

Metaphor

A Computational Perspective



Synthesis Lectures on Human Language Technologies

Editor

Graeme Hirst, *University of Toronto*

Synthesis Lectures on Human Language Technologies is edited by Graeme Hirst of the University of Toronto. The series consists of 50- to 150-page monographs on topics relating to natural language processing, computational linguistics, information retrieval, and spoken language understanding. Emphasis is on important new techniques, on new applications, and on topics that combine two or more HLT subfields.

Metaphor: A Computational Perspective

Tony Veale, Ekaterina Shutova, and Beata Beigman Klebanov
2016

Grammatical Inference for Computational Linguistics

Jeffrey Heinz, Colin de la Higuera, and Menno van Zaanen
2015

Automatic Detection of Verbal Deception

Eileen Fitzpatrick, Joan Bachenko, and Tommaso Fornaciari
2015

Natural Language Processing for Social Media

Atefeh Farzindar and Diana Inkpen
2015

Semantic Similarity from Natural Language and Ontology Analysis

Sébastien Harispe, Sylvie Ranwez, Stefan Janaqi, and Jacky Montmain
2015

Learning to Rank for Information Retrieval and Natural Language Processing, Second Edition

Hang Li
2014

Ontology-Based Interpretation of Natural Language

Philipp Cimiano, Christina Unger, and John McCrae
2014

[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

Copyright © 2016 by Morgan & Claypool

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means—electronic, mechanical, photocopy, recording, or any other except for brief quotations in printed reviews, without the prior permission of the publisher.

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	xi
1	Introducing Metaphor	1
2	Computational Approaches to Metaphor: Theoretical Foundations	5
	2.1 The <i>What, Why</i> and <i>How</i> of Metaphor	5
	2.2 The “Meaning” of Metaphor	9
	2.3 Paraphrasing Metaphor	11
	2.4 Metaphor and Simile	15
	2.5 Metaphor and Analogy	17
	2.5.1 Domain Representations in Metaphor and Analogy	19
	2.6 Conceptual Metaphor Theory	20
	2.7 Conceptual Integration and Blending	24
	2.8 An Integrated Perspective	30
3	Artificial Intelligence and Metaphor	33
	3.1 Corrective Approaches: Metaphor Diagnosis and Repair	33
	3.2 Analogical Approaches: Mapping Meanings between Domains	39
	3.3 Schematic Approaches: Families of Metaphor	45
	3.4 Common Themes and Future Prospects	52
4	Metaphor Annotation	55
	4.1 Metaphor Identification in Corpus Linguistics	55
	4.2 The MIP Overall	56
	4.3 Specifications of Auxiliary Concepts	58
	4.3.1 Sense Inventories	58
	4.3.2 Criterion for Sense Distinctiveness	59
	4.3.3 Lexical Unit	60
	4.4 The MIP in Languages other than English	63
	4.5 Annotated Datasets	64
	4.5.1 Categorized Collections of Metaphors	64

	4.5.2 Annotations of Specific Constructions or Target expressions	66
	4.5.3 Full-Text Annotations	67
	4.5.4 Summary of Annotated Data Available for Public or License-Based Use	70
	4.5.5 Additional Annotated Materials	70
5	Knowledge Acquisition and Metaphor	73
	5.1 WordNet and other Lexical Ontologies	74
	5.2 Extracting Knowledge from Free Text	79
	5.2.1 Simile Patterns	80
	5.2.2 Categorization Patterns	82
	5.2.3 Arbitrary Relations	85
	5.3 Conceptual Metaphors	86
	5.4 Summary	87
6	Statistical Approaches to Metaphor	89
	6.1 Association Measures and Selectional Preferences	89
	6.2 Supervised Classification	94
	6.3 Clustering	99
	6.4 The Topical Structure of Text	102
	6.5 Vector Space Models	104
	6.6 Concreteness	105
7	Applications of Metaphor Processing	109
	7.1 Metaphor in Culture, Communication, and Education	109
	7.1.1 Common Metaphors in Cultures and Discourse Communities	110
	7.1.2 Metaphor as a Framing Device in Political Communication	110
	7.1.3 Metaphor Comprehension and Use as a Skill	111
	7.2 Creativity Applications in NLP	112
	7.2.1 Creative Writing Tools	112
	7.2.2 Creative Information Retrieval of Figurative Possibilities	114
	7.2.3 Narrative Possibilities: From Metaphor to Metamorphosis	115
8	Conclusions	119
	Bibliography	123
	Authors' Biographies	147

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

2 1. INTRODUCING METAPHOR

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 biting 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

4 1. INTRODUCING METAPHOR

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.