on a sheet of paper, but a horse or pig can do the same thing. You can train a horse or a pig to recognize the different marks. Since I know how much cultures can vary, let me ask, was it culturally odd for me to pick out horse and pig as my examples? Do these animals have special meanings in Chinese that I shouldn't have used? No? Am I OK? Good!

Yes, a dog or a cat, or a pigeon, can be conditioned to distinguish marks on paper. But that's not what the young woman is doing. She is clearly doing something quite different from the horse, pig, dog, cat, or pigeon. In the conceptual integration network, there is one input in which a woman is alone and she is looking at a material object, but there is another input in which her fiancé is present and she can hold a conversation with her fiancé. These are very different mental spaces. But they get blended together selectively, so that there is emergent structure in the blend. Now there can be a conversation in the blend, but it's a special kind of conversation. He cannot answer questions that she would put to him. But in the blend, one can imagine what the other person would ask, and so you can say, but we learn in the blend that we can write, we can imagine what the other person is thinking, and she can write "of course you will not be surprised that I am wondering when you will come home..." So the blend can try to build in some ability for compressed question-and-answer, back-and-forth.

This is really, really amazing. Notice that in a real conversation, there would be audible sound from the fiancé. The fact that writing consists of words comes from the space of speaking. The specificity of those words and marks comes from the space with the specific marks on the paper, so there are some marks here and some marks here. In the blend, these are understood as different words because of the different specific marks. When she looks at the letter she receives from her fiancé, she doesn't just imagine what the fiancé is saying. She looks at the marks, and the marks specify what the fiancé is saying, so the specificity comes from the marks, while the general capacity to speak comes from what she knows about the fiancé.

Notice that the marks bear a complicated relationship to the speech. If you see a mark like *boy*, b-o-y, you know that is supposed to connect to a sound, [bɔi], or more precisely, there is a whole category of marks, very many of them, that are supposed to be blended with the sound [bɔi]. They are all different marks, but we take them as all belonging to a category, and as connected to a category of sounds, because different people can pronounce the word *boy* in different kinds of ways. The result is an astonishing compression.

A proficient reader ends up with a general blending template for writing and reading. It seems to us as if we are just expressing ourselves through reading and writing, but in fact these are elaborated double-scope compression networks and of immense cultural importance to us.

So we see that writing is possible only because we can construct and learn double-scope integration networks.

But now that we see this, we can take one step back, and, lo and behold, see that speech operates in the same way! Consider the scene in which the woman is actually listening to the speech of her fiancé. That seems like a basic human scene. It really works for us. We are human beings. We can do this. Suppose the fiancé comes home from the war unscathed, and she and her fiancé are having coffee in the kitchen. From one perspective, what is happening is just that longitudinal waves of air are striking her ear. I can make many different sounds, specific sounds that you distinguish, but that do not count for you as language. The making of distinguishable sounds is one scene. Dogs can understand that you are making various sounds for them to recognize. That's not language. But in the other scene, she is understanding language. The sounds are blended with the language. So I can say "boy" (loud) and "boy—" (very loud and long) and "boy" (normal) and "boy"(whisper), fine, all these different sounds. They can all be blended with a particular word, *boy*. The sounds and the words are very different kinds of things, but in the blend, we fuse them.

Sign language is similar. I will just refer you to Scott Liddell's work. Liddell has elaborate examples of the way in which double-scope blending makes sign language possible.

One phenomenon that never gets talked about in cognitive linguistics—this is my last point for today—is the concept of style. But yesterday, after the lecture, someone asked me brilliantly, "how do we include all the history of pragmatics?" Every human

performance is done in a style. It is impossible to act pure and simple, without a style of action. We typically do not recognize the styles in which we speak, or the styles in which we walk. But just watch: go out of this lecture hall onto the campus and watch people walk. Everyone has a style of walking.

It is the same with speaking, and conversing. Style drives communication. I will refer you to a book, *Clear and Simple as the Truth: Writing Classic Prose*, by Francis-Noel Thomas and Mark Turner. It shows that style is a result of double-scope integration. In a style, or at least classic style, one takes a particular scene of communication and blends it with much more diffuse networks. The basic scene of classic style is conversation in a scene of joint attention. All writing is assimilated to that scene. It is a highly human-scale scene. Vast and diffuse ranges of communication are compressed by borrowing the compression that comes with the scene of joint attention. We are all used to inhabiting particular scenes, as a way of getting along in communication. You all, for example, for the last hour and a half, have been inhabiting the role of audience. And you have been a great audience. Thank you very much!

Lecture Ten

Grammar

So this is in fact the last of the lectures, and I am surprised it has gone so fast. I think we should have another set of lectures. But we must return to the C Major of this life, and recognize that this is the last lecture. There are some things to wrap up.

Let us thank our institutional hosts! We are very grateful to the Beihang University, the Beijing Language and Culture University, the Beijing Forestry University, the Beijing Foreign Studies University, Peking University, and Tsinghua University, where we are today.

We would also like to thank the Foreign Language Teaching and Research Press, which is publishing all of these lectures. I would also like to thank all of the volunteers who organized this conference: Hu Ya'nan, Wu Shan, Li Lingmin, Ma Sai, Wang Fan, Xiong Liqin, Yang Jie, Yin Shuying, Yuan Wenjuan, and Zheng Lingyan. If I have left anybody out, please forgive me: I have not met all of you, but I will meet all of you on Sunday night at the great farewell dinner.

And of course the biggest thank you of all has to go to the *fons et origo*, the source, the driver of this initiative, who is of course Professor Li Fuyin. Thomas, thank you.

Today we are going to talk about grammar. Yesterday, I outlined a theory of the nature and origin of language, as a relational system of form-meaning pairs, used as prompts to ask for the construction of meaning. These constructions prompt for various kinds of integrations. Since the conceptual integrations to be achieved are often very many, with ranges of words and expressions attached to the various mental spaces, language must make it possible to achieve these combinations.

Consider “The beach is safe”. This expression asks you to construct a conceptual integration network, but it does not ask you to compose the meanings supposedly possessed by the words. Very little meaning is encoded in words. Words and expressions do not mean. They are prompts to people to construct meaning. Languages underspecify very seriously the meaning that you have to bring to bear to construct a meaning. As we saw, it is not the case that “safe” carries a meaning that it predicates of something. “Safe” asks us to construct a counterfactual space and then compress outer-space relations between the two spaces that are counterfactual to each other so that the counterfactuality is compressed to a feature, absence of harm, for something in the blend. The relevant situation is now to be understood as having this feature. There are many similar examples of compounding that we have discussed and that other cognitive linguists have discussed: *likely story*, *possible solution*, *eligible bachelor*, *fake gun*, but also *red ball* and *brown cow*.

Language is a relational system of form-meaning pairs, that is constructions, that themselves integrate for the purpose of prompting for the construction of conceptual integration networks. As we discussed yesterday, the reason that we can have form-meaning pairs is that we have the ability for double-scope conceptual integration. Advanced conceptual integration makes it possible to have language, but language is

useful for prompting people to do more conceptual integration. Language is a complex dynamic system. This is a notion with which you are familiar from many areas of Cognitive Linguistics. Language can also be thought of as a complex adaptive system. The primary statement of the role of conceptual integration in grammar is presented in *The Way We Think*. See <http://blending.stanford.edu> for details. There is a review of *The Way We Think* by Professor Li Fuyin in *Contemporary Linguistics*.

Complex grammatical structures prompt for integrations. Let us look at some of them.

Y-of Constructions

Consider a sentence like *Ann is the boss of the daughter of Max*. Now, English, in a way that is almost too good to be true, provides us with a linguistic construction whose purpose is to prompt for just the kinds of integrations that we have been looking at. It is the *Y-of construction*. The integrations that it asks us to do are not compositional. Indeed, compositionality is not an appropriate expectation for meaning. Rather than compositionality, we have selective projection and emergent structure. However, the *forms* can be compositional. Indeed, the *y-of construction* can compose repeatedly: *y1 of y2 of y3 of . . .* This composition of forms asks us to construct meanings, and the meanings are not compositions of the meanings encoded in the various expressions.

When we encounter a form, we construct a meaning, and it is an easy but inaccurate compression to think that the form carries the meaning. We see a symbol and say, “Oh, this means Tsinghua University”. But no, it is just a form. We are the ones who are coming up with meanings. Or I show you a little picture, two inches square, of only a head, and I say, “This is my son”. We see the form—the photograph—and come up with meaning, and again it is an easy compression to think that the form carries the meaning. We are not deluded, but sometimes it is easy to fail to separate form and meaning, because, of course, to be native speakers of the language, what we must do is compress and collapse them all the time.

