# Metaphor

## A Computational Perspective

# Synthesis Lectures on Human Language Technologies

Automated Grammatical Error Detection for Language Learners, Second Edition
Claudia Leacock, Martin Chodorow, Michael Gamon, and Joel Tetreault
2014

Web Corpus Construction
Roland Schäfer and Felix Bildhauer
2013

Recognizing Textual Entailment: Models and Applications
Ido Dagan, Dan Roth, Mark Sammons, and Fabio Massimo Zanzotto
2013

Linguistic Fundamentals for Natural Language Processing: 100 Essentials from Morphology and Syntax
Emily M. Bender
2013

Semi-Supervised Learning and Domain Adaptation in Natural Language Processing
Anders Søgaard
2013

Semantic Relations Between Nominals
Vivi Nastase, Preslav Nakov, Diarmuid Ó Séaghdha, and Stan Szpakowicz
2013

Computational Modeling of Narrative
Inderjeet Mani
2012

Natural Language Processing for Historical Texts
Michael Piotrowski
2012

Sentiment Analysis and Opinion Mining
Bing Liu
2012

Discourse Processing
Manfred Stede
2011

Bitext Alignment
Jörg Tiedemann
2011

Linguistic Structure Prediction
Noah A. Smith
2011

Learning to Rank for Information Retrieval and Natural Language Processing
Hang Li
2011

Computational Modeling of Human Language Acquisition
Afra Alishahi
2010

Introduction to Arabic Natural Language Processing
Nizar Y. Habash
2010

Cross-Language Information Retrieval
Jian-Yun Nie
2010

Automated Grammatical Error Detection for Language Learners
Claudia Leacock, Martin Chodorow, Michael Gamon, and Joel Tetreault
2010

Data-Intensive Text Processing with MapReduce
Jimmy Lin and Chris Dyer
2010

Semantic Role Labeling
Martha Palmer, Daniel Gildea, and Nianwen Xue
2010

Spoken Dialogue Systems
Kristiina Jokinen and Michael McTear
2009

Introduction to Chinese Natural Language Processing
Kam-Fai Wong, Wenjie Li, Ruifeng Xu, and Zheng-sheng Zhang
2009

Introduction to Linguistic Annotation and Text Analytics
Graham Wilcock
2009

Dependency Parsing
Sandra Kübler, Ryan McDonald, and Joakim Nivre
2009

Statistical Language Models for Information Retrieval
ChengXiang Zhai
2008

Metaphor: A Computational Perspective

Tony Veale, Ekaterina Shutova, and Beata Beigman Klebanov

www.morganclaypool.com

ISBN: 9781627058506 paperback
ISBN: 9781627058513 ebook

DOI 10.2200/S00694ED1V01Y201601HLT031

A Publication in the Morgan & Claypool Publishers series
*SYNTHESIS LECTURES ON HUMAN LANGUAGE TECHNOLOGIES*

Lecture #31
Series Editor: Graeme Hirst, *University of Toronto*
Series ISSN
Print 1947-4040 Electronic 1947-4059

# Metaphor

## A Computational Perspective

Tony Veale
University College Dublin

Ekaterina Shutova
University of Cambridge

Beata Beigman Klebanov
Educational Testing Service

*SYNTHESIS LECTURES ON HUMAN LANGUAGE TECHNOLOGIES #31*

MORGAN & CLAYPOOL PUBLISHERS

## ABSTRACT

The literary imagination may take flight on the wings of metaphor, but hard-headed scientists are just as likely as doe-eyed poets to reach for a metaphor when the descriptive need arises. Metaphor is a pervasive aspect of every genre of text and every register of speech, and is as useful for describing the inner workings of a "black hole" (itself a metaphor) as it is the affairs of the human heart. The ubiquity of metaphor in natural language thus poses a significant challenge for Natural Language Processing (NLP) systems and their builders, who cannot afford to wait until the problems of literal language have been solved before turning their attention to figurative phenomena. This book offers a comprehensive approach to the computational treatment of metaphor and its figurative brethren—including simile, analogy, and conceptual blending—that does not shy away from their important cognitive and philosophical dimensions. Veale, Shutova, and Beigman Klebanov approach metaphor from multiple computational perspectives, providing coverage of both symbolic and statistical approaches to interpretation and paraphrase generation, while also considering key contributions from philosophy on what constitutes the "meaning" of a metaphor. This book also surveys available metaphor corpora and discusses protocols for metaphor annotation. Any reader with an interest in metaphor, from beginning researchers to seasoned scholars, will find this book to be an invaluable guide to what is a fascinating linguistic phenomenon.

## KEYWORDS

metaphor, simile, analogy, blending, figurative language processing

# Contents

# Preface

The aim of this book is to introduce metaphor research to the wider NLP community, and to survey the state-of-the-art in computational methods in a way that may also be helpful to those approaching metaphor from a perspective that is not principally informed by work in Artificial Intelligence (AI). We focus on the history, methods, and goals of past research into this fascinating phenomenon in the hope of making metaphor a more accessible topic of future research, thereby pushing it further up the NLP community's wait-list. Our treatment will provide a condensed history of metaphor research that introduces the main theories of metaphor that survive, in one form or another, in contemporary analysis. Our coverage will include the main AI contributions to the field, which are modern attempts to give algorithmic form to views on metaphor that range from the ancient to the contemporary. And, just as contemporary AI research has taken on a distinctly web-colored hue, we shall explore the role of the Web in metaphor research, both as a source of data and as a computational platform for our metaphor-capable NLP systems. Computational linguistics and AI alike have each embraced statistical models as a means of improving robustness, exploiting rich veins of user data, and reducing a system's dependence on hand-crafted knowledge and rules. Metaphor research offers no exception to this trend, and so our book will also explore the role of statistical approaches in the analysis of metaphorical language. Since such approaches are ultimately only as good as the data over which they operate, we shall also focus on the contributions of corpus linguistics to the construction of annotated metaphor corpora. Finally, we shall draw these strands together to offer an application-oriented view of metaphor, asking whether there is a killer application for metaphor research, and whether (and how) computational approaches to metaphor can help advance not only the field of NLP, but other fields as well, such as the social sciences and education.

The ultimate goal of this book is not to make you believe, as we do, that metaphor is the very soul of language, though the growing field of metaphor research is always eager to welcome new converts. We will consider this book a success if readers take away a desire to address metaphor head on, in some form or another in their research, and find in this book the necessary tools to make this engagement a practical reality.

Tony Veale, Ekaterina Shutova, and Beata Beigman Klebanov
January 2016

CHAPTER 1

# Introducing Metaphor

Language would be a dull and brittle thing without metaphor. It is metaphor and its figurative kin—simile, analogy, blending, irony, understatement, hyperbole, and the like—that lend language its vitality and elasticity. It is metaphor and its kin that allow us to suggest much more than we actually say, and to invent new ways of saying it, when conventional language shows us its limits. It is metaphor and its kin that allow us to communicate not just information, but also real feelings and complex attitudes. Metaphor does not just report the result of personal insights, but also prompts and inspires listeners to have these insights for themselves. Each metaphor is a concise but highly productive *way of saying* that communicates a new and productive *way of seeing*.

But what exactly are metaphors and where do we look for them? Although metaphors are products of our faculty for creative thinking, metaphors in the wild can range from the scintillating to the banal. Just as repeated usage dulls the blade of a trusty knife, or repeated telling robs a once-funny joke of its ability to raise a laugh, repetition takes the bloom off a metaphor and turns it from an eye-catching flower into just another piece of the undergrowth. So metaphors are everywhere, in language, in film, in music, or in any system of signs that allows us to express ourselves creatively. These metaphors range from the novel to the conventional, and indeed some are so conventional as to escape our attention altogether. Nonetheless, even highly conventionalized metaphors retain a spark of the creativity that forged them, and in the right hands this spark can be fanned into a roaring fire. Consider the following example, which was used as the title of a popular science book by German rocket scientist Wernher von Braun: *"I Aim For The Stars!"* Since the heavens are filled with stars in any direction we care to look, von Braun could not have intended to use the word "aim" literally, which is to say *"I look and/or point in the specific direction of the stars"* or even (for he was a master pragmatist) *"I intend to move myself physically closer to the stars."* Rather, von Braun—who was attempting to popularize the concept of space travel with the American public, and thereby win state funding for his expensive scientific work—employed a basic metaphor that allows speakers to treat Purposes as Destinations. Thus, von Braun saw the achievement of interstellar space travel as his goal, and expressed this goal using the language of physical destinations.

But von Braun, an ex-member of the Nazi party, was also the controversial creator of the German V-2 missiles that rained down on Britain during World War II. (He had been spirited away to the U.S. as part of *operation paperclip* at the end of the war, when the American and Russian militaries competed to round up the brains of the German rocket program.) The comedian

Mort Sahl took full advantage of this fact, and of von Braun's use of the PURPOSES ARE DESTINATIONS metaphor, to cheekily propose an alternative title for the scientist's book: *"I aim for the stars but I keep hitting London."* Clearly, the generic and highly reusable metaphor PURPOSES ARE DESTINATIONS is of a different character than the specific utterances that are constructed from it. Researchers refer to the former as a *conceptual metaphor*, for it resides at the level of thinking and ideas, and to the latter as a linguistic metaphor. Explicating the relationship between the former and the latter is one of the goals of contemporary cognitive and computational approaches to metaphor, and so we shall return to this relationship many times in this book. We'll also meet the conceptual metaphor PURPOSES ARE DESTINATIONS again in a later chapter.

We each use conventional metaphors every day, perhaps without even realizing it, yet we each have the ability to elaborate on these standard-issue constructions in our own way, to inject our own voice and personality into what we say. Consider the following quote from Twitter co-founder Biz Stone, who writes in his autobiography of how he came to be an early employee of Google:

> *I didn't know it at the time, but behind the scenes Evan [Williams] had to pull strings in order to hire me [at Google]* (Stone [2014]).

The idea of exerting influence on others by *pulling [their] strings* is deeply entrenched in the English language. We describe a bargain or an offer of special treatment as *"no strings attached"* when we believe the giver is not seeking to unduly influence us, which is to say *"to pull our strings."* An emotional appeal that hits its mark is said to *"tug at our heartstrings,"* while we might describe a master manipulator as being able to *"play someone like a violin."* A mother's continued influence on an adult child is often given metaphoric form with the phrase *"apron strings,"* and any effort (by mother or child) to curtail this influence is described as *"cutting the apron strings."* Stone's use of the *strings* metaphor in the context of the idiom *"behind the scenes"* might also bring to mind images of the pulleys and ropes with which stagehands lower and raise the curtain in a theatrical production (indeed, the word *"scene"* gets its meaning from the piece of cloth that was draped behind the stage in ancient theatres). Metaphors are much more than the stuff of fancy wordplay, and we use them to do much more than give our messages an attractive sheen: metaphors engage fully with the mechanisms of thought, allowing us to spark associations, insights, and other metaphors in the mind of a hearer even when we are using the most conventionalized of figures. We can think of these figures as being made of clay; convention has given them their shape, but we can add fine detail of our own, or bend them further to our own meanings. Let's look at the larger context of Stone's metaphor:

> *I didn't know it at the time, but behind the scenes Evan had to pull strings in order to hire me. Actually, they were more like ropes. Or cables—the kind that hold up suspension bridges* (Stone [2014]).

