

Specifying Viewpoint and Information Need with Affective Metaphors

A System Demonstration of the *Metaphor Magnet* Web App/Service

Tony Veale

Web Science and Technology Division,
KAIST, Daejeon,
South Korea.

Tony.Veale@gmail.com

Guofu Li

School of Computer Science & Informatics,
University College Dublin,
Belfield, Dublin D4, Ireland.

Yanfen.Hao@UCD.ie

Abstract

Metaphors pervade our language because they are elastic enough to allow a speaker to express an affective viewpoint on a topic without committing to a specific meaning. This balance of expressiveness and indeterminism means that metaphors are just as useful for eliciting information as they are for conveying information. We explore here, via a demonstration of a system for metaphor interpretation and generation called *Metaphor Magnet*, the practical uses of metaphor as a basis for formulating affective information queries. We also consider the kinds of deep and shallow stereotypical knowledge that are needed for such a system, and demonstrate how they can be acquired from corpora and the web.

1 Introduction

Metaphor is perhaps the most flexible and adaptive tool in the human communication toolbox. It is suited to any domain of discourse, to any register, and to the description of any concept we desire. Speakers use metaphor to communicate not just meanings, but their feelings about those meanings. The open-ended nature of metaphor interpretation means that we can use metaphor to simultaneously express and elicit opinions about a given topic. Metaphors are flexible conceits that allow us to express a position while seeking elaboration or refutation of this position from others. A metaphor is neither true or false, but a conceptual model that allow speakers to negotiate a common viewpoint.

Computational models for the interpretation and elaboration of metaphors should allow speakers to exploit the same flexibility of expression with machines as they enjoy with other humans. Such a goal clearly requires a great deal of knowledge, since metaphor is a knowledge-hungry mechanism *par excellence* (see Fass, 1997). However, much of the knowledge required for metaphor interpretation is already implicit in the large body of metaphors that are active in a community (see Martin, 1990; Mason, 2004). Existing metaphors are themselves a valuable source of knowledge for the production of new metaphors, so much so that a system can mine the relevant knowledge from corpora of figurative text (e.g. see Veale, 2011; Shutova, 2010).

One area of human-machine interaction that can clearly benefit from a competence in metaphor is that of information retrieval (IR). Speakers use metaphors with ease when eliciting information from each other, as e.g. when one suggests that a certain CEO is a tyrant or a god, or that a certain company is a dinosaur while another is a cult. Those that agree might respond by elaborating the metaphor and providing substantiating evidence, while those that disagree might refute the metaphor and switch to another of their own choosing. A well-chosen metaphor can provide the talking points for an informed conversation, allowing a speaker to elicit the desired knowledge as a combination of objective and subjective elements.

In IR, such a capability should allow searchers to express their information needs subjectively, via affective metaphors like “X is a cult”. The goal, of course, is not just to retrieve documents that make explicit use of the same metaphor – a literal matching of non-literal texts is of limited use – but to

retrieve texts whose own metaphors are consonant with those of the searcher, and which elaborate upon the same talking points. This requires a computer to understand the user’s metaphor, to appreciate how other metaphors might convey the same affective viewpoint, and to understand the different guises these metaphors might assume in a text.

IR extends the reach of its retrieval efforts by expanding the query it is given, in an attempt to make explicit what the user has left implicit. Metaphors, like under-specified queries, have rich meanings that are, for the most part, implicit: they imply and suggest much more than they specify. An expansionist approach to metaphor meaning, in which an affective metaphor is interpreted by generating the space of related metaphors and talking points that it implies, is thus very much suited to a more creative vision of IR, as e.g. suggested by Veale (2011). To expand a metaphorical query (like “company-X is a cult” or “company-Y is a dinosaur” or “Z was a tyrant”), a system must first expand the metaphor itself, into a set of plausible construals of the metaphor (e.g. a company that is viewed as a *dinosaur* will likely be *powerful*, but also *bloated*, *lumbering* and *slow*).