Let’s stay with this point for a moment. What you are hearing right now is just a bunch of sounds, the gestures you are seeing are just some photons striking your retina, but it feels as if you are hearing meaning. In fact, it feels as if I am inserting meaning directly into your thought, as if you are more or less passive, doing no work. On the contrary, you are in the presence of forms, visual and auditory, and you must do all the fabulous semantic backstage cognition in order to attribute meaning to that. It doesn’t see to you as if you are doing much of anything. That is because you are double-scopers, cognitively modern human beings.

So here we go. When you encounter a *Y of* construction in English, it’s just two words, but that form is paired with a meaning, and the meaning is an instruction to you to perform an elaborate mental operation. A *Y-of* construction says, “Hey, I am a *Y-of* construction, and have a big job for you to do, and I want you to do it really fast and I want you not even to notice that you are doing it. I want you to use your advanced

double-scope cognition to put together a certain kind of conceptual integration network”. You are to call up an input space that contains a frame that contains *Y*. *Y* prompts for some frame. It can prompt more than one, in which case, the *Y-of* expression can be found to be ambiguous. But you know that it is your job to come up with some frame that plausibly can be called for by the form *Y*. You know additionally that you are supposed to construct a blended space, but you don’t know what it is yet. You are supposed to project the element *y* from the input space to create an element *y’* in the blend. You are supposed to come up with a *w* that is not mentioned at all. It’s not mentioned in the language at all, but you have to come up with that *w* in the input space that contains *y* that will bear some appropriate relationship to the *y*. You know that you are asked to project the *w* to create an element *w’* in the blend. You know that you must project the *y-w* relationship onto *y’-w’* in the blended space selectively. You know that you must construct open-ended connectors from *y’* and *w’* in the blend. We expect these connectors to make connections at some point, but at present, they are just hanging. And we expect the open-ended connector from *w’* in the blend to connect to something picked out by the noun phrase that will follow *of*. To summarize, a “*Y of*” expression prompts us to perform the following operations:

- Call up an input space for the relational frame containing *y* (the element named by *Y*).
- Construct a blended space.
- Project the element *y* to create an element *y'* in the blend.
- Provide for a *w* in the input space that will bear an appropriate relationship to *y*.
- Project that *w* to create an element *w'* in the blend.
- Project the *y-w* relationship onto *y'-w'* in the blended space.
- Provide open-ended connectors from *y'* and *w'* in the blend. We expect these connectors to make connections at some point.
- Expect the open-ended connector from *w'* in the blend to connect to something picked out by the noun phrase that will follow "of."

None of these instructions is encoded in the language, of course. We just say *Y-of* and you know what to do. A *Y-of* expression does not encode a meaning; rather, it asks you to go through a process to construct a meaning. And you do.

Here are some examples.

- Paul is the father of Sally.
- Necessity is the mother of invention.
- Tsinghua University is the MIT of China.
- This is the top of the building.
- Every minute now should be the Father of some Strategem. (Shakespeare)
- Zeus is the father of Athena.
- Joseph is the father of Jesus.
- I’m your father for today. (I’m the father of Sally for today.)
- You are my long-lost daughter.
- He was the Einstein of the fifth century B.C.
- The adjective is the banana peel of the parts of speech.

- Valet of the secretary of the president.
- The wife of the secretary of the president is the most important position in the government.
- Fear is the father of violence.
- Language is the fossil poetry of the soul.

Take *Paul is the father of Sally*. Here, the usual choice for the missing *w* is daughter. That's easy, because "father" is one of those words asking you to come up with a relationship between two things. Where there is a father, there is a child. But notice that the child, the daughter, is not mentioned, and does not have to be. You find a *w*, and you project from *y-w* into a *y'-w'* in the blend, and that *y'-w'* is compressed with Paul and Sally from a separate mental space, so now in the blend *Paul is the father of Sally*.

Consider now *necessity is the mother of invention*. The meaning here is quite different from the one we construct for "Paul is the father of Sally". "Necessity is the mother of invention" feels highly metaphoric. But note that the *Y-of* construction is prompting you to do the same mental work in both cases. Consider *Tsinghua University is the MIT of China*. No problem—everybody understands immediately. So the missing *w* for MIT is the United States. I didn't have to mention it. You have to go find it. Sometimes, it's not so easy to find the missing *w*, but you know you are supposed to find one, and you do find one. It doesn't have to be mentioned at all.

Note also that I don't even have to say what is supposed to be projection from Tsinghua University and China and MIT in the United States into the blend. You do all that work. I just say it to you, and you go do all the work, and it seems to you as if this form carries all the meaning you constructed, which it doesn't. Instead, the expression is something very small and thin. It prompts for a robust construction, and you oblige. You do all this work, but do not notice the work you are doing. Language is in this way much like vision: you look at the world, and it seems as if you are not doing any work to see.

We easily understand Shakespeare's "*Every minute now should be the Father of some Strategem*" to mean that we should come up with cunning plans all the time. *Zeus is the father of Athena. Joseph is the father Jesus*.

I once heard, "I'm your father for today," said by a friend of the family to a girl take-your-daughter-to-work day. This is a kind of *y-of* construction ("I am the father of you for today"), but using a different form: In the United States, there is, or was, a day on which some groups suggested that you take your daughter to work, so she can see what work is like. But sometimes, fathers travel, or they are absent for some other reason. So one thing that may happen is that a friend of the family or an uncle might take the daughter to work. One of these people said to a seven-year-old girl, "*No problem, I'm your father for today*". And nobody was confused. Everybody knew what was supposed to be projected from what frames. Notice, no one thought that this meant the man was the girl's father outside of the blend.

Consider *You are my long-lost daughter*. Again, this is like a *y-of* construction, "You are the long-lost daughter of me," but using a different form. The setting is this:

There is a man who had a daughter who died at age ten. He is quite sad. He has a younger daughter, Sally, but he is traumatized by the loss of his older daughter. As a result, he is fairly cold, and somewhat unreceptive to Sally. But as time goes on, Sally grows to be 12, 13, 14, and he realizes that Sally has now reached an age that his older daughter never reached. He begins to contemplate life and what this means and he begins to behave much more warmly and receptively to Sally. And she says, “What’s going on, Dad? What’s happening?” And he says, “*You are my long-lost daughter*”. This is the equivalent of *you are the long-lost daughter of me*, and is a variant of the *Y-of construction*. Of course she has been his daughter all along, so the structure prompted for in the blend by “long-lost” is not kin relation. It can receive projection from both the relationship of the father to the second daughter and the relationship of the father to the first daughter, including the daughter in a counterfactual space in which she is still alive. This is quite complicated.

Other examples are

- *He was the Einstein of the fifth century B.C.*
- *The adjective is the banana peel of the parts of speech, because you can strip it off.*

The missing *w* in the first example is *twentieth century* and in the second example is *banana*.

Consider *The valet of the secretary of the president*. This involves frames for roles. This example shows that *y-of* constructions can compose at the formal level.

Here are other examples we will discuss:

- *The wife of the secretary of the president is the most important position in the government.*
- *Fear is the father of violence.*
- *Language is the fossil poetry of the soul.*
- *Las Vegas is the American Monte Carlo.*
- *Washington DC is the American Beijing.*

I made that one up. It is very easy for us to invent instances of these *x is the y of z constructions*.

- *Las Vegas is the American Monte Carlo.*
- *Social movements are at once the symptoms and the instruments of progress.*
- *As poetry is the harmony of words, so conversation is the harmony of minds.*


Counterparts in these networks not be metaphoric: “the nation of England,” “the island of Kōpīpi,” “the stigma of cowardice,” “the feature of decompositionality,” “the condition of despair.”

Charles Fillmore gives the example: “*One needn’t throw out the baby of personal morality with the bathwater of traditional religion.*”

So what is happening in all such cases? Well, *Y of* says, “Call up as an input space a mental space of meaning that can plausibly be prompted for by *Y*. You are going to have to find a missing *w*. It is unlikely that *w* will be mentioned. You are prompted to project elements from that mental space selectively into the blend, with open-ended connectors from *y’* and *w’* in the blend. And you do not know what they are going to

connect to until you have some prompt or some knowledge of what they are to connect to.

But if I say, if I say “*Paul is the father of Sally*”, then *Paul* and *Sally* serve as an *x* and *z*, so we have not just *y-of*, but a full *x is the y of z*: *Paul is the father of Sally*.