This elaboration should dispel any doubts about whether the *pulling [one's] strings* metaphor is nothing more than an arbitrary idiom that speakers learn to repeat in whatever context that suits.

Rather, a metaphor—even a highly conventional metaphor—establishes a frame of thought that encourages us to think in a particular way. Once it has our attention and draws us in, we are free to explore it, question it, and customize it as we see fit. For instance, we might ask how pulling on metaphorical strings might influence another person. If this other person is a weak-willed *puppet*, a *lightweight* player, or a minor *cog* in the machine (notice how effortlessly one metaphor leads to others), then not much effort is needed to exert influence, and so its strings will be light and easy to pull. But if this other person is a significant cog—a so-called *big wheel* or a *heavyweight* player—greater effort is needed to achieve any influence, requiring strings of greater thickness and tensile strength. Notice how the metaphor encourages us to think in the *source* domain (the domain of strings, cogs, pulleys, puppets, etc.) and to transfer our insights from here into the *target* domain (the domain of corporate decision-making). In Stone's example, his metaphor leads us to believe that his friend Evan needed to perform Herculean efforts on his behalf, to influence some very powerful people at Google by pulling on some very heavy-duty strings. As we'll see repeatedly throughout this book, even the most innocuous metaphors conceal a wealth of complexity, both in terms of their underlying representations and the cognitive/computational processes that are needed to understand them. This hidden complexity is a large part of what gives metaphors their allure for the computationalist.

However, although metaphor has a long and illustrious history of academic study, in both philosophy and linguistics, it remains a niche area in the computational study of language. For although most Natural Language Processing (NLP) researchers would readily acknowledge the ubiquity of metaphor in language, metaphor is a complex phenomenon and a hard engineering problem that continues, for the most part, to be wait-listed by the NLP community. To the application-minded, there are simply too many other problems of more practical and immediate interest—concerning syntax, semantics, inference, sentiment, co-reference, under-specification, and so on—that jump to the top of the community's collective to-do list. That metaphor touches on all of these problems and more is often seen as beside the point, although ultimately it is very much to the point: the problem of metaphor is just too big, too unwieldy, and too knowledge-hungry to be tackled first. Better to get a handle on all the other problems first, to obtain an algorithmic understanding of the workings of language that can later be enriched by a computational model of metaphor. This makes good engineering sense, if little philosophical or cognitive sense.

The purpose of this book is to demonstrate that this conventional wisdom is predicated upon a false dichotomy: researchers can build figurative-language processing systems that are practical and efficient *and* cognitively plausible, and which also reflect an understanding of the profound philosophical issues involved. Indeed, it is difficult for the computationally-minded researcher to explore aspects of metaphor that have *not* been previously visited by philosophers or psychologists or by earlier computationalists in its long and illustrious history of academic analysis. If there is little in the field of metaphor that one can truly call "virgin territory," it is nonetheless a field of many interesting landmarks that rewards careful viewing and repeat visits. We have written

this book to be a comprehensive guide to the major landmarks in the computational treatment of metaphor and hope the reader will find it a useful map to this fascinating phenomenon's many attractions.

CHAPTER 2

# Computational Approaches to Metaphor: Theoretical Foundations

## 2.1 THE *WHAT*, *WHY* AND *HOW* OF METAPHOR

Metaphor is both pervasive and evasive: ubiquitous in language, yet remarkably hard to pin down in formal terms. Yet the fact that there exists no single, definitive perspective on metaphor is very much in keeping with its chimerical nature. For metaphor is a highly productive mechanism that allows us to create a panoply of viewpoints on any concept we care to consider, including metaphor itself. Indeed, it scarcely seems possible to say anything meaningful at all about metaphor without first resorting to one kind of metaphor or another. So we talk of tenors and vehicles, sources and targets, spaces and domains, mappings and projections, metaphors that are living or dead, those that are fresh but soon stale, or of seeing one idea with the aid of another. Many of our metaphors for metaphor are metaphors of seeing, for seeing often leads to knowing, or at least to an impression of knowing. So a metaphor can be a viewfinder, a magnifying glass, a microscope, a lens, a window, a pair of conceptual spectacles, a distorting fun-house mirror, and even a set of conceptual blinders. Max Black, the philosopher who more than anyone else kick-started the modern resurgence in metaphor research, described a metaphor as working much like a blackened piece of glass onto which a specific pattern of lines has been etched [Black, 1962]. Looking through this bespoke glass, we see only those parts of a target domain that the etched pattern allows us to see, and see these parts in clearer relation to each other precisely because we are not distracted by the many parts we cannot see. Metaphor prompts us to look, helps us to see, and then controls what we can see.

Metaphor has existed longer than our need to describe its function, and much longer than our need to name it as a distinct linguistic or cognitive phenomenon. Indeed, the word "metaphor" is itself founded on a metaphor, to carry (*-phor*) above and across (*meta-*) a signifier from one realm of experience to another. The Greek philosopher Aristotle, to whom we can trace the earliest use of the term, effectively describes metaphor as a form of semiotic displacement, in which our agreed signifier for one object, idea, or experience is deliberately displaced onto another. But metaphor is not the only kind of semiotic displacement in language. Punning, for instance, involves the temporary replacement of one signifier with another, based on the phonetic similarity of the

two, as in "mail" and "male" (e.g., "Q: Is that the *mail* plane coming in to land? A: Of course it's a *male* plane, can't you see its undercarriage?"). Synecdoche licenses the displacement of a signifier of a part onto the whole to which it belongs (e.g., "brains" stands for clever (or *brainy*) people in "the American and Russian militaries competed to round up the *brains* of the German rocket program"), while metonymy licenses a more general displacement of a signifier onto an object or idea that is functionally or conceptually related to its conventional signification (e.g., when Raymond Chandler's private investigator Philip Marlowe says "Trouble is my business," he means "My business concerns other people's troubles, which invariably cause trouble for me;" so the metonymy tightens the relation between Marlowe and trouble). Aristotle is thus careful to enumerate the particular kinds of displacement that constitute a metaphor, as opposed to any other kind of semiotic play, and in his *Poetics* he offers the following schematic specification (as translated by Hutton [1982]):

> *Metaphor is the application to one thing of the name belonging to another. We may apply:*
>
> (a) *the name of a genus to one of its species, or*
>
> (b) *the name of one species to its genus, or*
>
> (c) *the name of one species to another of the same genus, or*
>
> (d) *the transfer may be based on a proportion.*

The target of the name shift is typically called the *topic*, the *tenor*, or simply the *target* of the metaphor, while the source of the shift—the name belonging to another—is typically called the *vehicle* or simply the *source* of the metaphor. The designations *source* and *target* are generally preferred in the analogical literature (e.g., Gentner [1983]), although *vehicle* is more in keeping with Aristotle's view of metaphor as a carrier of meaning across domains. Throughout this book we shall use the terminology of *source domain* and *target domain* to denote the space of ideas associated with the source (or vehicle) and the target (or tenor) respectively. The word *domain* has no formal definition; we use it, as others do, to designate the cluster of related concepts, properties, and norms that attach to a particular source or target so that we may later speak of transferring, mapping, or projecting content from one domain onto another. With this in mind, Aristotle offers an illustrative use-case for each of his four displacement strategies.

(a) Genus to species: "*Here stands my ship*" employs the generic (genus) term "*to stand still*" as a replacement for its specialization (species) "*to be at anchor*."

(b) Species to genus: "*Truly ten thousand noble deeds hath Odysseus done*" uses "*ten thousand deeds*" as a specialization of the vague genus term "*large number*."

(c) Species to species: "*Drawing off the life with bronze*" and "*Cutting off the water with unwearied bronze*" each employ a specialization of the genus "*to take away*."

(d) Proportional analogy: "*The wine cup is to Dionysus as the shield is to Ares*" is an obvious statement of the proportion *Dionysus:Cup::Ares:Shield*, yet it supports more subtle allusions, such as the poetic use of "*the cup of Ares*" to refer to a battle shield, or "*the shield of Dionysus*" to refer to a wine cup.

So, metaphor distinguishes itself from other forms of semiotic displacement by its reliance on the vertical operations of generalization and specialization. Metaphor allows us to move signifiers up and down a hierarchy of concepts, and supports the sideways movement of signifiers only insofar as they achieve such movement by climbing up from one domain so as to climb down into another. In effect, Aristotle provides a strong account of the "*what*" of metaphor, and a reasonable account of the "*how*," but offers very little on the "*why*" of metaphor, other than to suggest that it is done in a spirit of semiotic sport. The strategies in (a) to (d) above focus mainly on those uses of metaphor where the vehicle is explicitly given (e.g., the *cup of Ares* or *ten thousand deeds*) and a listener must work out the identity of the implied tenor (e.g., a *shield*, or a *large number*). But what of metaphors where tenor/target and vehicle/source are both explicitly given, as in the copula metaphor "*marriage is slavery*?"

Such metaphors appear to instantiate strategy (c) in Aristotle's scheme, in which one species of an unstated genus is mapped to another of the same genus. Presumably it is the identification of this unstated genus term—perhaps *cruel practice* or *social institution* in the case of *marriage as slavery*—that offers the key to the metaphor. Yet, if so, the metaphor degenerates into an example of strategy (b) and communicates only as much information as this genus can offer. In fact, if the listener does not possess the same conceptual hierarchy as the speaker—for instance, the listener may not view even slavery as a cruel practice, much less marriage—then the metaphor will fail. However, one need not agree with a metaphor to understand its purpose, and many metaphors are designed to change a listener's conceptual structures rather than to merely reflect them. So, in an important sense, the "*why*" of metaphor is precisely this: speakers often use metaphors to align the conceptual systems of others to that of their own.

