Metaphor Suggestions based on a Semantic Metaphor Repository

Gerard de Melo
Rutgers University – New Brunswick
Department of Computer Science
Piscataway, NJ, USA
gdm@demelo.org

Abstract

Metaphors are not only remarkably pervasive in both informal and formal text, but reflect fundamental properties of human cognition. This paper presents an algorithmic model that suggests metaphoric means of referring to concepts. For example, given a word such as *government*, the method may propose expressions such as *father*, *nanny*, corresponding to common ways of thinking about the government. To achieve this, the model draws on MetaNet, a manually created repository of conceptual metaphor, in conjunction with lexical resources like FrameNet and WordNet and automated interlinking techniques. These resources are connected and their overall graph structure allows us to propose potential metaphoric means of referring to the given input words, possibly constrained with additional metaphoric seed words, which may be provided as supplementary inputs. The experiments show that this algorithm greatly expands the potential of the original repository for this task by enabling new connections to be drawn.

Keywords: Metaphor, lexical resources, graph structure

1. Introduction

Whenever one says that issues become *clear*, stock markets *go up*, or time is *spent*, language is arguably being used in a non-literal, metaphorical manner, at least with respect to the original senses of the words. Corpus studies have found that metaphorical phenomena are very pervasive even in formal language (Steen et al., 2010; Shutova and Teufel, 2010). Not only is such metaphorical use of language one of the primary means for creative linguistic expression. It has been widely stipulated that our reliance on metaphor is a natural consequence of the way our brains reflect on and reason about the world.

This paper presents a model that can be used to suggest both well-entrenched and novel metaphoric means of referring to a given input word or set of related input words. For example, given a word such as *government*, the method may propose expressions such as *father*, *nanny*, corresponding to ways of thinking about the government.

While metaphor has been studied extensively in linguistics, cognitive science, as well as NLP, the task of automatically suggesting metaphors has received only little attention. Young (1987) relied on simple relational database queries to find related words. Abe et al. (2006) proposed a method that takes a noun and a set of adjectives as input (e.g., *person* and *young*, *innocent*, *fine*) and use corpus topic models to find other nouns with these properties. Veale and Hao (2007) extend this idea to Web-scale knowledge by using the Google search engine to find relevant adjectives describing a noun. Terai and Nakagawa (2010) present an alternative model for this, based on semantic similarity, which considers both adjectives and verbs as relevant noun properties. Approaches of this sort excel at finding novel poetic metaphors, e.g. *hope is like a lightbulb*, as discussed by Terai and Nakagawa (2010). The system discussed in this paper, in contrast, is biased towards finding novel variations of more fundamental conceptual metaphors that shape human thinking.

The model achieves this by constructing a graph to capture relationships between metaphors, words, as well as mental schemas. For this, it draws on the MetaNet respository (Dodge et al., 2015), a manually created database of metaphor. The approach additionally relies on lexical resources such as FrameNet (Baker et al., 1998; Ruppenhofer et al., 2006) and WordNet (Fellbaum, 1998) and automated interlinking techniques. These resources are connected and their overall graph structure allows us to propose potential metaphoric means of referring to the given input words, possibly constrained with additional metaphoric seed words, which may be provided as supplementary inputs.

2. Metaphor and Cognition

Metaphor is often regarded as a process that allows us to think of one thing in terms of another (Lakoff and Johnson, 1980). The following sentences provide examples of **linguistic metaphors**.

(1) Their spirits were *high*.
(2) You lifted me *up* when I was *down*.
(3) That really *raised* their morale.

Although these example sentences involve different metaphorical expressions, it is evident that they share in common the notion that words relating to elevated positions can be invoked in describing an emotional state. It turns out that these three individual instances of linguistic metaphors can be viewed as instantiations of a more general **conceptual metaphor** HAPPY IS UP. Sentence (2) simultaneously also exemplifies the related metaphor SAD IS DOWN, highlighting that even these more general conceptual metaphors can be generalized and related even further to each other.