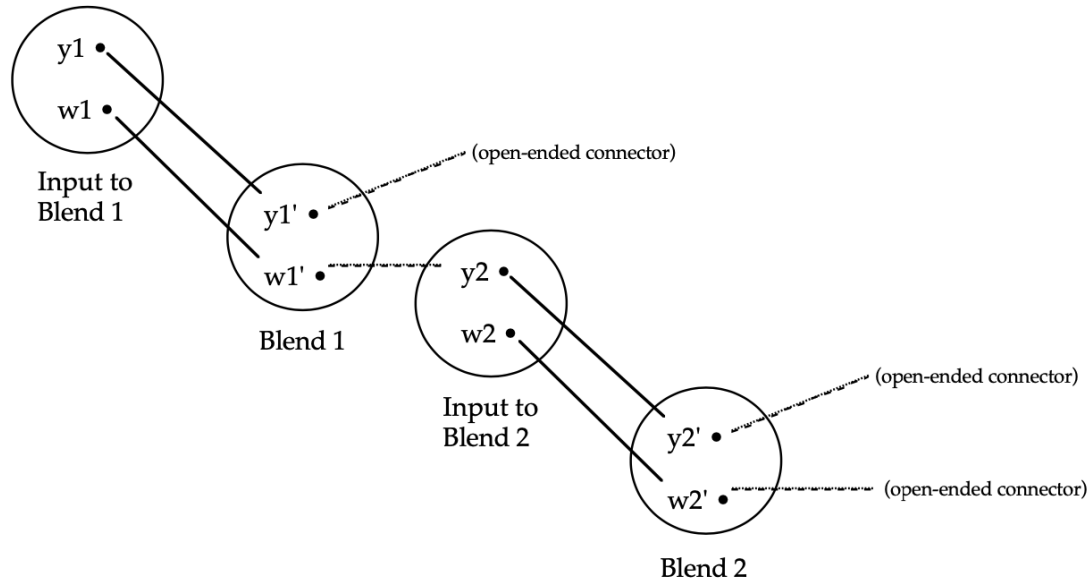
In the blend, Paul and Sally are related: Paul is the father of Sally. As we have discussed, we have a new role in the blend that is not available from either input space: father of Sally. This could be a very interesting social role. You could say, for example, *He is Sally’s father this year*, meaning that two years ago, the mother was married to somebody else, and now she’s married to this man. This role, father of Sally, can have social ontology, social consequences. He might, for example, now have legal responsibility for her college expenses or her monthly bills.

Composition

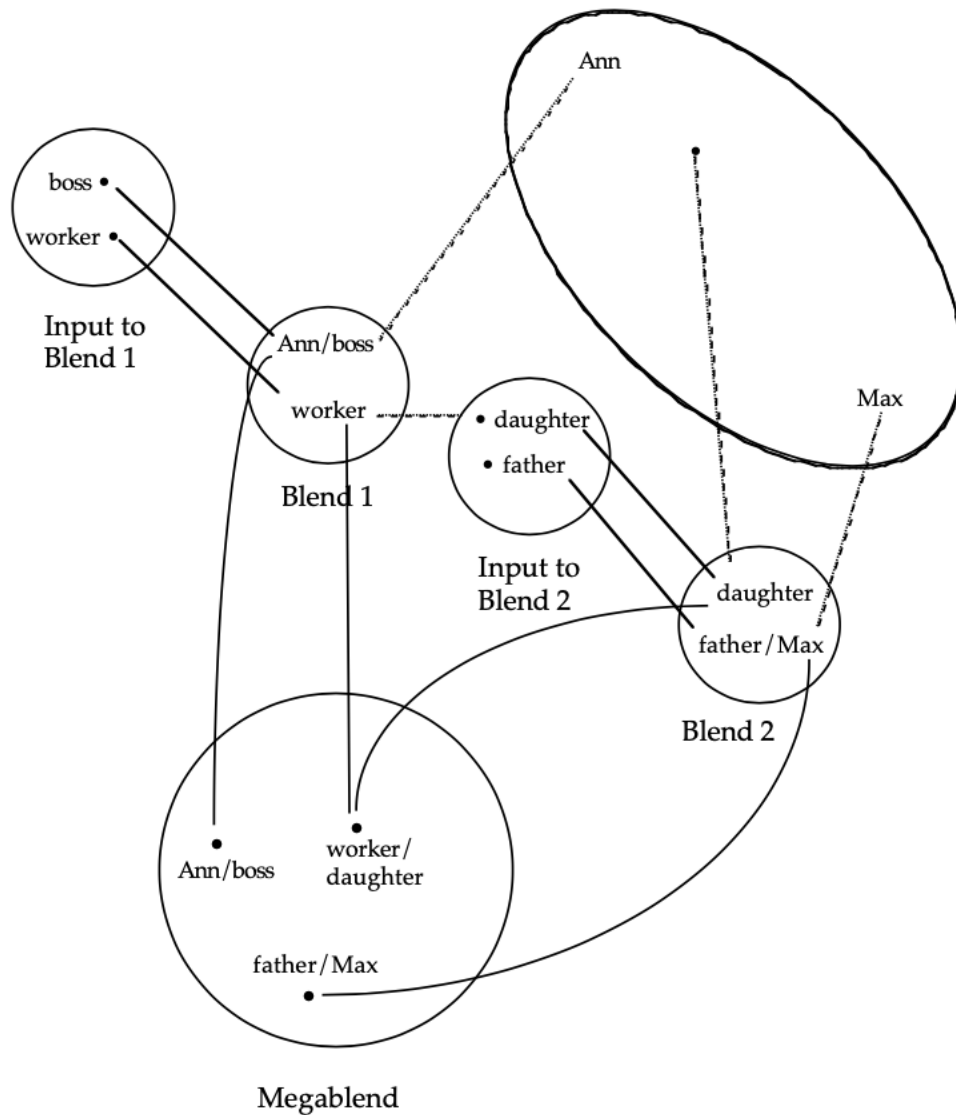
Y-of expressions themselves can be composed at the formal level. Specifically, what follows the *of* in the first *Y-of* expression can be another *Y-of* expression, and this can be repeated for as long as we like. We can say “*The doctor of the sister of the boss of Hieronymous Bosch.*” Nobody has any trouble with the composition, provided we do not create a cognitive load too high to be carried. You use this, and nobody has any trouble. In such a composition, one does not know how the open-ended connectors will be closed off until we reach the end. In such a composition, there are many missing *w*s. An expression of the form *the y1 of the y2 of the y3 of the y4 of...* is not uncommon. Because they have open-ended connectors, we can add *y-of* expressions repeatedly. The open-ended connectors of one *y-of* network can attach not just to an element but also to an entire *y-of* network. Accordingly, the formal composition can prompt for repeated conceptual mapping and blending operations. Do not confuse the compositionality of the forms with compositionality of meaning. The forms are

compositional, but the meanings are in general not. Putting together the meaning prompted for by a composition of *y-of* expressions is a highly creative act, even though in consciousness we might not recognize any of that creative work. The meanings are not the result of cut-and-paste composition.

Here, in the slideshow, is a representation of the *Y-of* squared network.

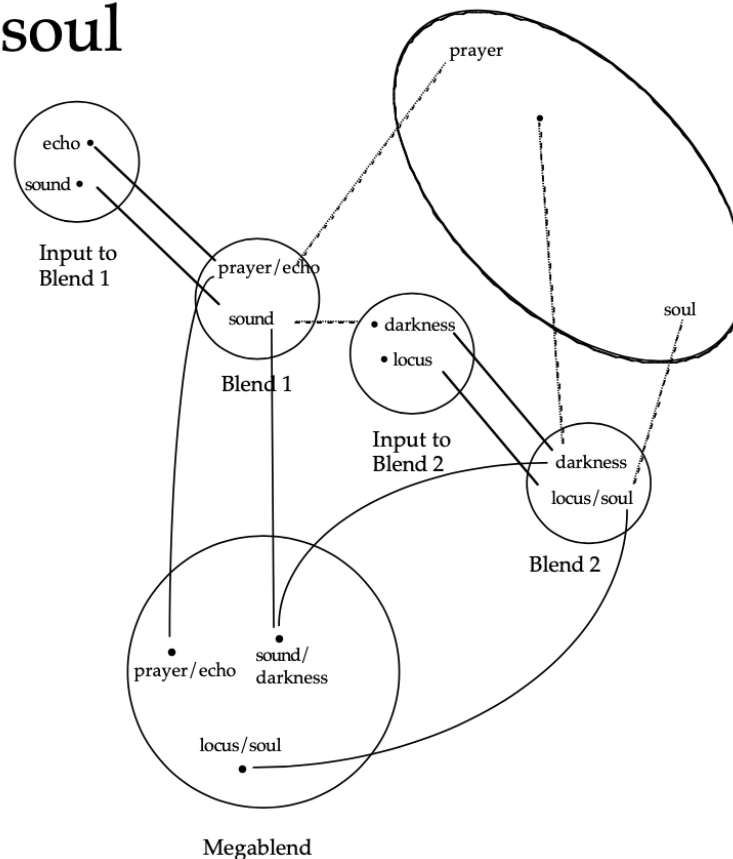


This is the abstract network prompted for by something like *the valet of the secretary of or the wife of the president of*. Let's consider an example of a *Y-of* squared network: *Ann is the boss of the daughter of Max*.



Notice that we can use the same form for quite a different meaning, as in *Prayer is the echo of the darkness of the soul*. Very poetic, but what could that mean? I'm not so sure, but we can all come up with interpretations, prompted for by the same form, calling for the same patterned network.