For a metaphor is not a statement of fact, and most are more than the mere expression of a subjective viewpoint. Rather, a metaphor is an invitation to build a joint meaning space, as framed by the given viewpoint. The metaphor "*marriage is slavery*" is a communicative gambit that asks the listener to think about, and perhaps argue about, the idea that our concept of marriage can be meaningfully framed by our concept of slavery. In other words, the metaphor aims to establish what Brennan and Clark [1996] call a *conceptual pact*, and is a request that can be paraphrased as follows: "*let us agree to talk about marriage using the words and ideas we normally associate with slavery; to this end, let us go so far as to suspend disbelief and assume that marriage really is a kind of slavery, and thereby explore the extent to which our concept of marriage can be subsumed under our concept of slavery*." The longwindedness of this paraphrase goes some way toward conveying the true utility of metaphor as much more than a flourish of language or a playful exercise in semiotic sport.

The metaphor "*marriage is slavery*" is an effective prompt to build a complex meaning space because "*slavery*" is itself an effective shorthand for a rich body of cultural knowledge and expectations. Wrapped up in our concept of slavery is a knowledge of the history of the practice, of the physical and emotional suffering of those who are caught in it, and of the motivations of those who practice it. In other words, slavery is a *dense descriptor*, an idea that is easily referenced and that brings with it a slew of properties, feelings, norms, and expectations. The use of a dense descriptor in a metaphor prompts a listener to unpack those aspects of the descriptor that seem relevant to the target/tenor. This conciseness comes at a price, as the listener may associate different feelings and properties with a concept like *slavery*, or may unpack the descriptor to focus on different aspects than those of interest to the speaker. For instance, a speaker may employ the concept of slavery to focus on the domestic abuse, emotional and physical, suffered by one who is trapped in a bad marriage, while the listener may instead unpack the reference to focus on the institutionalized nature of marriage and slavery. Yet the indeterminate nature of metaphor, a natural consequence of its concision, is often a desirable quality of a conceptual pact. One rarely uses metaphor to communicate a rigid set of facts or beliefs. Rather, metaphors introduce a set of *talking points* for a target, so that speaker and listener might converge toward a mutual understanding, if not of a target idea then of each other's views of the target. The indeterminacy of metaphor makes it a flexible form of communication. While literal language commits a speaker to a fixed interpretation, and offers little scope for the listener to contribute to the joint construction of meaning, metaphorical language suggests a looser but potentially richer meaning that is amenable to collaborative elaboration by each participant in a conversation.

The densest descriptors are the familiar stereotypes that get used time and again in our metaphors and our similes. Although every concept has the potential to be used figuratively, casual metaphors tend to draw their dense descriptors from the large pool of stereotypes shared by most speakers of a language (see Taylor [1954]). Because so many familiar stereotypes have polarizing qualities—think of the endearing and not-so-endearing qualities of babies, for instance—they serve as ideal vehicles for a metaphor that aims to convey a strong affective stance toward a topic. Even when stereotypes are not used figuratively, as in the assertion "*Steve Jobs was a great leader*," they are often likely to elicit metaphors in response, such as "*yes, a pioneer*," or "*a true artist!*" or even "*but what a tyrant!*" Familiar proper-named entities can also be used as dense descriptors, as when Steve Jobs is compared to the fictional inventor Tony Stark, or Apple is compared to Scientology, or Google to Microsoft. Metaphors are flexible conceits that draw on dense descriptors to concisely suggest a position on a target idea while seeking elaboration or refutation of this position from others. Our computational models for the interpretation of metaphors must thus allow speakers to exploit the same flexibility of expression when interacting with machines as they enjoy with other humans. Such a goal clearly requires a computational system to possess, or acquire, a great deal of knowledge, usefully clustered into dense pockets of properties, feelings, and expectations that can be evoked with a single word or idea.

## 2.2 THE "MEANING" OF METAPHOR

If metaphor is as much a process as a product, what then is the meaning of any given metaphor? We can point to the linguistic rendering of a metaphor and say that this string is a linguistic metaphor. We can point to the conceptual structure underpinning this linguistic rendering—and to other renderings of the same conceit—and say that this deep structure is a conceptual metaphor. But to what can we point and say: this is the *meaning* of this metaphor? The problem of metaphor meaning can be approached from a number of perspectives. We can talk of the truth value of a metaphor, much as we can talk of the truth value of any proposition. We can talk of a specific interpretation of a metaphor. Or, we can talk of the specific inferences that a metaphor licenses in the mind of the hearer, as well as the emotions for the target that it is likely to engender in a hearer.

Although most metaphors are literally false, truth value is a remarkably poor guide to the metaphorical nature of any statement. The assertion that "*Barack Obama is a Muslim*" is no more a metaphor because it is false than "*Al Gore is not a robot*" is not a metaphor because it is true. For logicians and philosophers, meaning resides not in truth values but in *truth conditions*, the set of criteria that, if true, would make a statement true as a whole. For instance, in semantic theory in which "*Barack Obama*" denotes the President of the United States in 2015, and in which "*Muslim*" denotes the set of all people who profess faith in the religion of Islam, the key truth condition for "*Barack Obama is a Muslim*" is whether the former individual is a member of the latter set. Different speakers may bring different definitions and denotations to bear on an utterance, and so it is sufficient for some political critics to define the truth condition of "*being a Muslim*" as membership in the set of people born of at least one Muslim parent, or as membership in the set of people born or raised in predominantly Muslim countries. However one finesses the denotations, "*Barack Obama is a Muslim*" has different truth conditions than "*Barack Obama is a socialist*," as each hinges on membership in different (albeit possibly overlapping) sets, regardless of whether both statements ultimately have the same truth value.

As if defining the literal truth conditions of a predication such as "*being a Muslim*" or "*being a socialist*" weren't hard enough, metaphor poses a further, rather special challenge to the truth-conditional view of meaning. What, for instance, are the truth conditions of the statement "*Barack Obama is the Cicero of the 21st Century*?" One often needs as much creativity to assign truth conditions to a metaphor as to invent one in the first place. For truth conditions are neither as authoritative or as absolute as they seem, nor are they designed by experts to be shared by all. Rather, truth conditions are often speaker-relative and context-sensitive. One speaker may think that to truthfully call one a socialist, it is enough that this person openly espouse socialist values. Another may think it sufficient for one to act like a socialist, whether one actually thinks of oneself as a socialist or not. These speakers are likely to disagree over the truth conditions of the statement "*Barack Obama is a socialist*," yet this is not a problem for metaphor or for language. The truth conditions of a metaphor often arise from a tacit negotiation between speaker and hearer of the meaning of the metaphor, rather than the other way around.

Philosophers who are otherwise bullish about truth conditional semantics are often bearish toward the idea in the context of metaphor. Donald Davidson, for instance, questions the usefulness of truth conditions for arriving at the meaning of a metaphor. For Davidson [1978], a metaphor does not so much communicate a meaning as inspire a meaning, leading him to argue that "*the attempt to give literal expression to the content of the metaphor is simply misguided*." Thus, in the metaphor "*my car drinks gasoline*," only some of the truth conditions associated with literal uses of the verb *to drink* will be applicable in the metaphorical context of a drinking car, but none of these carry the weight of the interpretation that the speaker hopes to inspire in the listener (for instance, that my car requires *too much* gasoline to operate). These truth conditions include, for instance, the expectation that the drinker is an animate creature, that the substance consumed is a potable liquid, and that the drinker both consumes the liquid and derives some chemical benefit from it. Just two of these truth conditions are sensible for the metaphor "*my car drinks gasoline*," namely that my car consumes gasoline and derives some chemical value from it. In effect, this is the truth-conditional equivalent of Aristotle's strategy (b), whereby the species *to drink* is displaced onto its genus *to consume*. However, to suggest that this banal observation is the speaker's intended meaning is to fly in the face of Grice's *maxim of informativeness* (see Grice [1978]), since the claim of gasoline consumption holds for most cars and is hardly worth mentioning in the context of this car. Rather, as we can expect most listeners to know that cars consume gasoline, that people generally only drink what they want to drink (and thus drink what they like), then the speaker's most likely meaning is *my car likes to drink gasoline and thus consumes a lot of gasoline, rather more than I would prefer*. This meaning cannot be derived from the truth-conditional semantics of the words in the metaphor, but can only emerge from a pragmatic understanding of the world knowledge behind the words.

Black [1962] proposes an *interactionist* view of metaphor that explains how the meaning of a metaphor emerges from a heuristic, common-sense consideration of possible interactions between the ideas raised by the metaphor. For Black, the meaning of a metaphor is more than the sum of the meaning of its parts, and a surprising meaning may thus emerge from the juxtaposition of two rather familiar ideas (in this sense, Black foreshadows the notion of emergent meaning developed by Fauconnier and Turner [2002] in the context of their theory of *conceptual blending*). The use of common-sense knowledge—such as cultural associations, folk beliefs, conventional implicatures, and so on—offers a much richer and open-ended basis for exploring the associations and beliefs that emerge in the mind of the listener in response to a metaphor than a purely truth-conditional approach based on denotations and set membership. As in the more contemporary take on Black's interaction view of metaphor offered by Indurkhya [1992], metaphor is not so much a user of pre-existing similarity as a creator of new similarities, insofar as the meaning that emerges from a metaphorical juxtaposition reveals the source and target to be more similar than the listener may have previously realized. Davidson and Black (and later researchers such as Indurkhya) agree that what makes metaphor special is not the truth conditions imposed by its words on the semantics of the specific utterance, but the system of ideas and associations that are

evoked—in response to the metaphor—in the mind of a "suitable" listener in a "suitable" context. A key question for Davidson is whether this system of ideas and associations can be considered the *secondary* or *special meaning* of the metaphor (where the literal interpretation serves as the primary meaning), or indeed whether it is well-founded to consider a metaphor as having such a secondary meaning. As Davidson [1978] points out, similes do their work without the need for a secondary meaning—they typically mean what they purport to mean on the surface—so why should metaphors be any different?

Specifically, Black and Davidson disagree on the extent to which it is sensible to speak of the special meaning of a metaphor (as opposed to the literal meaning that resides on the surface) as a parcel of cognitive content that is communicated from speaker to listener. Davidson [1978] explicitly denies the position (which he attributes to Black) that a metaphor "asserts or implies certain complex things by dint of a special meaning and *thus* accomplishes its job of yielding an insight" (his italics). For Davidson, to speak of the meaning of a metaphor as a distinct message or insight is simply wrong-headed, as misguided as speaking about the definitive meaning of a poem, a joke, or a theatrical play. For instance, what does it mean to speak of the meaning of a play? At one level, the only level at which Davidson argues that it is sensible to speak of the meaning of a metaphor, a play means just what its words and sentences purport to mean, although it may also nudge a viewer to think certain thoughts or consider certain possibilities about the world. The deeper meaning of a play, what one might call its message to the audience, is not contained in the text itself, but is inspired in the mind of a viewer. Davidson's view is reminiscent of the angry response by the Irish playwright Brendan Behan to an interviewer who had asked him to summarize the message of his most recent work: "*Message? What the hell do you think I am, a postman?*" However, the fact that two viewers may find themselves in agreement in their responses to a play, and find that their responses agree with those intended by the playwright, suggests that playwrights may nonetheless be successfully asserting some secondary meaning via their work. Likewise, the fact that listeners may agree as to the insights delivered by a metaphor, and agree with the speaker as to the nature of these insights, speaks to the reality that speakers often do aim to assert something with their metaphors that is not to be found in the literal meaning of their statements. Indeed, as Black [1979] notes, the fact that listeners may disagree with each other, or with a speaker, in their response to a metaphor suggests that speakers do wish to assert some measure of propositional content with their metaphors. The difficulty of precisely pinpointing and circumscribing this propositional content (or the impossibility of doing this *in principle* for all metaphors) should not lead us to infer that it never exists, or give up on the task of trying to pinpoint (formally or computationally) at least part of this meaning.