The system described in this paper, *Metaphor Magnet*, demonstrates this expansionist approach to metaphorical inference. Users express queries in the form of affective metaphors or similes, perhaps using explicit + or – tags to denote a positive or negative spin on a given concept. For instance, “*Google is as –powerful as Microsoft*” does not look for documents that literally contain this simile, but documents that express viewpoints that are implied by this simile, that is, documents that discuss the negative implications of Google’s power, where these implications are first understood in relation to Microsoft. The system does this by first considering the metaphors that are conventionally used to describe Microsoft, focusing only on those metaphors that evoke the property *powerful*, and which cast a negative light on Microsoft. The implications of these metaphors (e.g., *dinosaur*, *bully*, *monopoly*, etc.) are then examined in the context of Google, using the metaphors that are typically used to describe Google as a guide to what is most apt. Thus, since Google is often described as a *giant* in web texts, the negative properties and behaviors of a stereotypical giant – like *lumbering* and *sprawling* – will be considered apt and highlighted.

To perform this kind of analysis reliably, for a

wide range of metaphors and an even wider range of topics, requires a robustly shallow approach. We exploit the fact that the Google n-grams (Brants and Franz, 2006) contains a great many copula metaphors of the form “X is a Y” to understand how X is typically viewed on the web. We further exploit a large dictionary of affective stereotypes to provide an understanding of the +/- properties and behaviors of each source concept Y. Combining these resources allows the *Metaphor Magnet* system to understand the implications of a metaphorical query “X as Z” in terms of the qualities that are typically considered salient for Z and which have been corpus-attested as apt for X.

We describe the construction of our lexicon of affective stereotypes in section 2. Each stereotype is associated with a set of typical properties and behaviors (like *sprawling* for *giant*, or *inspiring* for *guru*), where the overall affect of each stereotype depends on which subset of qualities is activated in a given context (e.g., *giant* can be construed positively or negatively, as can *baby*, *soldier*, etc.). We describe how *Metaphor Magnet* exploits these stereotypes in section 3, before providing a worked example in section 4 and screenshots in section 5.

2 An Affective Lexicon of Stereotypes

We construct the lexicon in two stages. In the first stage, a large collection of stereotypical descriptions is harvested from the Web. As in Liu *et al.* (2003), our goal is to acquire a lightweight common-sense representation of many everyday concepts. In the second stage, we link these common-sense qualities in a *support graph* that captures how they mutually support each other in their co-description of a stereotypical idea. From this graph we can estimate positive and negative valence scores for each property and behavior, and default averages for the stereotypes that exhibit them.

Similes and stereotypes share a symbiotic relationship: the former exploit the latter as reference points for an evocative description, while the latter are perpetuated by their constant re-use in similes. Expanding on the approach in Veale (2011), we use two kinds of query for harvesting stereotypes from the web. The first, “as ADJ as a NOUN”, acquires typical adjectival properties for noun concepts; the second, “VERB+ing like a NOUN” and “VERB+ed like a NOUN”, acquires typical verb behaviors. Rather than use a wildcard * in both

positions (ADJ and NOUN, or VERB and NOUN), which yields limited results with a search engine like Google, we generate fully instantiated similes from hypotheses generated via the Google n-grams. Thus, from the 3-gram “a drooling zombie” we generate the query “drooling like a zombie”, and from the 3-gram “a mindless zombie” we generate “as mindless as a zombie”.

Only those similes whose queries retrieve one or more web documents via Google are considered to contain promising associations. But this still gives us over 250,000 web-validated simile associations for our stereotypical model. We quickly filter these candidates manually, to ensure that the contents of the lexicon are of the highest quality. As a result, we obtain rich descriptions for many stereotypical ideas, such as *Baby*, which is described via 163 typical properties and behaviors like *crying*, *drooling* and *guileless*. After this filtering phase, the stereotype lexicon maps 9,479 stereotypes to a set of 7,898 properties and behaviors, to yield more than 75,000 pairings.