While metaphor is often used as a creative, if not poetic linguistic device, Lakoff and Johnson (1980) have convincingly argued that metaphor is a more fundamental cognitive process, the primary function of which is in fact understanding. For instance, when speaking of time, most humans conceive of time in terms of a motion along a path.

Linguistic instantiations of this metaphor include sentences like *We have exciting times ahead of us.* or *That happened way back in the 1980s.* Most modern speakers of English also rely on the TIME IS MONEY metaphor, invoking expressions such as *spending time* without any particular conscious realization of this fact. This metaphor is so pervasive that the mere thought about time, for present-day speakers of English, is likely to invoke the metaphor and its entailments (e.g., that time is a limited and valuable resource). Psychological experiments have confirmed that metaphors in natural language covertly influence the way humans reason about things, even when they are not aware of the metaphor (Thibodeau and Boroditsky, 2013).

A metaphor such as TIME IS MONEY thus allows us to make sense of a **target domain** such as TIME in terms of a **source domain** such as MONEY. In conceptual metaphor theory, this is regarded as a directional mapping. Fauconnier and Turner (2008) have argued that the on-line interaction between source and target domain is best thought of as involving a form of conceptual blending of the two domains. Grady et al. (1999) explain that we can think of conventional metaphors, described in terms of source–target mappings, as launching blends, in the sense of providing inputs to and constraints on them. Typically, the source domain is more concrete than target domain. The notion of time is clearly quite abstract, while money has traditionally been more physical, i.e., something you might carry around in your pocket. It is thus argued that metaphor allows the more basic concrete world, e.g., things that can be physically experienced, grasped, or manipulated, to facilitate our understanding of more abstract concepts.

In conceptual metaphor theory, the target and source domains are often thought of as so-called **schemas**. A schema is an established cognitive structure reflecting a particular aspect of the brain?s interactions with the world. One can distinguish the following two types of schemas:

1. Cogs: Lakoff has proposed the notion of cogs to refer to concepts directly grounded in bodily experience. Following Gallese and Lakoff (2005), cogs can be neurally simulated in a secondary area (e.g., the premotor cortex) without active connections to a primary area (e.g., the motor cortex). It is claimed that such simulation can be used for reasoning and that cogs often correspond to the meaning of grammatical constructions, e.g., verb aspect as in *she is about to run* (Narayanan, 1997). Primitive image schemas (e.g., containment, source-path-goal, force dynamics, and orientation schemas) are assumed to be prime examples of cogs ?- see also Dodge and Lakoff (2005).

2. Frames: Frames are taken to include all other concepts, i.e., in particular those that stem from one?s cultural interactions. This is compatible with the notion of frames used in the theory of frame semantics (Fillmore, 1985), which encompasses traditional event representations (Rouces et al., 2015b) but also regards other sorts of entities as being manifested as frames.

As the source domain of a metaphor tends to be more concrete than the target domain, cogs frequently serve as source schemas. For example, the notion of object manipulation (e.g., *grasping*, *holding*) can be applied to target schemas such as the THINKING domain (e.g., *grasping an idea*, *holding views*). However, not all source schemas are cogs.

3. Exploiting the MetaNet Repository

The MetaNet project[1] (Dodge et al., 2015) has been developing a repository of conceptual metaphors that is both human-readable and machine-readable (Hong and Dodge, 2013). The information captured in such a resource is more systematic and formal and thus better-suited for computational processing than previous work on documenting conventionalized metaphors such as the Master Metaphor List (Lakoff and Schwartz, 1991).