Prayer is the echo of the darkness of the soul



As an example of a missing *w* would be the one that is needed for *echo of*. The *w* people usually find is *sound*.

An “x is the y of z” form is a syntactic grammatical form related to others in a more general category of what I call *xyz constructions*. One doesn’t always have the word *of*, for example, as in “Sex is the poor man’s opera”. There are other prepositions possible, and a variety of compositions, as in “A leader is a dealer in hope” or “Adversity is the first path to truth”. Chapter 9 of *Reading Minds: The Study of English in the Age of Cognitive Science* gives many other examples and related forms.

XYZ constructions are extremely powerful and can call for highly creative conceptual integration. To count as a native speaker of English, one must be able to follow the prompts provided by an xyz construction to build a network. Of course, if you lack some of the cultural frames of meaning, you may have difficulty constructing a meaning, but that is not because you do not understand what the form is prompting you to do.

Quite interestingly, the same form, the same set of prompts can work over essentially any conceptual domain. The highly usual “*Paul is the father of Sally*” has the same syntactic form as the highly unusual “*Vanity is the quicksand of reason*”,

which some people may find unintelligible. *x is the y of z* can now be used across all kinds of conceptual integration networks, including simplex, mirror, frame-compatible, single-scope, and double-scope. It is a general form for prompting for a great deal of conceptual integration.

Let's have a look at "*Vanity is the quicksand of reason*". The pattern of the network is no different from the one we find in *Ann is boss of the daughter of Max*. In one space, you have some kind of conceptual structure that can be prompted for by using words like *x* and *z*. In this case, the *x* is *vanity*. Your job is to find an input space that contains *x* and *z*. One obvious one is that you have human traits like reason and vanity. The *y* is *quicksand*. You know that your job is to call up a frame that contains *y*. There could be many of them, but a standard one is travel that involves falling into quicksand. This is the frame in which people and animals and jeeps and so on get trapped in quicksand. Note that the construction of the meaning does not require that you believe the meaning is true. For example, I have some experience of deserts and know that quicksand is a mixture of sand and water. That's why you fall into it. But there seems to be a more widespread notion that quicksand occurs in areas that are completely arid. You see this notion used in the movie *Lawrence of Arabia*. In that notion, people can fall into quicksand in the driest part of the desert. So let's use that frame for quicksand—the one in which dry desert swallows the traveler up.

In this situation, you can say to yourself, "Aha, the missing element *w* is going to be a traveler." Notice that the traveler it's not mentioned in "*vanity is the quicksand of reason*". Then you project, selectively, *quicksand* and *traveler* and the relationship between them to the blend, and blend them with *vanity* and *reason*. One very standard interpretation of "*vanity is the quicksand of reason*" is that if you are vain, if you have a big ego, it is going to diminish your ability to reason: your reason will not work well. This conceptual integration network can recruit elaborate structure from other integration networks. Eve Sweetser, in *From Etymology to Pragmatics*, brilliantly analyzes the conceptual integration network of the Mind as a body moving in space. Using this blend, we can view reason as a traveler, who makes progress, encounters impediments, and so on, as in "I am advancing towards a solution." The conceptual integration network for the mind as a body moving in space is elaborate and entrenched. If we use it, then we can see vanity as an impediment to reason, the traveler.

It is important to see that there are major conflicts between the vanity-reason frame and the quicksand-traveler frame. I discuss these conflicts in detail in an article titled "Figure," which you can download from the Cognitive Science Network: <http://ssrn.com/author=105812>. The projection to the blend is highly selective. The blend takes parts of each frame, blends them, and makes a new frame with emergent structure.

It is important to see that double-scope blending lets us blend frames that are highly conflicted. But this does not mean that conceptual integration is unconstrained. On the contrary, it is a very highly constrained operation. Do not be misled by the amazing power of double-scope integration into saying, "Oh, I see, conceptual integration just lets us do anything. We just throw things together" No. As you have

seen, there is a set of constitutive principles. Something that does not meet those constitutive principles fails as a successful network. And there is a set of governing principles, like topology, web, good reason, and on and on and on. And these constrain the work you can do, but conceptual integration is not all algorithmic. If you were to take two frames and blindly compute all the possible combinations of structure from the two frames, the result would be extremely high. Almost none of those blends will ever occur in human thought. Very, very few of the in-principle combinations will ever become conceptual integration networks. But we are not built to see all the combinations that we do not make. We can look at some of the ones we do make, and their range is amazing to us, so it seems as if conceptual integration has sweeping, unconstrained powers. On the contrary, the constraints are powerful and many.

It is easy, for example, for a blending theorist to give examples of bad blends, blends that do not work for people. Let's force ourselves to construct a bad blend here, for "*vanity is the quicksand of reason.*" I mentioned that the usual choice for a missing *w* is traveler. But how about something else? Suppose I tell you to use *bacteria* as the missing *w*. Now try blending the vanity with quicksand and reason with bacteria. The answer to this request is usually, "What do you mean?" Most people can't find a place to start. One constraint on building this xyz network is that the missing *w* must be available from the frame containing *y*. If it is not, then you are unlikely to be able to make the blend. But what is available depends upon the situation you are in and activation, that is, what is active in your mind. So suppose I say, "Did you know that some bacteria can live only in quicksand and they depend on it for everything?" Well, now you have conceptual structure involving biological relationships and habitat for quicksand and bacteria. It is part of a more general frame you have—for example, Torrey Pines are a kind of pine tree that live only in San Diego, California, near where I grew up. With this frame, you can interpret this way: there is a certain kind of bacteria that can live only in quicksand; this bacterium depends on it for everything; otherwise, it dies. For these bacteria, quicksand is not a trap; it is what they need to live at all. In the blend, well, for some people, vanity is the quicksand of reason. Their vanity gives them the self-confidence to think well. If they are not being vain, then they are not confident, and they just don't have any confidence, and they are intimidated, so they can't reason well.

This blend might be a little stretched, but now it is possible, because we have managed to find a suitable missing *w*.

Xyz constructions do not operate willy-nilly, but neither is their range narrow or their operation algorithmic.

From our discussion yesterday of *brown cow*, we are familiar with the idea that a given form can prompt for very many different constructions of meaning, including vast ranges of meaning for which that form has never been taken as a prompt. The range of conceptual structure that can be prompted for by "*My cow is brown*" is vast.

Conditionals

Let's consider some of this range of meaning construction from a single construct:

If I were you, I would quit my job. This construct is already an integration of many different constructions: all those words, the subjunctive form of the verb for the first part of the “*if then*” combination, the use of the conditional on “*would*” for the second, and so on. Many grammatical constructions must be integrated to make this construct. This is just the usual view taken by varieties of construction grammar. Now, it might seem that *If I were you, I would quit my job* prompts for a unique blend, but in fact it can prompt for very many different blends depending on the selective projection and the emergent structure. For example, I could follow *If I were you, I would quit my job* in these ways:

- . . . but I am independently wealthy; you shouldn't quit by any means.
- . . . but I am a hothead and would regret it later and would have to go on my knees begging for my job back.
- . . . and so should you.
- . . . but you shouldn't.
- . . . but that's only because the boss needs me so much he would offer me a raise to get me back.
- . . . since I couldn't live with myself knowing how badly I had treated me.
- . . . because being you would make me so utterly miserable I couldn't possibly get any work done.
- . . . since I would have a wealthy father.
- . . . since you have another job offer.
- . . . since I have another job offer.
- . . . since your beloved boss has another job offer and will be leaving soon.

Notice that in different cases, different people are quitting the job in the blend. In some cases it is the speaker, but in others it is the listener. In different cases, quitting the job is judged to be wise versus foolish. In the different blends, there are different reasons for quitting the job. Consider *If I were you, I would quit my job, since I couldn't live with myself knowing how badly I had treated me.* It may seem ungrammatical, but suppose we are in the context in which someone is saying this to his boss and they are disagreeing. Now, *myself* refers in the blend to the *boss*. So it is fine. In this blend, the boss has his the boss's conditions of life and the psychology of the worker. So the sentence becomes grammatical, because the *myself* is the boss while the person who is treated badly is *me*. You know where the *myself* goes and you know where the *me* goes.

Consider *If I were you, I would quit my job, since I would have a wealthy father.* In this case, the addressee is the one that has the wealth. But, by contrast, in “. . . but *I am independently wealthy; you shouldn't quit by any means*”, I was the one with the wealth. We get quite different meanings in these two blends, both prompted for by *If I were you, I would quit my job*.