We construct the second level of the lexicon by automatically linking these properties and behaviors to each other in a support graph. The intuition here is that properties which reinforce each other in a single description (e.g. “as *lush and green* as a jungle” or “as *hot and humid* as a sauna”) are more likely to have a similar affect than properties which do not support each other. We first gather all Google 3-grams in which a pair of stereotypical properties or behaviors X and Y are linked via co-ordination, as in “*hot and humid*” or “*kicking and screaming*”. A bidirectional link between X and Y is added to the support graph if one or more stereotypes in the lexicon contain both X and Y . If this is not so, we consider whether both descriptors ever reinforce each other in web similes, by posing the web query “*as X and Y as*”. If this query has non-zero hits, we also add a link between X and Y .

Let \mathbf{N} denote this support graph, and $N(p)$ denote the set of neighboring terms to p , that is, the set of properties and behaviors that can mutually support p . Since every edge in \mathbf{N} represents an affective context, we can estimate the likelihood that a property p is ever used in a positive or negative context if we know the positive or negative affect of enough members of $N(p)$. So if we label enough vertices of \mathbf{N} as + or -, we can interpolate a positive/negative valence score for all vertices p in \mathbf{N} .

To do this, we build a reference set $-\mathbf{R}$ of typi-

cally negative words, and a set $+\mathbf{R}$ of typically positive words. Given a few seed members of $-\mathbf{R}$ (such as *sad*, *disgusting*, *evil*, etc.) and a few seed members of $+\mathbf{R}$ (such as *happy*, *wonderful*, etc.), we find many other candidates to add to $+\mathbf{R}$ and $-\mathbf{R}$ by considering neighbors of these seeds in \mathbf{N} . After three iterations in this fashion, we populate $+\mathbf{R}$ and $-\mathbf{R}$ with approx. 2000 words each.

For a property p we can now define $N^+(p)$ and $N^-(p)$ as follows:

$$(1) \quad N^+(p) = N(p) \cap +\mathbf{R}$$

$$(2) \quad N^-(p) = N(p) \cap -\mathbf{R}$$

We can now assign positive and negative valence scores to each vertex p by interpolating from reference values to their neighbors in \mathbf{N} :

$$(3) \quad \text{pos}(p) = \frac{|N^+(p)|}{|N^+(p) \cup N^-(p)|}$$

$$(4) \quad \text{neg}(p) = 1 - \text{pos}(p)$$

If a term S denotes a stereotypical idea and is described via a set of typical properties and behaviors $\text{typical}(S)$ in the lexicon, then:

$$(5) \quad \text{pos}(S) = \frac{\sum_{p \in \text{typical}(S)} \text{pos}(p)}{|\text{typical}(S)|}$$

$$(6) \quad \text{neg}(S) = 1 - \text{pos}(S)$$

Thus, (5) and (6) calculate the mean affect of the properties and behaviors of S , as represented via $\text{typical}(S)$. We can now use (3) and (4) to separate $\text{typical}(S)$ into those elements that are more negative than positive (putting a negative spin on S) and into those that are more positive than negative (putting a positive spin on S):

$$(7) \quad \text{posTypical}(S) = \{p \mid p \in \text{typical}(S) \wedge \text{pos}(p) > 0.5\}$$

$$(8) \quad \text{negTypical}(S) = \{p \mid p \in \text{typical}(S) \wedge \text{neg}(p) > 0.5\}$$

2.1 Evaluation of Stereotypical Affect

In the process of populating $+\mathbf{R}$ and $-\mathbf{R}$, we identify a reference set of 478 positive stereotypes (such as *saint* and *hero*) and 677 negative stereotypes (such as *tyrant* and *monster*). When we use these reference points to test the effectiveness of (5) and (6) – and thus, indirectly, of (3) and (4) and of the

stereotype lexicon itself – we find that **96.7%** of the positive stereotypes in **+R** are correctly assigned a positivity score greater than 0.5 ($pos(S) > neg(S)$) by (5), while **96.2%** of the negative stereotypes in **-R** are correctly assigned a negativity score greater than 0.5 ($neg(S) > pos(S)$) by (6).