## 2.3 PARAPHRASING METAPHOR

Most discussions about the meaning of a given metaphor inevitably turn on our ability to produce a convincing paraphrase. Yet Davidson [1978] raises a cautionary note when he argues that "what we attempt in 'paraphrasing' a metaphor cannot be to give its meaning, as that lies on the

surface; rather, we attempt to evoke what the metaphor brings to our attention." So just what is a paraphrase then? If we take what "lies on the surface" of a metaphor to be its literal content in the source domain, and what it "brings to our attention" to be its literal focal point in the target domain, then what we mean by a literal paraphrase is a *target-domain* paraphrase. This distinction, although perhaps obvious, is an important one. After all, many metaphors—such as negative metaphors—are literally true. In the case of negative metaphors such as "*I am not Superman*," "*Money does not grow on trees*," and "*Your father is not an ATM*," the surface meaning from the source-domain of each is literally true and just as informative as the corresponding metaphor, but it cannot usefully be taken as a literal paraphrase of the metaphor. Likewise, a paraphrase that merely transforms a metaphor into a simile—such as "*it's like my car drinks gasoline*"—may yield a reformulation that is literally true, but it does nothing to shift the focus of our attention to the relevant area of the target-domain. (The relation of metaphor to simile, and the ability to expand/compress/paraphrase one as the other, is the subject of ongoing debate; see e.g., Glucksberg and Haught [2006]; Utsumi [2007]; and Barnden [2012] for a flavor of this debate.) In other words, it is not the literalness of a paraphrase that makes it a useful guide to a metaphor, but its literalness with respect to the target domain.

Because concise metaphors are often the most under-specified, they frequently offer the most open-ended interpretations and require the most complex target-domain paraphrases. Indeed, this richness of possible paraphrases may not be an accidental feature of metaphor but a driving force in its interpretation. Utsumi [2007, 2011] refers to this richness as the *interpretative diversity* of a metaphor. Utsumi argues that listeners often choose an interpretation strategy (such as, e.g., a comparison of source to target, or a categorization of the target under a category suggested by the source) that increases the diversity of the resulting interpretation. Consider the metaphor "*wishes are children*" from the musical *Into The Woods*. This metaphor, coupled with the narrative context in which it is used—a fairy-tale world in which wishes come true with unexpected consequences—can evoke a diversity of more-or-less mutually consistent thoughts about wishes and their fulfillment, which might be paraphrased thus: "*wishes are born of our deepest desires*;" "*wishes have lives of their own and grow to follow their own paths*;" "*wishes can bring us pain as well as joy*;" "*we give life to our wishes and so are responsible for their consequences*;" etc. Related figurative phenomena, such as metonymy, may also play a role in paraphrasing a metaphor. For instance, one might consider wishes to be the stuff of childish optimism, and thus consider wish-makers to be child-like in their approach to life. Childless couples often wish for children, and so their wishes—when they come true, as in the case of one sub-plot of *Into The Woods*—actually become children. However, for Davidson at least, the metaphor means just what it says on the surface: that is, in truth-conditional terms, for every wish $w$ there exists a child $c$ such that $w = c$. As such, none of those prior statements is a paraphrase of the metaphor, rather they are a prior justification for it or a posterior response to it.

Davidson's injunction against "*attempt[s] to give literal expression to the content of [a] metaphor*" further reminds us that metaphor is never reducible to a literal restatement, no matter how

detailed that restatement may be. This fact is tautological in its obviousness, for, by definition, any literal restatement is *not* a metaphor, does not resonate with the same semantic tension as a metaphor, does not tease and thrill with the same ambiguity and under-specification as a metaphor, and does not pose a comparable challenge to our conceptual systems. Moreover, any literal restatement is not a paraphrase of the actual meaning of the metaphor, but a paraphrase of just one interpretation. The meaning of a metaphor resides not in a single authoritative interpretation, but in a whole space of possible interpretations. The indeterminacy and interpretative diversity of metaphor lends it an elasticity that cannot be captured in literal language. No matter how detailed our paraphrase, a metaphor always holds out the promise of more—more detail, more mappings, more associations—if we would only deepen our search. Finally, a paraphrase does not propose the same conceptual pact to an audience, and so is unable to serve the same communicative role as the metaphor it replaces.

One important communicative role that is not captured by an explicit paraphrase is *plausible deniability*, wherein a speaker uses metaphor to communicate a tendentious claim, and subsequently appeals to the ambiguity of metaphor to deflect criticism if this claim is ill-received by an audience. Consider the case of Republican presidential candidate Donald Trump, who was challenged on his alleged record of sexism by FOX-News reporter Megyn Kelly during the first televised GOP debate in 2015. Trump rebutted the charge and complained of mistreatment at the hands of the debate moderators, later claiming in an interview that "*you could see there was blood coming out of her* [Kelly's] *eyes, blood coming out of her wherever.*" Many commentators in the media took Trump's comments to be a potent mixture of metaphor, metonymy, hyperbole, and bombast. Although Trump painted a picture of a woman possessed by a demon, bleeding uncontrollably from every orifice, and thus lacking control over her own physical and mental state, his under-specified reference to a woman bleeding from "*her wherever*" was widely taken to be a metonymic reference to menstruation. Trump, it was claimed, was simply giving new form to the old charge that a woman who behaves aggressively to a man must be in the irrational grip of premenstrual tension. In a daring move, Trump was seemingly fighting charges of sexism with more sexism, yet when pressed on the matter, he defended himself with this tweet: "*Re Megyn Kelly quote: 'you could see there was blood coming out of her eyes, blood coming out of her wherever' (NOSE). Just got on w/thought.*" He elaborated further in a subsequent statement, claiming that "*Only a deviant would think anything else.*" Although we can debate Trump's actual intentions for his blood metaphor, it does seem that its non-literal under-specification was an essential part of its communicative function, allowing Trump to push a provocative idea into the public sphere while denying that this idea was ever part of his intended message. Another issue with literal paraphrases then is that they sometimes render in plain speech an opinion that was never intended to be plainly expressed or openly interrogated in the first place.

Yet, for all this, there is little to be gained from being a radical skeptic on the matter of literal restatement. Paraphrases may be necessarily imperfect, but they possess a practical utility for researchers and speakers alike. For, in establishing a conceptual pact, a metaphorical utterance

may prompt a listener to produce a paraphrase in response, as the listener's own contribution to the construction of a joint meaning space. By doing so, the listener may propose a particular reading of a metaphor that enriches and elaborates upon the speaker's own viewpoint. In effect, this exchange of paraphrases serves to align the understanding of both speaker and listener, and allows each to arrive at a deeper appreciation of a metaphor. For similar reasons, and for others that are distinctly computational, the ability to produce a literal paraphrase of a figurative statement is also of some practical value in a computational setting. Not only does paraphrasing allow a computer to explain its interpretation of a metaphor, in a way that allows users to detect a failure of interpretation, it also allows a computer to directly assign a semantic representation to a metaphor that is consistent with the system's own axioms, so that then it can reason safely about the metaphor's contents. Indeed, in a wide range of NLP systems—such as for sentiment analysis, information retrieval and extraction, question answering, and summarization—literal paraphrasing can support a perfectly adequate approach to metaphor meaning that is just as deep as the system's treatment of literal texts and literal meanings.

As shown in our earlier discussion of the metaphor "*wishes are children*," we need not limit ourselves to purely literal statements when paraphrasing a metaphor, and so, perhaps, neither should our NLP systems. The value of any paraphrase lies in its ability to explain a challenging turn of phrase using more conventional and less taxing language. This is so whether the text to be paraphrased uses the non-literal language of poetry or the literal language of legalistic jargon (as found, say, in a legal contract or a patent application). Indeed, the latter demonstrates that the most natural and most useful paraphrase is not always the most literal, and less challenging forms of figurative expression, such as a familiar conventional metaphor, may better convey the intended meaning of a metaphor. For example, when paraphrasing a metaphorical description of a wine as "*muscular*," it is more helpful than not to exploit the conventional if metaphorical view of wines as possessing a "*body*." Or consider the metaphor "*to pick up a touch of the flu*," which in turn contains a pair of tactile metaphors, "*to pick up*" and "*touch*." The first of these is perhaps more naturally paraphrased using the conventional metaphor "*to catch*," which is very commonly used with infections, rather than the strictly literal and somewhat dramatic "*become infected by*" or "*to suffer from*." The second is perhaps best paraphrased using the modifier "*mild*" for "*flu*," even though the idea of mildness is likely a metaphor itself in this context. Indeed, it is hard to find a concise literal rendering for the tactile metaphor "*touch*:" a touch suggests a glancing contact, which implies a contact of short duration that lacks force. So a "*touch of flu*" implies a short-lived infection that does not manifest the worst symptoms of the ailment.

An acceptance of the occasional conventional metaphor in a supposedly literal paraphrase is especially useful in the context of statistical models of metaphor, for it is in the nature of conventional metaphor to be used habitually and to assume the status of normative, literal language over time, to the extent that the two are hard to separate on statistical grounds alone. Indeed, much of what we consider literal was once a newly minted figurative form that, as the philosopher Nietzsche put it, has become the base metal of literal coinage through oft-repeated usage. An NLP

system that aims to remove all metaphor from language will find much—perhaps too much—to remove, and will be left with precious little with which to compose its paraphrases. In contrast to the divide between the figurative and the literal, the unconventional/conventional divide has an observable reality in large text corpora that is conducive to statistical analysis (e.g., Shutova et al. [2012], Bollegala and Shutova [2013]). It is a quantifiable reality that allows a paraphrasing system to learn to do as human paraphrasers do: to generate helpful conventional paraphrases from unconventional metaphors. If such paraphrases contain conventional metaphors, as they are very likely to do, this need not pose a serious challenge to the semantic representations of an NLP system. Conventional metaphors are part of the furniture of language, and can easily be accommodated—with some design forethought—in a system's core representations. For instance, both the *MIDAS* system of Martin [1990] and the *ATT-META* system of Barnden and Lee [2002], and Barnden [2006] put a semantic representation of conventional metaphors at the heart of their metaphor interpretation systems. *MIDAS* and *ATT-META* each show, through their different uses of schematic structures, how an NLP system can support a wider range of inferences, and, thus, a fuller interpretation of a metaphor, by reserving a place for the metaphor in its semantic and conceptual representation of an utterance.