In *If I were you, I would quit my job, since you have another job offer*, it is the addressee in the blend who has the addressee's conditions and the speaker's reasoning and quits the addressee's job. But in *If I were you, I would quit my job, since I have another job offer*, it is the speaker in the blend who has the addressee's conditions and

the speaker's reasoning and who quits the addressee's job. For another contrast, *If I were you, I would quit my job, since your beloved boss has another job offer and will be leaving soon* has a blend in which the addressee has the addressee's conditions and the speaker's judgment and quits the addressee's job. There are many possibilities, open-ended, depending on context and varieties of projection and compression. But that does not mean that blending throws together just anything. Most of the in-principle combinations never occur to us. The constraints are severe, but the meaning process is not compositional or algorithmic. Human beings are creative in these ways, all the time, in everyday thought, action, and language, and the theory must be true to those phenomena. Words do not mean. Words do not have or carry meaning. Words, and constructs, are prompts that we use to call up mental work to construct meaning. We use basic mental operations we already possess to work on things that for the most part we already know.

As we discuss in *The Way We Think*, in a theory of meaning, activation does not come for free. The existence of frames, knowledge, experience, scenarios, and memories does not come for free. Ease of activation and degree of entrenchment by themselves impose very strong constraints on the imagination and the use of language. Linguists, logicians, and for the most part even psychologists tend to focus on the entrenched cases. There is a common-sense view that, methodologically, we must first explain the supposedly very simple cases, like *brown cow*, or *red ball*, or *the cat is on the mat*. In this common-sense view, we should not bother ourselves with cases like *Prayer is the echo of the darkness of the soul* or *Language is fossil poetry* or *I'm your father for today*. The idea is that we will get to that in due time, maybe a decade from now, or a hundred years from now, or three millennia from now. But the truth is that we never get to it. "We'll get to that later" has meant "We will get to that never." The reason we never get to it by this additive, incremental approach is that the approach is false to the nature of meaning and language. Double-scope blending is not something that is added as children develop. Children do it from the "get-go," from the start. They do not build up from very simple grammar that doesn't have any complexity: the forms they can produce may be few, but they use them with equipotentiality, and double-scope blending is going full speed from the start. That is indeed how they can acquire complex grammars. It might seem as if our scientific task would be simpler—I don't actually think it would be—if in fact children started with "simple" cases and simple operations and only later worked up higher-order operations. But the simple stuff isn't simple at all. In order to understand something like *brown cow*, you have to have double-scope integration going straightaway to integrate forms with meanings in the way human beings do it rather in the limited ways in which other species can use communicative prompts. In order to understand the *Y-of*, you have to be going full bore with double-scope integration, not just to do the conceptual integration, but also to understand how to put together constructions into a construct. You can't start with the supposedly easy stuff as if it is limited when you are dealing with mind and language. Starting with the entrenched cases does not make our job easier or our science more respectable. Quite the contrary. We are lulled by

the familiarity of these cases into blindness: we don't see the complexity going on; we slip in lots of complicated meaning construction as if it can simply be assumed, or not even recognized. We take for granted what should not for a second be taken for granted.

We lull ourselves in the entrenched cases into thinking that the meanings are predictable from the forms. The lulling consists of making a compression and then thinking it is scientifically accurate. It's a useful compression, but bad science.

The Way We Think discusses, for example, the very many ways in which a word like *house* can be used. It seems to us that in understanding an everyday use of *house* we are not doing any conceptual blending. This is probably why linguists, logicians, and analytic philosophers of language have often been blind to the complex creativity involved in everyday language and so felt comfortable excluding visibly inventive, figurative, creative, and literary examples from their data sets. On the contrary, this mistaken view that only predictable composition of meaning can be scientifically tractable and important, and that only predictable composition of meaning can support genuine rational thought is fundamentally false. When you begin by looking at only the entrenched cases, it's not just that you are missing what is going on in the obviously complex cases. You are also missing what is going on in the very simplest cases. It takes a lot of work to see what's going on at all in human thought and language of any variety. We are not built to look into it, and we need the benefit of visibly provocative data to begin to have insight to the general patterns underlying cognition and language.

You might have thought that the most boring thing in the world is the *Y-of* construction. In fact, it turns out to provide a spectacular demonstration of the creativity of conceptual integration.

Individual words can also be prompts for elaborate conceptual integration networks. We saw this for a word like *safe*. If you ask whether the child on the beach is in danger, I can respond, "*The beach is safe*", or I can respond, "*The child is safe*". And I can mean the identical thing by saying those two sentences. So if you thought that *safe* somehow predicates of feature of the subject, that's wrong, because if that's true, then predicating that feature of the beach as opposed to predicating that feature of the child should produce two quite different meanings. But of course, *safe* does not predicate a feature. Instead, either sentence in context prompts us to call up the scene with the beach and the child, create a counterfactual space by blending the current scene with a harm scenario to create a blend in which the child is harmed, understand that this scene is counterfactual to the current scene, and then compress the counterfactual link so that in the blend of the newly-understood current scene, we have the current scene with *absence of harm to the child*. Indeed, perhaps obviously by now, the single word *danger* also prompts for an elaborate conceptual integration network.

Lucky also prompts for a conceptual integration network. If we say, "We are lucky," then in the blend that counts as the new understanding of the current scene, there is a new feature—we are lucky. This, too, arises by compression of a counterfactual.

It is easy to imagine that we know what is in a scene, but, as we discussed before,

we do not conceptualize all the absences. If I say, “I’m safe,” in this lecture room, you can confirm it, but it is not as if the *absence of harm* was already part of your conceptual structure for conceiving of the current situation. In saying, “I’m safe,” I am prompting you to build that structure explicitly into your understanding of the current scene. Your new understanding is not the same as your old one. Your new understanding requires a complex blend. If I say, “*I’m really lucky to be talking to such a great audience at Tsinghua University*”, I am asking you to make a new blend for conceiving of the present situation. Previously, we were just talking. Now, in the new understanding, there’s no big difference in what is happening, but there is a counterfactual scene, which takes blending to achieve, in which I am not talking to such a great audience in Tsinghua University. You are to understand that this space is counterfactual to the original mental space, and compress that counterfactual link so that your new understanding of the present scene includes the fact that I am lucky.

Simple words and expressions like *accident, safe, in fact, brown, red*, all of them, prompt for the constructions of integration networks. We can go through various complexities in each case, just as we did for the *y-of* construction.

Consider, as an example of this complexity, “*All the jewels are safe*”. What I might be trying to do is to prompt you to construct a conceptual network according to which, in the blend, there is absence of harm to the owner through the destruction or theft of the jewels. But suppose I am shipping the jewels, and I say, “the packaging is safe”. I do not mean, or I do not have to mean, that the packaging is not going to come to harm. On the contrary, the packaging itself might sustain considerable harm, as a way of preventing damage to the jewels. I can say, “the packaging is safe” to mean that the jewels won’t get harmed. Coulson and Fauconnier give similar analyses for expressions like “fake gun”. We have also talked about words like *missing* and *gap*, which prompt for conceptual integration networks.

Clausal Constructions

We saw how the same clausal syntax—noun phrase, verb phrase, noun phrase, prepositional phrase—, a caused-motion syntax, can prompt for an integration in which we use the frame of caused-motion, in which an agent takes some object and performs an action on the object that causes the object to move in the direction. *Throw* is a standard caused-motion construction.

But the caused-motion construction can be used as a prompt to ask for compression of elaborate and diffuse structure to a human-scale scene, which borrows the compression from the caused-motion frame that is part of the caused-motion construction. The clausal syntax attached to that frame in the construction can be projected down to the blend and now pick out structure in the blend. Accordingly, I do not need new language to express new meaning that emerges in the blend from the integration of the caused-motion frame with the elaborate and diffuse conceptual structure that needs to be compressed. Again, this is the method by which we solve the central problem of language—how to use relatively few grammatical constructions with equipotentiality, to express vast and open-ended ranges of meaning.

For example, we can use the caused-motion construction not only with caused-motion verbs, like *throw*, but also with other verbs, as in Goldberg's example, "Paul sneezed the napkin off he table," and other verbs, as in "He floated the toy boat to me." Note that *float* is not a caused-motion verb like *throw*. Yet you understand that, in saying, "*He floated the boat to me*", I am prompting for a caused-motion scene. *Read* is not a caused-motion verb, but I can say, "*I read my son to sleep*", where here the caused-motion is metaphoric, to a metaphoric location that is the goal, a state of being asleep. I am the causal agent, and what I did was to read. The effect was that there was an event of motion in the direction of a goal, but, in this case, the motion and direction to a goal integrates with change of state. So his metaphoric motion is moving from one state to another.

We can even say, "*We blocked him from the door*", despite the fact that *block* is a verb for indicating the stopping of motion rather than the causing of motion. We can perform conceptual integration with the caused-motion frame as in input—prompted for by the caused-motion syntax—with meanings prompted for by the various words. We call up caused-motion, but also some frame that can be suggested by the word *block*. In the blend, we have an agent who performs an action, not of causing an object to move in a direction, but of stopping something from moving in a direction. The blending here is highly double-scope because the two frames clash, but it presents no problem for you. You do not say, "The caused-motion construction asks me to call up a scene of caused-motion, but the verb conflicts with that, so I am defeated." No, you have double-scope blending powers. You are not a chimpanzee or bonobo. You are a double-scoping cognitively modern human being, so these kinds of integrations are easy for you. You do it all the time. It is not cognitively costly. It is the way we think.