3 Expansion/Interpretation of Metaphors

The Google n-grams are a rich source of affective metaphors of the form *Target is Source*, such as “politicians are crooks”, “Apple is a cult”, “racism is a disease” and “Steve Jobs is a god”. Let $src(T)$ denote the set of stereotypes that are commonly used to describe T, where commonality is defined as the presence of the corresponding copula metaphor in the Google n-grams. To find metaphors for proper-named entities like “Bill Gates”, we also analyze n-grams of the form *stereotype First [Middle] Last*, such as “tyrant Adolf Hitler”. Thus:

$$src(racism) = \{problem, disease, joke, sin, poison, crime, ideology, weapon\}$$

$$src(Hitler) = \{monster, criminal, tyrant, idiot, madman, vegetarian, racist, \dots\}$$

We do not try to discriminate literal from non-literal assertions, nor do we even try to define literality. We simply assume each putative metaphor offers a potentially useful perspective on a topic T.

Let $srcTypical(T)$ denote the aggregation of all properties ascribable to T via metaphors in $src(T)$:

$$(9) \quad srcTypical(T) = \bigcup_{M \in src(T)} typical(M)$$

We can also use the $posTypical$ and $negTypical$ variants in (7) and (8) to focus only on metaphors that project positive or negative qualities onto T.

(9) is especially useful when the source S in the metaphor $T is S$ is not a known stereotype in the lexicon, as happens when one describes *Apple as Scientology*. When the set $typical(S)$ is empty, $srcTypical(S)$ may not be, so $srcTypical(S)$ can act as a proxy representation for S in these cases.

The properties and behaviors that are salient to the interpretation of $T is S$ are given by:

$$(10) \quad salient(T, S) = \frac{|srcTypical(T) \cup typical(T)|}{|srcTypical(S) \cup typical(S)|}$$

In the context of $T is S$, the metaphorical stereotype

$M \in src(S) \cup src(T) \cup \{S\}$ is an apt vehicle for T if:

$$(11) \quad apt(M, T, S) = |salient(T, S) \cap typical(M)| > 0$$

and the degree to which M is apt for T is given by:

$$(12) \quad aptness(M, T, S) = \frac{|salient(T, S) \cap typical(M)|}{|typical(M)|}$$

We can construct an interpretation for $T is S$ by considering not just $\{S\}$, but the stereotypes in $src(T)$ that are apt for T in the context of $T is S$, as well as the stereotypes that are commonly used to describe S – that is, $src(S)$ – that are also apt for T:

$$(13) \quad interpretation(T, S) = \{M | M \in src(T) \cup src(S) \cup \{S\} \wedge apt(M, T, S)\}$$

In effect then, the interpretation of $T is S$ is itself a set of apt metaphors for T that expand upon S . The elements $\{M_i\}$ of $interpretation(T, S)$ can now be sorted by $aptness(M_i, T, S)$ to produce a ranked list of interpretations ($M_1, M_2 \dots M_n$). For any interpretation M, the salient features of M are thus:

$$(14) \quad salient(M, T, S) = typical(M) \cap salient(T, S)$$

If $T is S$ is a creative IR query – to find documents that view T as S – then $interpretation(T, S)$ is an expansion of $T is S$ that includes the common metaphors that are consistent with T viewed as S . For any viewpoint M_i , $salient(M_i, T, S)$ is an expansion of M_i that includes all of the qualities that T is likely to exhibit when it behaves like M_i .

4 Metaphor Magnet: A Worked Example

Consider the query “*Google is Microsoft*”, which expresses a need for documents in which Google exhibits qualities typically associated with Microsoft. Now, both *Google* and *Microsoft* are complex concepts, so there are many ways in which they can be considered similar or dissimilar, either in a good or a bad light. However, the most salient aspects of Microsoft will be those that underpin our common metaphors for Microsoft, i.e., stereotypes in $src(Microsoft)$. These metaphors will provide the talking points for the interpretation.