## 2.4 METAPHOR AND SIMILE

Since the worldview communicated by a metaphor is often shaped by a speaker's perception of a situation or an interlocutor, simile—the figurative device most often called upon to simultaneously communicate and explain our perceptions—can be of particular use when paraphrasing a metaphorical worldview. Consider the metaphor "*marriage is slavery*:" while there is little insight to be gained by paraphrasing this metaphor with the simile "*marriage is like slavery*," there is much to be gained from the paraphrase "*this marriage is like the relationship between a slaver and his slave: you act like you own me, and treat me like your slave.*" The use of similes imbues this paraphrase with three interesting qualities. First, the similes bring a desirable emotional distance to the description of an undesirable relationship, for the assertion "*you act like you own me*" conveys a very different affect than "*you own me*" or "*you are my owner*." Second, the paraphrasing similes are not freighted with the same semantic tension as the original metaphor, since similes merely assert the similarity of two ideas and openly admit—through their use of "*as*" or "*like*"—to the counterfactuality of a viewpoint. That is, every simile "A is like B" asserts not just that *A is similar to B*, but that *A is very much not B*. Finally, the similes not only explain the meaning of the metaphor, but offer a rationale for it too: the paraphrase suggests that the metaphor is the speaker's way of making sense of the bad behavior of others.

The implicit negation in every simile makes simile a particularly good choice for paraphrasing a negative metaphor. For, just as explicit similes contain tacit negations, explicit negations often suggest tacit similes, especially when the negation emphasizes the figurative qualities of a descriptor. Since most negative metaphors are trivially true in a literal sense, insofar as the underlying positive metaphor is literally absurd, then each negative metaphor is also its own literal

paraphrase. Thus, it really is the case that no man is an island; that your wife is not your maid; that your overbearing boss does not own you; that your college fund is not an ATM; etc. Yet there is more to a negative metaphor than the negation of an obvious falsehood. A negative metaphor is not so much a disavowal of an explicit metaphor, or of the tacit conceptual pact that it implies, as it is an explicit repudiation of someone else's implicit simile. Why else would we need to say something that is so obviously true? So, the frustrated wife who cries at her boorish spouse "*I am not your maid!*" is in fact saying "*Do not treat me like your maid!*" The moody teenager who snaps at a concerned neighbor "*You are not my father*" is in fact saying "*Stop acting like my father.*" And the man who needs to be told that "*No man is an island*" is no doubt acting like someone who believes the opposite. Similes allow a speaker to be abundantly clear as to the perceptual foundations of a figurative viewpoint, and clear in ways that are hard to achieve in metaphor.

Giora et al. [2010] argue that negation, when used to convey such apparently obvious facts, is a *metaphor-inducing operation*. So, like metaphors and similes, negations of the false or the absurd are much more likely to activate the figurative aspects of a descriptor than any literal qualities. Metaphor involves highly selective inference (see Hobbs [1981]) and the assertion "*I'm no Superman*" selectively activates the qualities of the cultural icon (and very dense descriptor) *Superman* that are most often projected by a metaphor or a simile, such as *strength*, *resilience*, and *speed*, rather than any literal quality related to appearance (e.g., wearing a red cape, or red underwear on the outside) or behavior (e.g., fighting crime, working in disguise as a reporter, etc.). Precisely what these figurative qualities are for a given descriptor is something that one would have difficulty discerning from only metaphorical uses of the descriptor, as metaphors rarely make explicit the qualities that are projected onto the target. This brings us then to another aspect of similes that makes them so useful to the computational modeling of metaphor. Not only do similes explicitly mark their figurative status with "*like*" or "*as*," and frequently indicate the perceptual roots of their viewpoint with qualifiers such as "*act like*," "*smell like*," "*look like*," etc., they also often explicitly identify the qualities that are transferred from the source/vehicle to the target/tenor.

The scholar Quintilian saw the difference between metaphor and simile as the difference between an implicit and an explicit comparison. Not only does a simile mark the comparison between a *comparandum* and its *comparatum*, it may also provide a third element, a *tertium*, to state the reasons why both are seen as similar [Roberts, 2007]. In the English simile frame "*X is as Y as Z*" the tertium is given by the *[Y]* element, although a speaker may omit the tertium to say simply that "*X is like Z*." But when the tertium is explicit, as it always is in "*as-as*" similes, a listener can acquire from the similes of others a sense of the properties of a *[Z]* that are most likely to be activated in a figurative comparison. In this way, similes become an important vector for the transmission of cultural knowledge via language; as Charles Dickens puts it in the opening page of *A Christmas Carol*, "*the wisdom of our ancestors is in the simile*." Dickens was referring to the swirl of stock similes that were common currency in the language and culture of his day, but the modern scholar, or modern NLP system, now has access to a long-tail of diverse similes—

and diverse tertia—for a wide range of useful descriptors on the Web. Consider again the dense descriptor *Superman*: a Google search for the phrasal query "*as * as Superman*" (here * denotes a wildcard that can match any token in a text) returns web documents that fill the tertium * position with the following values: *strong*, *powerful*, *fast*, *mighty*, *invulnerable*, *resilient*, *cool*, and *American*. As these properties are frequently highlighted in figurative similes, it seems reasonable to assume that a metaphorical use of Superman will draw on the same properties.

The relationship between metaphor and simile is itself best described as a simile. Metaphor is like simile but it is not simile; nor is it wholly reducible to simile. Nonetheless, each concerns itself with the similarity of two ideas, and each focuses on qualities that readily extend across domains. So, just as similes can be used to make explicit the qualities that go unsaid in a metaphor, they can also be used to learn those very qualities, to provide—as we shall see later—the large set of dense descriptors that is required by any metaphor processing system.

## 2.5 METAPHOR AND ANALOGY

From Aristotle to Cicero to Quintilian, the earliest treatments of metaphor put similarity at the heart of the phenomenon, but agree on no single approach to formalizing similarity. Recall that Aristotle's four-fold scheme posits three ways in which inter-category similarity can be used to form a metaphor (strategies (a) to (c)) and a single way in which analogical similarity can be used (strategy (d)). Yet these strategies are not claimed to be mutually exclusive, and so, we can imagine a metaphor simultaneously exploiting both categorization and analogy together. It is thus useful to see the metaphor "*marriage is slavery*" as both a categorization statement (putting marriage in the same category of unpleasant situations as slavery) and as an analogy (so husbands are to wives, or wives are to husbands, as owners are to slaves). Yet the dichotomy between, on one hand, similarity arising from shared category membership, and on the other, similarity based on analogical proportion, is one that still separates modern theories of metaphor production and interpretation.

Aristotle's category membership approach survives, in a more finessed form, in Glucksberg's [1998, 2008] *category inclusion* view of metaphor. Whereas Aristotle's scheme (c)—in which the name of one species is applied to another of the same genus—presumes that source and target share a pre-existing membership in a common genus, Glucksberg argues this sharing is a *result* of the metaphor and not a cause. Moreover, since the Aristotelian scheme is easily trivialized when applied within a well-connected category hierarchy that ultimately places every concept under an all-embracing category root, Glucksberg's category inclusion view argues that the source and target categories cannot simply share, or be made to share, just any genus category. The point of metaphor, after all, is not just to assert that two ideas happen to share a common category; at the very least, a metaphor must also suggest a means of naming the specific category that the speaker has in mind. For Glucksberg, then, the source of a metaphor does not so much represent itself but the broader category of which it is a highly representative member. As such, "*slavery*" does not literally denote either the legal or historical sense of *slavery* in the metaphor "*marriage is slavery*,"

but any institutionalized system of exploitation and oppression. In contrast to the classical view of categorization, which models categories as simple mathematical sets, modern cognitive psychologists adopt a textured view of category structure in which some members are more central than others. The most central members may be so associated with a category—such as *shark* and *wolf* for the category of *ruthless predators*, *jail* for *oppressive and confining situations*, *slavery* for *cruel and exploitative relationships*—that they offer a more evocative, concise, and convenient way of naming that category than any literal alternative.

Yet Aristotle's strategy (d), analogical proportion, suggests how we might side-step this search for a mediating category altogether. For there is no need to find a common genus to unite the source and the target if the two can be reconciled by means of direct, unmediated relational similarities. By this reckoning, metaphors are made from the same stuff as scientific analogies: one observes a pattern of relationships in one domain, the tenor or target domain, and is reminded of a parallel set of relationships in another, the vehicle or source domain. But, as argued by the structure-mapping school of analogy [Gentner, 1983], a good analogy is more than a set of observed correspondences between domains. Rather, these correspondences—systematically linking two different conceptual structures—are just the starting point of an analogy. Having anchored elements of the source domain to the target domain, these correspondences then guide the transfer of additional material from the source into the target. For example, the Bohr/Rutherford analogy of atomic structure views the nucleus of the atom as occupying a relationally parallel position to the sun in a solar system, so that the electrons that orbit around the nucleus are the nanoscopic equivalent of the planets that orbit around the sun. Having established these correspondences, the analogy can now suggest a causal explanation for the orbit of electrons, by importing the causal explanation for the orbit of planets around a sun: for, just as the sun keeps speeding planets in orbit via a cosmic force, gravity, a comparable force must be attracting speeding electrons to stay in orbit around a nucleus. Gentner argues that many metaphors also work in this way: a source is chosen because of its analogical parallels to the target, and because it contains additional conceptual material that a speaker also wishes to assert of the target. A listener unpacks the metaphor by uncovering much the same parallels, and by then using these as a guide to the transfer of additional material from the source.

Consider again the metaphor "*marriage is slavery.*" A systematic analogy between *marriage* and *slavery* will identify abusive behaviors in both domains and create mappings between the protagonists and agonists of these behaviors. Thus, abusive husbands may be slave-owners and abused wives their slaves, while matchmakers can be mapped to slave dealers, family homes to plantations, wedding rings to shackles, marriage licenses to ownership papers, etc. A particular mapping is systematic if its elements are well-connected to each other, especially if they are connected via parallel causal relationships that explain how the various parts of a domain influence other parts of the same domain. Having established correspondences between ideas in the source and target domains, an analogy can now transplant onto the target any relationships that link these ideas in the source domain, to provide new insights as to the cause of these behaviors in the target

domain. For instance, slave owners beat their slaves while abusive husbands beat their wives. But slave owners beat their slaves because they believe they own them, and see them as mere physical objects to mistreat as they see fit. By analogy, abusive husbands may beat their wives because they, too, believe they own them, and because they, too, view them as mere physical objects.