The story is similar for the Resultative clausal construction in English. An example is "*he boiled the pan dry*". Notice that *he* didn't boil. A stereotypical example would be "*Kathy painted the wall white*". That's a resultative — noun phrase, verb phrase, noun phrase, adjective. "*Catherine painted the wall white*" doesn't mean she painted the wall *because* it is white or *although* it is white or *because* she was white. No. But why not? The answer is that we know that *white* needs to be a result someone brought about by the agent. *White* needs to be a result for the *wall*. You know that from the construction, this little prompt.

Similarly, in "*he boiled the pan dry*", he did not boil and he did not boil the pan. You build the meaning according to which he performed an action, but you do not know for certain what the action was. Perhaps he turned a knob on the stove and then forgot about it. Perhaps he punched some buttons. You do not know what action he performed. *Boil* comes in from the event of the water's being heated. Notice that water is not even mentioned in the sentence. The Resultative construction in this case is taking in a verb from the event that is caused—namely, he performs some action that results in the boiling of the water: the water boils. *Boiling* is not an action the man is performing; it is an event caused by whatever actions he is performing. But we have

a compression in the blend. We can perform the same kind of compression for a transitive: *I boiled the water* or *I boiled it*. The agent is not boiling; he is performing an action that results in an event of boiling. But in the blend, there is a compression, and we can use the word “boil” for that compressed causation. In “*he boiled the pan dry*”, we construct the meaning according to which there is a result, *dry*, for the object, *pan*. *He* is the causal agent, and somehow boiling—which is not an action performed on the pan and not an event that occurs to the pan—leads to the result for the pan.

Other examples of the Resultative construction include:

She bled him dry.

She kissed him unconscious.

Last night’s meal made me sick.

He hammered it flat.

The earthquake shook the building apart.

Roman imperialism made Latin universal.

These are all cases in which grammatical constructions, form-meaning pairs, are able to exist because of double-scope conceptual integration, and integrating the forms prompts us to integrate the conceptual inputs. Because we are integrators, we can take the prompt and know that is asking us to put together a mental space network with conceptual integration, just as in *Ann is the boss of the daughter of Max* or *Paul is the father of Sally*. We can manage the diffuse conceptual network by working from the compressed, human-scale blend that can be prompted for by the integration of the linguistic forms.

The Ditransitive construction is similar. Consider “*Mary poured Bill some wine*”. The classic conceptual frame for the ditransitive involves *transfer*, as in “I hand you the controller.” May I have it back? Thank you. Notice that you can engage in these little transfers repeatedly with little children. They find it very entertaining. You ask the child if he or she would like the ball, and then hand it over. Then ask, “May I have it back?” It’s fun. A very basic human-scale scene. Goldberg, in studying construction, notes that a verb like *hand*—a denominal verb—is a basic verb of transfer.

Consider “*Mary poured Bill some wine*.” We take to call for a ditransitive frame because of the clausal syntax—noun phrase, verb phrase, noun phrase, noun phrase. It is very natural for you to project *reception* of the wine from the ditransitive frame of transfer, even though there is no mention that he received it, and I can ask you to cancel that projection. I can say “Mary poured Bill some wine and left it on the table, but he did not notice and didn’t pick it up.” Adele Goldberg has worked on the Ditransitive construction, including examples like “She refused him that courtesy”, which is interesting for the Ditransitive construction in the way that “We blocked him from the door” is interesting for the Caused-Motion construction. In both cases, the frame for the verb conflicts with the frame for the clausal construction. In “She refused him that courtesy,” the clausal construction calls for *transfer*, but the verb, *refuse*, indicates stoppage or denial of that transfer. Using double-scope integration, we can blend them in a compressed scene. Similarly for “She denied him the job”.

Let me give an example of what I call the elaborate ditransitive. Frequently, the

recipient in a ditransitive frame receives not just the object but also the benefit of receiving it. If I hand you a United States Dollar or 1000 RMB, you get not just the money, but also the benefit of the money. It's not that the ditransitive always provides a benefit. When my son hands me a banana peel, I don't receive a benefit. But often there is a reception of a benefit. So one of the standard uses of a ditransitive is to prompt for reception not of the patient of the action but rather the benefit that comes from that action. Goldberg's example is "Slay me a dragon" or "He slew me a dragon." In these examples, the princess does not receive the dragon, but she does receive the benefit of the dragon's having been slain. James Taylor sings a song in which he asks a member of the band to "Slide me a bass trombone." I pick this one because it is literally possible to slide someone a bass trombone so that they actually receive the bass trombone. You could slide it across the floor or across a table. But that is not the meaning we construct. What he means is that the trombonist should play the trombone—which involves moving the slide—so that James Taylor will receive the benefit of the trombone's having been played in this manner. *Slide* comes in as the event action that you do and what I get is not the bass trombone. What I get is the benefit of your having slid the bass trombone.

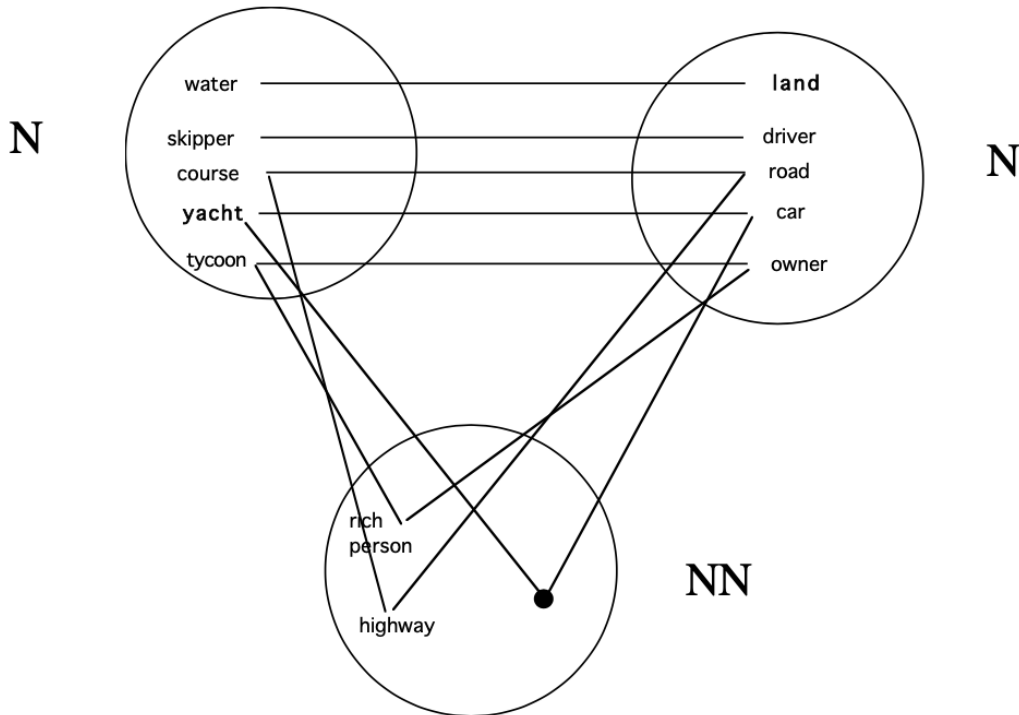
Nominal Compounds

We have talked about the way nominal compounds are prompts to do integration. We looked at examples like *boat house* and *house boat*. Typically, the two nouns in the nominal compound are not counterparts in the mapping that connects the inputs in the integration network. We talked about nominal compounds like *jail bait* and *land yacht*. Nominal compounding is an extremely common pattern. You take one noun from one space, one noun from the other space and that's all you need.

Let's look at *land yacht*. I pick this one because I think you are less likely to know how it is used in the United States. Most people in the United States do not know what a land yacht is, but if you ask them what the phrase might mean, they can sort of figure it out.

A land yacht is a car. It's a big car, not a specific big car—not a particular make or model, but a big car with certain features. It's a big car that wealthy people drive and it's not very good at handling. It sort of sways. There might be opera windows in the backseat and plush leather. It's not sexy at all. It's very corporate and important. It's a land yacht.

In the slide show, you see a diagram indicating an integration network for "land yacht".



Land yacht

Notice that in the integration network, the land and the yacht are not counterparts. The counterpart of *land* is *sea*, and the counterpart of *yacht* is *car*. But what you get down in the blend is something that you can refer to, something that is driven on roads that are nowhere near the water. The integration network is rich in cultural meaning; it is vaguely insulting. The person who owns the land yacht wants to look prestigious, but people who refer to the car as a “land yacht” are making fun of the owner.