The Google n-grams yield up the following metaphors, 57 for Microsoft and 50 for Google:

$$src(Microsoft) = \{king, master, threat, bully, giant, leader, monopoly, dinosaur \dots\}$$

$src(\text{Google}) = \{\textit{king}, \textit{engine}, \textit{threat}, \textit{brand}, \textit{giant}, \textit{leader}, \textit{celebrity}, \textit{religion} \dots\}$

So the following qualities are aggregated for each:

$srcTypical(\text{Microsoft}) = \{\textit{trusted}, \textit{menacing}, \textit{ruling}, \textit{threatening}, \textit{overbearing}, \textit{admired}, \textit{commanding}, \dots\}$

$srcTypical(\text{Google}) = \{\textit{trusted}, \textit{lurking}, \textit{reigning}, \textit{ruling}, \textit{crowned}, \textit{shining}, \textit{determined}, \textit{admired} \dots\}$

Now, the salient qualities highlighted by the metaphor, namely $salient(\text{Google}, \text{Microsoft})$, are:

$\{\textit{celebrated}, \textit{menacing}, \textit{trusted}, \textit{challenging}, \textit{established}, \textit{threatening}, \textit{admired}, \textit{respected}, \dots\}$

Thus, $interpretation(\text{Google}, \text{Microsoft})$ contains:

$\{\textit{king}, \textit{criminal}, \textit{master}, \textit{leader}, \textit{bully}, \textit{threatening}, \textit{giant}, \textit{threat}, \textit{monopoly}, \textit{pioneer}, \textit{dinosaur}, \dots\}$

Suppose we focus on the metaphorical expansion “Google is king”, since *king* is the most highly ranked element of the interpretation. Now, $salient(\textit{king}, \text{Google}, \text{Microsoft})$ contains:

$\{\textit{celebrated}, \textit{revered}, \textit{admired}, \textit{respected}, \textit{ruling}, \textit{arrogant}, \textit{commanding}, \textit{overbearing}, \textit{reigning}, \dots\}$

These properties and behaviors are already implicit in our perception of Google, insofar as they are salient aspects of the stereotypes to which Google is frequently compared. The metaphor “Google is Microsoft” – and its expansion “Google is king” – simply crystallizes these qualities, from perhaps different comparisons, into a single act of ideation.

Consider the metaphor “Google is -Microsoft”. Since **-Microsoft** is used to impart a negative spin (+ would impart a positive spin), $negTypical$ is here used in place of $typical$ in (9) and (10). Thus:

$srcTypical(\text{-Microsoft}) =$
 $\{\textit{menacing}, \textit{threatening}, \textit{twisted}, \textit{raging}, \textit{feared}, \textit{sinister}, \textit{lurking}, \textit{domineering}, \textit{overbearing}, \dots\}$

$salient(\text{Google}, \text{-Microsoft}) =$
 $\{\textit{menacing}, \textit{bullying}, \textit{roaring}, \textit{dreaded} \dots\}$

Now $interpretation(\text{Google}, \text{-Microsoft})$ becomes:

$\{\textit{criminal}, \textit{giant}, \textit{threat}, \textit{bully}, \textit{victim}, \textit{devil}, \dots\}$

In contrast, $interpretation(\text{Google}, \text{+Microsoft})$ is:

$\{\textit{king}, \textit{master}, \textit{leader}, \textit{pioneer}, \textit{partner}, \dots\}$

More focus is achieved with the simile query “Google is as **-powerful** as Microsoft”. In explicit similes, we need to focus on just a subset of the salient properties, using e.g. this variant of (10):

$\{p \mid p \in salient(\text{Google}, \text{Microsoft}) \cap N(\textit{powerful})$
 $\wedge neg(p) > pos(p)\}$

In this **-powerful** case, the interpretation becomes:

$\{\textit{bully}, \textit{giant}, \textit{devil}, \textit{monopoly}, \textit{dinosaur}, \dots\}$

5 The Metaphor Magnet Web App

Metaphor Magnet is designed to be a lightweight web application that provides both HTML output (for humans) and XML (for client applications). The system allows users to enter queries such as *Google is -Microsoft*, *life is a +game*, *Steve Jobs is Tony Stark*, or even *Rasputin is Karl Rove* (queries are case-sensitive). Each query is expanded into a set of apt metaphors via mappings in the Google n-grams, and each metaphor is expanded into a set of contextually apt qualities. In turn, each quality is then expanded into an IR query that is used to retrieve relevant hits from Google. In effect, the system allows users to interface with a search engine like Google using metaphor and other affective language forms. The demonstration system can be accessed using a standard browser at this URL:

<http://boundinanutshell.com/metaphor-magnet>

Metaphor Magnet can exploit the properties and behaviors of its stock of almost 10,000 stereotypes, and can infer salient qualities for many proper-named entities like *Karl Rove* and *Steve Jobs* using a combination of copula statements from the Google n-grams (e.g., “*Steve Jobs is a visionary*”) and category assignments from Wikipedia.