Any model of metaphor that hinges on the identification of a common genus, whether it has the simplicity of Aristotle's earliest scheme or the sophistication of Glucksberg's category inclusion model, operates by reducing the specific to the generic. Such approaches are best suited to the generation and interpretation of familiar metaphors, or to apparently novel metaphors that dress up old conceits in new ways. In contrast, the structure-mapping approach does not aim to see, or need to see, the generic in the specific, and so, *ceteris paribus*, is just as capable of seeing the analogical proportions between the source and target ideas of a truly original metaphor as it is for those of a wholly conventional pairing. However, this also means that the analogical approach will fail to capture the sense of familiarity that one experiences when faced with a timeworn combination of ideas, just as it must fail to capture the thrill that accompanies a truly original pairing. Of course, even the most hackneyed metaphor was once fresh and novel, and, although time does not change the substance of our metaphors, it does change the way we see them, and perhaps even the way we process them.

With two alternate approaches to choose from, each of which is better suited to a different kind, or historical *stage*, of metaphor, it makes sense to view these not as competing but as complementary approaches. As proposed by Bowdle and Gentner [2005] in their *career of metaphor* hypothesis, fresh metaphors that do not obviously instantiate a familiar conceit are more naturally understood as analogies. Structure-mapping analysis allows listeners to identify the parallel relationships in source and target domains that contribute to a new metaphor, thus allowing them to abstract from the source those parts of its meaning that are most likely to contribute to other metaphors. As speakers become habituated to a given metaphor, and to other metaphors that use the same source for similar ends, they come to think of the source as a vehicle for those parts of its conceptual structure that are transferable across domains. In effect, they form an abstracted representation of the source that ultimately allows the metaphor, and similar metaphors with the same source, to be understood in category inclusion terms.

The earliest account of metaphor, as offered by Aristotle, acknowledges that metaphor is a complex phenomenon that calls for a varied approach. Although any mixed model that recognizes the *career* of metaphor will inevitably lack the parsimony of a one-size-fits-all approach, it is surely preferable to tailor our models to the phenomenon than to cut the phenomenon to fit our models.

### 2.5.1 DOMAIN REPRESENTATIONS IN METAPHOR AND ANALOGY

As much as language unites us in our use of words to describe the world to each other, we are all free to think what we will about the world, or at least those aspects of the world that interest us the most. Different speakers may thus employ very different representations of the same domains, leading to strong disagreements about the most natural interpretation of a metaphor or an analogy

involving those domains. Consider this exchange from the film *Jurassic Park*, between the owner and operator of the park (Hammond) and a mathematician (Malcolm) who has been tasked with evaluating the park:

> **John Hammond**: All major theme parks have delays. When they opened Disneyland in 1956, nothing worked!
>
> **Dr. Ian Malcolm**: Yeah, but, John, if The Pirates of the Caribbean breaks down, the pirates don't eat the tourists.

Here we see Hammond attempt to mitigate problems with his park by comparing them to the teething difficulties experienced in perhaps the most prototypical theme park of all, *Disneyland*. His representation of both domains (*Disneyland* and *Jurassic Park*) is thus high-level and generic, and so the resulting analogy lacks detail. If it is not so very different from an analogy comparing *Amazon.com* to *Barnes and Noble*, this blandness speaks to Hammond's larger goal in using this analogy. In response, Malcolm focuses on the specific problems at hand (namely, the hungry dinosaurs running amok all about them), and so brings to mind a very specific representation of *Disneyland* to suit the very specific representation of *Jurassic Park* that circumstances have forced upon him. Had Malcolm made an additional joke about both parks suffering very different kinds of *teething* problems, this pun would require an even more detailed representation of the source and target domains (e.g., technical glitches in the representation of *Disneyland* could map to the problems wrought by the dinosaurs' sharp teeth in the representation of *Jurassic Park*). Analogies and metaphors that are used to persuade are not simple one-shot efforts at communication. Rather, as argued in Cameron [2007], they form part of a larger framework of negotiation and alignment that allows speaker and listener to focus on the same aspects of a domain (or to at least see the other's point of view). As highlighted by this example, the representation of the domains in an analogy (and thus any metaphor based on an analogy) is open to negotiation by the speakers during the formation and interpretation of the analogy. It is not realistic to assume that these representations are pre-formed before the analogy is made and simply retrieved from memory to participate in a process of structure mapping. Rather, it seems more accurate to suppose that the search processes required by analogy involve more than a search through the space of possible alignments between two representations, but also a search through the space of possible representations to align.

## 2.6 CONCEPTUAL METAPHOR THEORY

Black and Davidson each tell us that the meaning of metaphor is not to be found in word meanings, but in the world itself, or at least in our models of the world. Lexical semantics can only take an NLP system so far, and the crucial ingredient in any metaphor-processing system is its array of conceptual representations. The field of cognitive linguistics, as exemplified by the work of Mark Johnson and George Lakoff [Johnson, 1987, Lakoff and Johnson, 1980], goes further

still in emphasizing the role of world knowledge. A metaphor-processing system—whether human or machine—must possess an *embodied* understanding of the world: it must understand the mind's relation to the body and the body's relation to the world. These relationships are not literal, but are themselves governed by foundational metaphors that are grounded in our physical understanding of space, orientation, movement, and containment. These schematic structures––embodied concepts such as *Path* and *Container* and the metaphors built on top of them—are independent of language. Indeed, Lakoff and Johnson argue that metaphor is not primarily a linguistic phenomenon at all. Language is merely the stage on which metaphor does its most noticeable work.

The conceptual perspective sees metaphors as more than the finished products of creative language but as the very building blocks of thought itself. Complex world-views can be constructed from these conceptual metaphors, just as complex metaphors can themselves be constructed from simpler, lower-level metaphors. Grady [1997] argues that physical embodiment is the key to acquiring the latter kind of metaphor, which he dubs "*primary*." For instance, from a young age we come to associate—or as Johnson [1999] puts it, *conflate*—affection with physical closeness and shared bodily warmth, and continue throughout our lives to express affection through physical touch. It is unsurprising then that we develop the primary metaphor (Emotional) Affection is (Physical) Warmth, which allows us to talk of close relationships as *warm* relationships, of distant relationships as *cold* relationships, of giving someone "*the cold shoulder*" or "*a warm welcome*." Primary metaphors are not arbitrary: rather, just as they are born of past physical experience they also offer a useful prediction of future physical experience. Thus, we can expect a *warm welcome* to involve a physical embrace, while we can also expect those who give us the *cold shoulder* to keep their distance from us.

So the physical world serves as a sand-box in which a cognitive agent learns to associate cause and effect and to fit cognitive structures to concrete reality. For instance, we learn that objects are usefully grouped into piles; that adding to a pile makes it bigger and higher; that taking from a pile makes it smaller and lower; and thus, more generally, that more of anything implies upward accumulation, while less of anything implies downward reduction. In other words, we learn that More is Up and Less is Down. Likewise, our interactions with physical containers informs us of their affordances and limitations, and allows us to generalize a more abstract notion of containment that is conducive to cross-domain metaphorical reasoning. To appreciate the power of conceptual metaphors in a physical, non-linguistic setting, consider this tale of the Parsis and how they first came to settle in India. *Time Out Guide to Mumbai* offers a concise version of the legend:

> *They* [The Parsis] *arrived in Gujarat in the eighth or ninth century and sought asylum from the local king. He is said to have sent them away with a glass of milk full to the brim—his way of saying that his kingdom was full. The Parsi elders conferred, added some sugar to the milk and sent it back—to suggest that they would mix thoroughly and sweeten the life of the community.*

So, with the right objects to hand, one does not need words to communicate with metaphors. The glass proves to be a remarkably versatile vehicle of meaning, for both the king's message and that of the Parsi elders, because it instantiates a highly productive cognitive structure called the Container schema (see Johnson [1987]). The king uses the glass to represent a country: a country, like a glass, is a container, not just of liquids but of people and places and resources. Like a typical container, a country has physical boundaries that mark its extent; like a typical container, one can put things into a country or one can take them out again; and, like a typical container, a country can be full (of people) or it can be empty. Thus, the king uses a full glass of milk to convey the conceptual metaphor A Country is A Container, but with the kicker "*this one is full*." As the glass shuttles back and forth, it is used to carry subtly different meanings, each pivoting on the figurative affordances of the Container schema. The Parsis fashion their riposte from a different Container metaphor, Life is a Container, one that can be figuratively filled with diverse events, relationships, and feelings. In contrast to the physical affordances of the Country is a Container metaphor, which implies that countries are physical containers with physical limits on the amount of physical contents they can hold, Life is a Container is a more abstract metaphor that places no such limits on its abstract contents. If Life is a Container, it is a container with an unbounded capacity for emotional and cultural possibilities. The king and the Parsis are effectively playing a game of tropes, where the Parsis trump the king's use of the Container schema with a more creative metaphor of their own.

The Container schema is ubiquitous in language and thought. We talk of minds as containers of beliefs (perhaps "*bursting with ideas*"), of lives as containers of events (allowing some to live "*a full life*"), of texts as containers of words, and words as containers of meanings. Our relationships can be figuratively conceived as containers, too (as in "*there were three of us in that marriage*"), although we most often anchor our metaphors for personal relationships in the Source-Path-Goal (or SPG) schema. Thus, we talk of relationships that "*are going nowhere*," of relationships that have "*gone through a rough patch*" and are still "*on the rocks*," or those that have "*hit the buffers*" or even "*come off the rails*." Purposeful activities of all kinds are conceptually anchored in the SPG schema [Johnson, 1987], allowing us to talk of "*career paths*," of being "*held back*," and of "*getting ahead*." Schemas such as SPG and Container allow us to talk of the big intangible ideas, such as Life and Love, in terms of physical objects that we can see and grasp and follow. When said out loud, SPG-based conceptual metaphors such as Life is a Journey sound like the most hackneyed of clichés, yet, as Lakoff and Johnson [1980] demonstrate, they exude a powerful influence on the way we see and talk about the world. They are evident in the way our thoughts and our texts often cohere around a single unifying metaphor. Although conceptual metaphors are most apparent in our clichés, they underpin even our most creative efforts. Whenever we riff on a creative metaphor, as the Parsis did with the king's figurative use of a physical container, we are instantiating and extending a conceptual metaphor.