Nominal compounds raise the subject of emergent syntax. One input space has a noun, “land,” associated with it. Another has a noun, “yacht,” associated with it. “Land yacht” is also a noun, or more exactly, a noun phrase. But, syntactically, this noun phrase can go where nouns go, for the most part. In creating a nominal compound, we did not create a new grammatical category. It’s still a noun phrase. The inside structure of *noun phrase* changes once a language develops nominal compounding, but the external syntax doesn’t change.

In this way, a new construction can arise that does not need to tear up the rest of language, because the emergent syntax can still sit inside existing syntax. We talked yesterday about derivative or blended syntax. We discussed the ways in which a verb like *risk*, because of blending, can take on alternative syntactic formation. You can *risk something to the waves*, in which case, you are calling in the exposure frame and the exposure syntax. You can *risk your money in the stock market*, in which case, you call in the investing frame and the investing syntax, namely *in*, because you invest *in* something. You expose *to*, but you invest *in*. You can say “*I risk 1000 RMB on the race horse*”, in which case you are calling in the betting frame because you are betting *on*.

So you get *risk to*, *risk in*, and *risk on*, emergent syntax from the blending of the frames, which brings with it blending of the forms.

The forms themselves can be combined into new syntactic forms, just as we saw in the case of *Y-of* constructions. The form of the *y-of construction* can be composed. We have the same kind of composition in the case of nominal compounding. Let's look at that. Consider the nominal compound *girl scout*. You probably have girl scouts in China because they seem to be everywhere. The Girl Scouts are an organization. Girl Scouts learn how to hike and camp and so on. We have Boy Scouts and we have Girl Scouts. Let's take *girl scout*. *Girl* is a noun, *scout* is a noun, *Girl Scout* is a noun, in the sense that *Girl Scout* can go into the spots in the language that nouns can go. Now take *ballet school*. *School* is a noun and *ballet* is a noun. That's a frame compatible network, because it's like *fan of bicycle racing*. People learn things in schools and ballet can be one of the things they learn, so the *ballet* frame is subordinated and nests in this case inside the *school* frame. Now take *lace curtain*. Those are two nouns and you get a nominal compound. Now, it may seem just funny and weird to say, "*Oh, she is a lace curtain ballet school girl scout.*" But in fact, you can understand it. Note how we have composed nouns and even nominal compounds, to get emergent syntax. What I might mean by "*Oh, she is a lace curtain ballet school girl scout*" is that she comes from a certain social demographic ("lace curtain," as in the phrase "lace curtain Irish") and she gets a certain kinds of education (ballet) and she is in a certain kind of organization. You might even recognize this as a social stereotype if you are in certain parts of the United States. Six nouns that go into three nominal compounds where the last of the nominal compounds *Girl Scout* now takes a role something like a head noun, and the other two drop into spots in the nominal compounding. They become modifiers.

Adjective + Noun

Consider adjectives, such as those in *guilty pleasures*, *likely candidate*, and *red ball*. Notice that a *likely candidate* is usually used to refer to somebody who is not a candidate. So you are not taking the meaning of *candidate* and the meaning of *likely* and putting them together. On the contrary, you are taking *candidate* from one mental space and *likely* from a particular kind of frame. "Likely candidate" can be taken as prompting you to construct a frame in which you have someone who is likely to be a candidate. This is not strange. This is the normal way to refer to someone who is likely to be a candidate. You don't say that there's *somebody who is likely to become candidate*. You say that *one likely candidate* is so-and-so. This phrase is not used to refer to an actual candidate who is likely to be elected, for example. Rather, it is used to refer to someone who is likely to be chosen to be a candidate. A likely candidate is not yet candidate.

Consider the phrase "guilty pleasures," as in "*Allow yourself these guilty pleasures*" and "*Chocolate is a guilty pleasure*". Of course, it is not the pleasure itself that is guilty. Rather, it is the person who has the pleasure who feels guilty. In this case,

there is a cause-effect vital relation between the input spaces—having the pleasure in one space causes the effect in the other. But now, that outer-space cause-effect relationship is compressed in the blend into a feature of the pleasure.

There are many similar examples, such as *grateful memories*. The memories are not grateful. The person who has the memories is grateful for the events to which the memories refer. But now that intentional relationship between you and your memories and the events to which they refer is compressed into a feature of the memories in the blend. Of course, you are not confused. You know exactly how to understand *grateful memories*. This is not an exotic or costly or unusual process. As we discuss in *The Way We Think*, it is the same process we use in stock, supposedly simple cases such as *red ball*.

Does anyone in the room know what a “silver honeymoon” is? No? Nobody? Good. You won’t believe it when I tell you. In the United States, there are vacation packages for a “silver honeymoon.” When you have been married for a long time,

OK, *likely candidate, silver honeymoon*. By the way, does anybody in the room know what a *silver honeymoon* is? No? Nobody? Good. You won’t believe it when I tell you. A “silver anniversary” is a twenty-fifth wedding anniversary. A “silver honeymoon” is a vacation for a couple that has been married for a long time, maybe even 25 years. There are tour agencies in the United States that will sell you a vacation package for a “silver honeymoon”, designed just for the specific couple. The company will arrange to take you back to just the kind of environment you were in when you were young adults. The old married couple dresses in clothes like the ones from their youth. They go to a dance with music from their memories. Perhaps they wear tennis shoes, go out for a cheeseburger and fries, and neck in a car. In one input space, we have the environment of their first honeymoon. In another, we have the old married couple with all of their resources for affording the trip, and the bodily conditions that go along with their actual age. The blend has an intimate, romantic vacation, which it gets from both input spaces. The “silver” comes from the present space while the “honeymoon” comes from the space of their first honeymoon. No one is fooled—the silver honeymoon is not actually the same as a first honeymoon, but then, it is not supposed to be. There is considerable emergent structure in the blend for considering, managing, guiding, and celebrating a life and a relationship.

Morphological integrations in a single word.

Consider the tunnel goes under what the English call “The English Channel,” connecting England and France. It is the Channel tunnel, which the English refer to as the “Chunnel.” It is an accident that the phonology and morphology of “Channel” and “tunnel” provide a good fit for a combination. Blending often exploits accidents. In French, naturally, the “English Channel” is not called the “English Channel”. It is called *La Manche*. So that’s *Tunnel sous La Manche*. So French does not present suitable phonological and morphological conditions for combining the forms. French does not make it easy to find a suitable blend of forms to indicate a blend of frames: *channel* and *tunnel*.

Blending in grammar often depends upon the structure that is already present in the language. If you ask, “Why don’t we do verbs in English the way the French do then in French?”, one answer will be that each language already has a system that is not easily disturbed. Blending is path-dependent, and whether the blend will be suitable depends on the path the language has gone through to get to its present state.

Consider “*McJobs*.” This phrase is used for a vast variety of employment, entry-level jobs with low pay, not much respect, and little chance of advancement. We can say, “Oh, the corporations in America are doing a terrible thing to the American workers, they are eliminating good jobs and replacing them with *McJobs*.” Using the compression of analogy and disanalogy to change for an element, we could even say, “They are *turning* good jobs into *McJobs*.” The form “McDonald’s” and the form “jobs” are integrated to prompt for an integration of the frames.

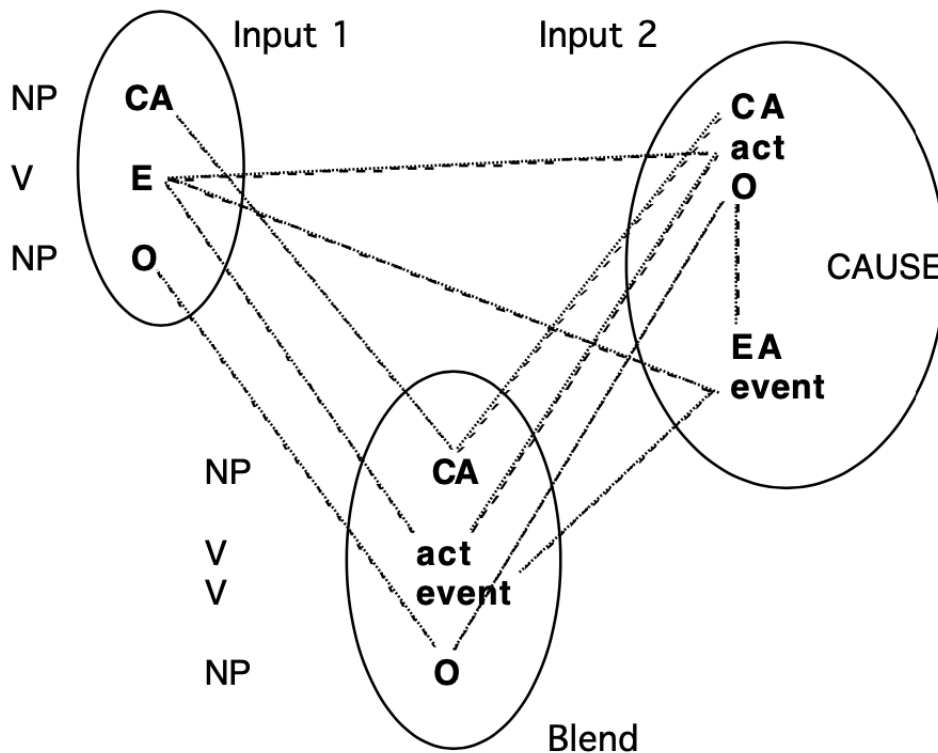
In all the cases we have considered in this talk, there are constructions in the language with stable grammatical patterns and those grammatical patterns prompt for the construction of conceptual integration networks. The indicated blending scheme—as for clausal constructions, words like “safe,” the *y-of* construction, and so on—carries with it particular kinds of compression. A more thorough analysis can be found in “Blending as a Central Process of Grammar”, by Fauconnier and Turner, which is available from the Cognitive Science Network.