The interpretation of the simile/query “*Google is as -powerful as Microsoft*” thus highlights a selection of affective viewpoints on the source concept, *Microsoft*, and picks out an apt selection of viewpoints on the target *Google*. *Metaphor Magnet* displays both selections as phrase clouds in which each hyperlinked phrase – a combination of an apt stereotype and a salient quality – is clickable, to yield linguistic evidence for the selection and corresponding web-search results (via a Google gadget). The phrase cloud representing *Microsoft* in this simile is shown in the screenshot of Figure 1, while the phrase cloud for *Google* is shown in Figure 2.

Source Metaphors: Microsoft

menacing:bully
perilous:battle twisted:devil
menacing:tough menacing:lackey
isolated:exception bewildered:victim
overbearing:king bloated:dinosaur
menacing:threat anxious:investor
troubling:question menacing:evil
threatening:enemy arrogant:master
twisted:joke domineering:leader threatening:competitor
menacing:criminal
threatening:problem
menacing:giant arrogant:intellectual
twisted:root conservative:corporation domineering:superior

Figure 1. A screenshot of a phrase cloud for the perspective cast upon the source “Microsoft” by the simile “Google is as –powerful as Microsoft”.

Metaphor Magnet demonstrates the potential utility of affective metaphors in human-computer linguistic interaction, and acts as a web service from which other NL applications can derive a measure of metaphorical competence. When accessed as a service, *Metaphor Magnet* returns either HTML or XML data, via simple *get* requests. For illustrative purposes, each HTML page also provides the URL for the corresponding XML-structured data set.

Acknowledgements

This research was partly supported by the WCU (World Class University) program under the National Research Foundation of Korea (Ministry of Education, Science and Technology of Korea, Project No: R31-30007), and partly funded by Science Foundation Ireland via the Centre for Next Generation Localization (CNGL).

References

Thorsten Brants and Alex Franz. 2006. *Web 1T 5-gram Version 1*. Linguistic Data Consortium.

Target Metaphors: Google

domineering:leader threatening:enemy
menacing:tough threatening:bully
menacing:lackey calculating:investor
menacing:criminal
dreaded:king
domineering:monopoly
complicated:web threatening:competitor arrogant:intellectual
vain:celebrity bloated:dinosaur
twisted:devil
threatening:problem
threatening:threat
isolated:exception twisted:root
dreaded:evil troubling:question

Figure 2. A screenshot of a phrase cloud for the perspective cast upon the target term “Google” by the simile “Google is as –powerful as Microsoft”.

Dan Fass. 1997. Processing Metonymy and Metaphor. *Contemporary Studies in Cognitive Science & Technology*. New York: Ablex.

Hugo Liu, Henry Lieberman and Ted Selker. 2003. A Model of Textual Affect Sensing Using Real-World Knowledge. *Proc. of the 8th international conference on Intelligent user interfaces*, 125-132.

James H. Martin. 1990. *A Computational Model of Metaphor Interpretation*. NY: Academic Press.

Zachary J. Mason. 2004. CorMet: A Computational, Corpus-Based Conventional Metaphor Extraction System, *Computational Linguistics*, 30(1):23-44.

Ekaterina Shutova. 2010. Metaphor Identification Using Verb and Noun Clustering. In *Proc. of the 23rd International Conference on Computational Linguistics*, 1001-1010

Tony Veale. 2011. Creative Language Retrieval. *Creative Language Retrieval: A Robust Hybrid of Information Retrieval and Linguistic Creativity*. In *Proc. of ACL’2011, the 49th Annual Meeting of the Association for Computational Linguistics*.