Unsurprisingly, Conceptual Metaphor Theory (CMT) has significant implications for the computational modeling of metaphor in language [Lakoff, 1994, Lakoff and Johnson, 1980].

CMT-inspired computational models of metaphor range from symbolic approaches that focus on high-level data structures and knowledge-representations (e.g., Martin [1990]), to cognitive and neurological approaches that focus on lower-level arrangements of neuron-like activation elements and their interconnections (e.g., Feldman [2006]). But, regardless of which view of CMT is implemented, it is not enough that an NLP system goes beyond mere lexical semantics to embrace rich conceptual representations of target and source domains. It must embrace the *right kinds* of conceptual representations. Most theories of metaphor that lend themselves to a computational realization are ultimately agnostic about conceptual structure. Although they stipulate that such structures are necessary, and provide the general form of these structures—whether the set-theoretic genus/species representations of Aristotle, the graph-theoretic representations of Gentner's analogical structure-mapping approach, or the textured category representations of Glucksberg's category inclusion model—they do not make ontological commitments to the representations of particular concepts, such as Container, Path, Orientation, Force, etc. CMT argues that, for any agent to make and understand metaphors like a human, it must have the same conceptual biases as a human, and anchor its model of the world in the same foundational schemas. However, while these schematic concepts have an embodied basis in human cognition, this is not to suggest that an NLP system must also be embodied in the same way. Developmental roboticists may find some value in embodying a computational realization of CMT in an anthropomorphic robot, but this is neither necessary nor useful for a metaphor-capable NLP system. Such schemas are the conceptual products of physical embodiment, and it is enough that our systems also contain representations of these schemas that afford the same kinds of conceptual inferences.

For example, the *MIDAS* system of Martin [1990] employs the basic schematic structures of CMT to anchor and guide its interpretation of conventional metaphors and variations thereof. A single conventional metaphor may be instantiated in a wide range of domain-specific variations, as Martin demonstrates in the context of a dialogue system for offering Unix advice. This domain offers ample opportunity for the use of physically motivated conceptual schemas to describe otherwise intangible phenomena, since, e.g., computer users conventionally conceive of software processes as living things (prompting the metaphor "*How do I kill a runaway process?*") and software environments as containers (prompting the metaphor "*How do I get out of Emacs?*"). In another CMT-inspired approach, Veale and Keane [1992] employ the schemas of Orientation and Container as semantic primitives in a lexico-conceptual representation they call a *Conceptual Scaffolding*. In contrast to other systems of semantic primitives, which conventionally employ such primitives as final, irreducible components of utterance meaning, these operators are intermediate placeholders that are designed to be replaced by domain-specific constructs during subsequent processes of interpretation and inference. Scaffolding operators are always replaced in this way, whether the speaker intends an utterance to be read literally or figuratively. Literal and metaphoric uses of image-schematic notions of containment (such as "*falling into a manhole*" vs. "*falling into debt*") and orientation (such as "*sliding down the hill*" vs. "*sliding down the polls*")

will thus use precisely the same operators in precisely the same way in their respective scaffolding structures. The versatility of the CMT schemas allows the scaffolding construction process to be agnostic as to the literal or figurative status of an utterance, and pushes a decision on this distinction into the realm of inference and common-sense reasoning, where Black [1962] and Davidson [1978] each say it belongs. Alternatively, one can use CMT's schematic structures as a cognitively grounded vocabulary for tagging linguistic metaphors in large corpora with the appropriate conceptual annotations. A large corpus of metaphors annotated in this fashion may allow a computational system to learn to recognize and categorize unseen variations on these metaphors in free text, or to, perhaps, even generate novel variations of their own.

## 2.7 CONCEPTUAL INTEGRATION AND BLENDING

If figurative phenomena blur the boundaries between very different concepts and domains, we should not be surprised if they also blur the boundaries between themselves. It can be quite a challenge, even for a seasoned researcher, to determine precisely what figurative mechanism is at work in any given example of non-literal language. Metonymy and metaphor are often tightly bound up together, with metonymy providing the subtle glue to seamlessly tie metaphors to their targets. Thus, in the news headlines "*Chipotle shares hit by E. Coli scare*" and "*Nike jumps on share buyback*," the metonymies $company \rightarrow share$ and $share \rightarrow share\_price$ are necessary for the respective metaphors, Hitting is Diminishing and Jumping is Increasing, to take hold (note also the punning humor of the verb "*jump*" in the context of Nike, a maker of basketball shoes; Brône and Feyaerts [2005] describe this phenomenon as *double grounding*). Different kinds of figurative phenomena are not always marked at the linguistic level, and so an implicit comparison, in which the source and target are not explicitly identified with each other in the same utterance, may be taken as a metaphor by some, a simile by others, or an analogy by dissenters to both of these views. Moreover, what seems like a comparison to some may be taken as a categorization by others. Consider again our earlier example from the movie *Jurassic Park*. In particular, consider again the quote from the park owner, John Hammond:

> **Hammond**: All major theme parks have delays. When they opened Disneyland in 1956, nothing worked!

It makes sense here to understand Hammond's observation as an implicit comparison between Disneyland and Jurassic Park, although the latter is not explicitly mentioned. Nothing is seemingly working in Jurassic Park, but nothing worked in 1956 at Disneyland either, yet the latter turned out to be a monumental financial and cultural success. With this statement, Hammond thus implies comparable future success for his own troubled venture. However, note the generalization that precedes this comparison: "*All major theme parks have delays*." Hammond here seems to be establishing a categorization, in which Jurassic Park and Disneyland are to be placed in the same category, namely the category for which Disneyland is prototypical: a major theme park that started life inauspiciously but went on to win the hearts and wallets of America.

If Hammond's metaphor is simultaneously both a comparison and a categorization (and not one *or* the other) then Malcolm's rejoinder seems to be an entirely different species of figurative beast:

> **Malcolm**: Yeah, but, John, if The Pirates of the Caribbean breaks down, the pirates don't eat the tourists.

In fact, Malcolm's conversational gambit is best described as a *conceptual blend* (see Fauconnier and Turner [2002]). Not only does it operate on the same domains as Hammond's gambit—Disneyland and Jurassic Park—it also thoroughly integrates a selection of elements from each of these domains to create an entirely new conceptual structure, one that contains emergent ideas and images that arise from neither input alone. Once again, Malcolm's remark is intended to be understood in the context of a juxtaposition of Disneyland and Jurassic Park, but it is much more than a comparison, categorization, or analogy. It involves analogy, of course, as a constituent process: thus, The Pirates of the Caribbean is implicitly aligned with the attractions of Jurassic Park, and so the animatronic pirates of the former are aligned with the genetically-engineered dinosaurs of the latter. But the salient behaviors of the latter—such as eating people willy-nilly—are also integrated with the protagonists of the former, to generate a delightful counterfactual image of mechanical pirates eating hapless tourists in mouse-earred caps. In the terminology of Fauconnier and Turner [2002], Malcolm is *running the blend*, which is to say he is conducting a mental simulation within the blend space that results from the integration of the inputs, to stimulate his imagination into recognizing the emergent possibilities that were hitherto just latent in the blend.

Metaphor has traditionally been modeled as a two-space phenomenon, in which conceptual content in one mental space—representing the vehicle or source domain—is projected onto, or mapped onto, the corresponding content in another space representing the tenor or target domain. Conceptual blending theory adds to the number of spaces that are implicated in figurative processing and further clarifies the notion of a mental space. Thus, a *generic space* provides low-level schematic structures (such as the SPG and Container schemas) that can serve to unite the disparate elements of the input spaces, while the *blend space* is a new space, distinct from any of the input spaces, in which content from the inputs is selectively projected and integrated. When Malcolm *runs* the blend to imagine pirates eating tourists, his mental simulation takes place here, in the newly constructed blend space. The set of all spaces that contribute to a blend, and the connections that link them, is called a *conceptual integration network*. Following Fauconnier [1994, 1997], a mental space is defined to be an ad-hoc bundle of domain-related information that a speaker or listener brings to the comprehension process. Notice, in the example above, Hammond and Malcolm both refer to the domain of Jurassic Park (qua theme park), but each builds a different mental space for this domain. For Hammond, the park is a business venture beset by temporary technical glitches; for Malcolm, the park is a place born of hubris and beset by rampaging carnivores. As shown by Malcolm's rejoinder, blending is especially suited to the construction and comprehension of counterfactual scenarios, since the blend space serves as a mental sand-

box in which one can playfully experiment with the consequences of any given mapping between input spaces.

Blending theory also provides a relatively seamless framework for exploring and understanding the interactions between metaphor, analogy, metonymy, and concept invention. Indeed, one might argue that it is this seamlessness that often makes these phenomena so difficult to tease apart in the context of real examples (see Barnden [2010], who explores the "*slippery*" linkage between metaphor and metonymy). Consider the compound noun "*dinosaur hunter*," commonly used in the media to describe those paleontologists that seek out the skeletal remains of prehistoric creatures. Is the word "*hunter*" used metaphorically here, as in the compound "*bargain hunter*," or—since dinosaurs are animals every bit as fierce as elephants, rhinos, and lions—is it used literally (if a trifle unconventionally)? We can sidestep this question, while still arriving at an acceptable interpretation, by viewing the compound as a blend, in which the space of paleontology is integrated with the space of hunting to produce a new kind of hunter:prey relationship, one in which the hunter is temporally separated from his prey by 65 million years. That these hunters never actually see their prey in the flesh is beside the point, since metonymy allows us to conflate these long-dead dinosaurs with their skeletons. Indeed, Fauconnier and Turner [2002] argue that a common side-effect of conceptual blending is *metonymic tightening*, whereby the chains of associations that link the input spaces are compressed to create a tighter integration of conceptual content. So, when dinosaur skeletons are displayed in natural history museums, they are typically displayed as full-fledged dinosaurs, not as the mere parts of dinosaurs. If such displays are expertly framed so as to compress the holonymic relationship between a dinosaur and its bones into an identity relationship, blends such as "*dinosaur hunter*" do as much with language to achieve the same effect.