Causatives in French and Hebrew

In French, sometimes you can use a single verb for an integration of events. It’s just a question of whether the single word has been created for the language. For example, *Pierre nourrit Paul* basically means “Pierre fait manger Paul”. Pierre has to do something, and Paul has to do something. We might characterize this situation as one in which Pierre feeds Paul.

But there is another construction in French, a double verb construction, a causative, in which you can use *faire* (which means something like “to do”) and a second verb, such *manger* (which means “to eat”). Then you can say *Pierre fait manger Paul*. Now the causation part of the complex of events belongs to Paul, who goes into the subject position. The agent of the caused event is Pierre. He is the agent of *manger*. There are two inputs in such cases: the causal action and the caused action. Each has a verb attached to it, and both verbs come down into the blend, where there is syntax for referring to the integrated event. The single-verb constructions available in French do not always allow enough room to include the needed prompts for construction. So French has created a new emergent syntax, a verb phrase with not one but two verbs, one verb taken from each of the inputs. We call *Pierre fait manger Paul* a transitive blend, because the basic clause that serves as the compressed input to the blend is a transitive construction: Noun-Phrase, Verb, Noun-Phrase:

Transitive Blend



Pierre fait manger Paul.

There are other double-verb causatives in French that arise from similar blending. The *transfer blend* arises when the basic single-verb construction that provides the compression to the blend is a transfer construction: NP V NP à NP, as in *Marie donne la soupe à Paul*. The *optional transfer blend* arises when the basic single-verb construction that provides the compression to the blend is the optional transfer construction: NP V NP (à NP) (par NP), as *Marie vend des livres (à Paul) (par un intermédiaire)*.

Suzanne Kemmer and Arie Verhagen in 1994 (“The grammar of causatives and the conceptual structure of events”. *Cognitive Linguistics* 5:2) wrote something quite brilliant: “Analytic causative constructions can best be described as extensions of simpler kinds of expressions, rather than as reductions from more complex underlying structures.” This is, I think, exactly right. Kemmer and Verhagen assert that there are cognitive models of causation based on force dynamics and interactions between participants, and these models relate to basic models, including transitive and ditransitive event structures, as we have seen. Gilles Fauconnier and I think this view is exactly right and that blending is the cognitive operation which allows the basic models to serve as inputs to the conceptual integration of more elaborate causal sequences and, in fact, to give us these kinds of blended syntactic structures. The result

in the case of the French double-verb causatives is that you get emergent syntax of verb plus verb. It's extremely complicated to account for the system, but the gist is, you have one verb from one space and one verb from the other; the first carries the causation and the second carries another event. The blend compresses the second input to the tight, basic, human-scale structure of the first input, but brings in the second verb from the second space. The single-verb constructions do not have enough slots for the prompts that you need to deliver, for the prompts needed to lead people to construct the right conceptual scene. But blending allows new, emergent syntax to provide a more suitable construction.

In summary, you have three basic single verb constructions in French: the transitive, the transfer, and the optional transfer. For each of them, there is a corresponding double-verb construction. All three of them have a causal agent who is attached to the first verb.

Building up from simpler structures is not a theoretical approach available to generative or relational theories of linguistics. So the exceptional amount of machinery those theories must import to try to solve or explain something like French double-verb causatives is immense. We think that the blending account is a much more straightforward and persuasive scientific generalization over the data. An additional strength of blending theory is that, on independent grounds, we must acknowledge the power and operation of double-scope blending, whether or not we take grammar and syntax into account at all. Blending theory does not need to invent ad hoc machinery. The mental operations that must be acknowledged independently provide the theory for the emergent grammar. We do not need additional machinery exclusive to just language in order to account for the emergence of complex syntactic patterns and simpler patterns.

There are also interesting where the prompt for causative blending is a single word. The relevant study here is Nili Mandelblit's work on binyan blends. Verbs in Hebrew are conjugated by blending a consonant rack with a set of vowels. The consonant rack prompts for a basic, compressed meaning. In this case, the causative can be formed as a single verb, because Hebrew offers this morphological form of blending. The concept prompted for by the consonant rack takes on other verbal meaning when those consonants are given a vowel-structure as dictated by the rules of the particular binyan. There are seven binyanim in Hebrew.

There is a morphological vowel-pattern in Hebrew that prompts for causation. When it combines with a consonant structure for a certain verbal concept, the result is a morphological blend that prompts for both the event and the causation. So a consonant rack for "run" and a causative vowel-structure creates a single verb meaning "cause to run." The result is that where French uses a two-verb solution for the causatives, Hebrew uses a one-word verb involving a morphological blend. So, in Hebrew, you do not need to bring in a separate verb for the causation, as in French *faire*.

Mandelblit explains that the consonant rack r-u-c means "run" and the vowel structure hi--i- means "cause" and their morphological blend *hiruic* (transliterated "heric") means "cause to run." Accordingly, you can have a sentence meaning *the commander made the soldiers run*, with the soldiers as the direct object, and a single

verb, *heric*, meaning “made run”.

Universal Grammar

After my lecture yesterday, I was asked about universal grammar. I refer you to this 2002 article by Hauser, Chomsky, and Fitch, “The Faculty of Language: What Is It, Who Has It, and How Did It Evolve?” They argue that human beings have many capacities for language and ask whether any of them belong to only language. Their assertion is that perhaps the only capacity that belongs exclusively to language is recursion; the rest of the capacities operate more broadly in cognition. Recursion, in their view, might be the only uniquely linguistic component of the faculty of language. But there is a suggestion that even recursion might not be exclusive to language: “We further argue that FLN [Faculty of Language, Narrow] may have evolved for reasons other than language, hence comparative studies might look for evidence of such computations outside of the domain of communication.”

This article caused a storm inside principles & parameters linguistics.

In cognitive linguistics or the precursors of cognitive linguistics, one finds a great deal of on syntactic combination, amalgams, and recursion. Of course, recursion is conceptual. We have seen that. I don’t understand why one would imagine that recursion, in the technical sense, belongs only to syntax or language. We have, at the conceptual level, *the mother of Paul, the mother of mother of Paul, the mother of the mother of the mother of Paul*. This is recursion. It is blending, and blending comes inherently with recursion. Similarly, at the level of meaning, we have *Aquinas believes Augustine believes Ambrose believes Paul believes Jesus believes . . .* This is recursive in the technical sense that the output of the operation is the input to the operation. It’s recursive. Consider a train of thought: *There is a particular film. Which film? The one that was panned by the reviewer. Which reviewer? The one who was kissed by the actress. Which actress? The one who was escorted by the director. Which director? The one who was insulted by the reviewer.* We do recursion in this form as an everyday operation, very commonly. The output of an operation can be an input to the operation. We have seen this conceptual recursion throughout our discussion of blending.

Of course, the recursion happens also at the level of form through blending. Now I am talking about the level of expressions: “the mother of the mother of the mother of Paul” and “Aquinas believes Augustine believes Ambrose believes Paul believes Jesus believes . . .” and “Lace stocking ballet school girl scout” and “The film [that was panned by the reviewer [who was kissed by the actress [who was escorted by the director [who was insulted by the reviewer]]]]. Of course, you get recursion in grammar. The reason you get recursion in grammar is because you get recursion in cognition, and the way you get it is by conceptual blending.

So is there universal grammar? Sure there is, in the sense that there are species-wide cognitive operations that make it possible for us to construct languages. I would never say that conceptual blending is a universal grammar, because it is the operation

of conceptual blending that we must emphasize, not particular products of conceptual blending. All human beings are equipped with this operation. It's not that particular blends belong to universal grammar, but rather that the operation of conceptual blending that belongs to universal grammar. And of course, blending is not the whole story. There are other cognitive operations, like attention. So what we would say is that the operation of conceptual blending is part of universal grammar. The operation of conceptual blending is part of, only part of, a non-uniquely-linguistic universal grammar.

Thank you very much.