Fauconnier and Turner's [2002] blending theory (or the theory of *conceptual integration networks* to give it its formal title) sets out to explain—via a general architecture of constraints and optimality principles—how blended spaces are constructed via the selective projection of elements from multiple input spaces. But, how are those input spaces constructed in the first place? In the case of blends such as "*dinosaur hunter*," it is not unreasonable to assume that the mental spaces for *dinosaur* and *hunter* are derived, in part at least, from the corresponding entries in our mental lexicons. In other cases, these inputs may be the products of lower-level processes, such as the processes of conflation and differentiation that lead to the development of primary metaphors [Grady, 2005]. Lakoff and Johnson [1999] argue, for instance, that the neurally grounded processes that give rise to primary metaphors are complemented by conceptual processes that blend these primary inputs into more complex metaphors. For example, the conceptual metaphor Life is a (Purposeful) Journey can be considered a blend of the primary metaphors Purposes are Destinations (which we saw earlier in Wernher von Braun's metaphor of "*aiming for the stars*") and Actions are Motions (as when, e.g., a plan becomes reality when it is "*put in motion*"). In yet other cases, the inputs to a blend may themselves be the products of conceptual blending, which is to say, the blend space of one conceptual integration network may serve as an input space

to another. Blends are conceptual products, after all, and successful blends may gain currency in a culture, so that they are given an evocative name or paired with a linguistic form to yield a reusable construction that achieves the status of a cultural trope.

Consider the following description of the film director Sam Mendes from a British newspaper in 2010: "*appearance: like the painting in George Clooney's attic.*" The newspaper was moved to sarcastic commentary after Mendes had been announced as the director of the next movie in the *James Bond* franchise, named *Skyfall*, and this sarcasm extended to its wry perspective on Clooney's carefully-constructed media image. The use of the definite article "*the*" before "*painting*" would suggest that this is a picture previously known to us, but, lacking any knowledge of Clooney's attic or of his taste in art, we must fall back on more general, cultural knowledge instead. Fortunately, the linguistic construction "*painting in [the] attic*" is strongly evocative of Oscar Wilde's morality tale *The Picture of Dorian Gray*, not because we remember the actual text of the story but because we are familiar with the phrase's repeated use in popular cultural as a shorthand for hidden excess and unnatural youthfulness. Recall that Wilde's novel centers around the gilded youth of the title and his bargain with fate: as Dorian remains forever young and beautiful on the outside, the ravages that time, sin, and excess should have wrought on his appearance are instead reflected in his portrait, which Gray wisely conceals in his attic. The story is so entrenched in the popular imagination that it takes just the merest mention of "*the*" painting in the attic to bring this particular picture to mind. Like the *Da...Dum...Da..Dum..DaDumDaDum* theme from the movie *Jaws*, or Monty Norman's signature theme from the *James Bond* movies, this simple construction evokes a wealth of unspoken narrative expectations. Of course, the fictitious painting lurking in Clooney's attic is not literally a painting of Dorian Gray; rather, it is a painting of Clooney as Gray, which is to say, the painting of a counterfactual Clooney who has committed to the same body-for-soul deal with fate as Wilde's anti-hero. In other words, "*the painting in George Clooney's attic*" is a cue to create a conceptual blend from the input spaces *George Clooney* and *Dorian Gray*. A visual representation of the resulting integration network is presented in Figure 2.1.

The blend space of this integration network now becomes available as input to the higher-level blend, where it is integrated with the target of the original simile, *Sam Mendes*. Because blends allow partial and highly selective projection of elements from their inputs onto their resulting blend spaces, what gets projected from this newly minted input space is the notion of a picture of George Clooney that differs significantly from the public perception of the actor. In the blend space of this new integration network, Mendes is identified with the figure in the painting—which, following the logic of Wilde's novel, we imagine as portraying Clooney as an older, wrinklier, flabbier, and generally more dissolute version of himself. So, in this blend space, Mendes is *not* identified with Clooney as we see him in the media and on the silver screen, but with this far less attractive counterfactual version of Clooney. Figure 2.2 shows the combination of both integration networks that a reader must construct to obtain the meaning of the original simile (which, at its most reductive, asserts that Mendes resembles an older, wrinklier, flabbier, and generally more dissolute version of Clooney). Note that we have elided from our analysis

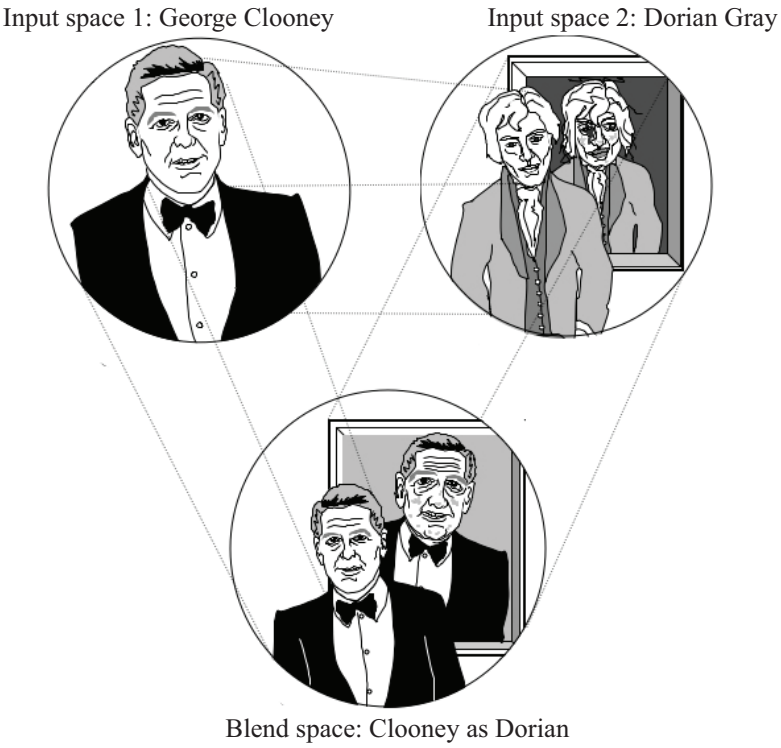


**Figure 2.1:** A conceptual integration of two mental spaces, one containing George Clooney and one containing Dorian Gray (reproduced from Veale [2012a]).

here the metonymy that is necessary to understand a comparison between a person and a painting: Mendes does not look like the painting itself (a flat painted surface in a decorative frame) but like the person that we imagine to be depicted in the painting. So, lurking beneath this simple-seeming simile is a blend within a blend, which in turn hinges on a metaphor (identifying Clooney and his imagined lifestyle choices with Gray and his soul-consuming narcissism), another simile (between Mendes and a less attractive version of Clooney), and a metonymy (allowing portraits to stand for their subjects).

Fauconnier and Turner [2002] effectively espouse a non-Davidsonian approach to figurative language, inasmuch as they propose that beneath the surface of every such utterance there is a cognitive content that one can call its *meaning*. So the integration network illustrated in Figure 2.2—with its component spaces and the many cross-connections between them—represents the cognitive content that listeners must construct for themselves in order to claim to have understood the meaning of "*appearance: like the painting in George Clooney's attic*" in the context of its target, *Sam Mendes*. Davidson argues that the meaning of a metaphor is essentially open-ended, but so too is the meaning of a blend. Listeners can, for instance, *run the blend* for themselves to

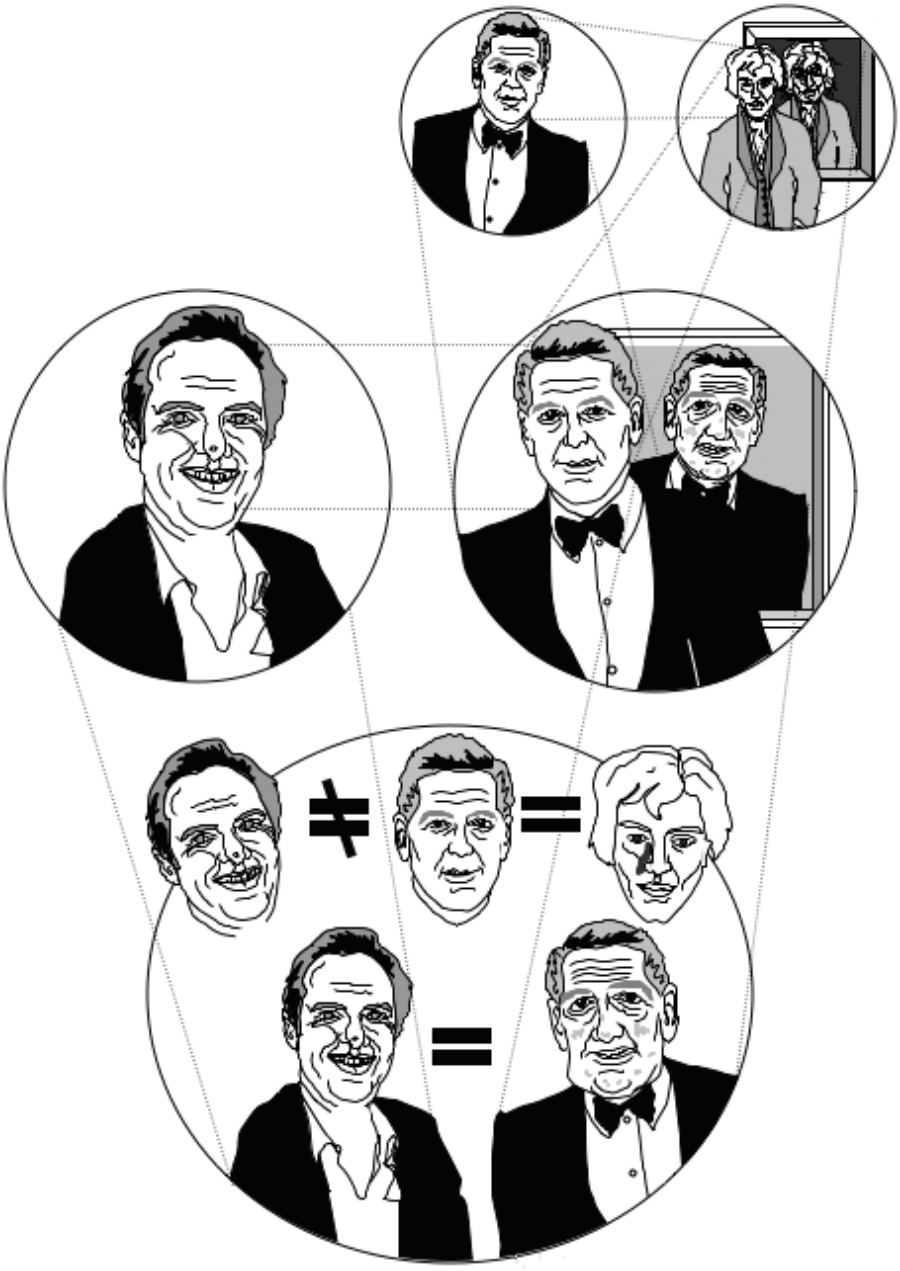

**Figure 2.2:** A conceptual integration of two mental spaces, one containing Sam Mendes and one containing the blend space from Figure 2.1 (reproduced from Veale [2012a]).