The repository[2] directly represents schemas and metaphors. Metaphors are assigned a human-readable label and described in terms of their target and source schemas, as well as entailments, among other things. Schemas can be described in terms of the involved semantic roles, and a description of a conceptual metaphor can explicitly capture the relevant bindings between the roles of the target and source schemas. Additionally, relationships between different schemas and between different conceptual metaphors can be captured. For instance, the ARRIVING schema is connected to more general schemas, all the way to the very abstract MOTION schema, and even further. The TRUST-RELATIONSHIPS ARE BUILDINGS metaphor is a subcase of the more general metaphor RELATIONSHIPS ARE PHYSICAL STRUCTURES.

The repository currently covers four languages (English, Spanish, Russian, and Farsi). Some schemas are linked to related frames defined in the FrameNet project (Baker et al., 1998; Ruppenhofer et al., 2006). However, in practice, these links are often still missing.

4. Metaphor Suggestions

We shall now see how such information about conceptual metaphors can be used in a graph-based framework to suggest linguistic metaphors.

4.1. Overview

The input to the algorithm will typically consist words from the target domain, i.e. the domain that we want to talk about. Our goal is essentially to go from these original words to potential metaphorical expressions from suitable source domains. For instance, we may provide the word *election* as input, and the system will suggest metaphorical words that may be suitable when talking about an election. The outputs could include words such as *headstart*, *race*, *front-runner* from the source domain RACE, as well as words such as *battleground*, *victory*, *allies* from the source domain WAR. Of course, not all proposed words will always be suitable. Given one or more terms from the target domain, the system will produce a ranked list of candidate

[1] `https://metanet.icsi.berkeley.edu/metanet/`

[2] Available online at `https://metaphor.icsi.berkeley.edu/pub/en/`.

terms that may serve as metaphorical expressions for the input terms, chosen from suitable source domains automatically. Optionally, one may also provide relevant seed words from the desired source domain to bias the answers towards source domains of interest.

In its standard form, the algorithm focuses on variations and entailments motivated by existing conceptual metaphors rather than entirely poetic uses that do not bear any relation to common metaphoric cognition processes. The MetaNet repository readily provides a substantial number of such metaphors for a non-trivial number of schemas. However, its lexical coverage is limited. Our algorithm thus adopts a graph-based approach that considers not just the MetaNet repository but also additional lexical resources for much greater lexical coverage.

In conceptual metaphor theory, it is assumed that metaphoric understanding may involve a cascade of activation through a network (Hong and Dodge, 2013). The approach presented in this paper involves setting up a large representational graph structure and then using linear constraints to reflect such activation mechanisms. However, it must be noted that while this method does draw on cognitively inspired resources and mechanisms, no assertion is made that this algorithm comes with any degree of plausibility for human cognition. The goal here is merely to produce useful outputs given the inputs.

4.2. Algorithm

The algorithm operates on a directed graph $G = (V, A)$ with a heterogeneous set of nodes V representing words, schemas, metaphors, and other entities. A directed arc in $(u, v) \in A$ reflects a dependency between node relevance scores, with an arc weight w_{uv} determining to what degree the relevance of u entails the relevance of v. In the following, $V_{\mathrm{m}} \subset V$ shall denote the subset of nodes that represent conceptual metaphors, $A_{\mathrm{t}\to\mathrm{m}} \subseteq A$ to denote the subset of arcs that represent arcs from target schemas to conceptual metaphors, and $A_{\mathrm{m}\to\mathrm{s}} \subseteq A$ to denote the subset of arcs that represent arcs from conceptual metaphors to source schemas. More specific details about the graph are provided later in Section 4.3..

For each node $v \in V$ in the graph, there are two variables to capture their relevance scores: t_v reflects the degree of relevance in the target domain, while s_v reflects the degree of relevance in the source domain.

Formally, the input consists of a set of target node constraints C_{T} of the form $(v, t_{\min}, t_{\max})$ and source node constraints C_{S} of the form $(v, s_{\min}, s_{\max})$. Each target node constraint specifies a desired interval $[t_{\min}, t_{\max}]$ for the target domain relevance t_v of node v. Similarly, each source node constraint specifies a desired interval $[s_{\min}, s_{\max}]$ for the source domain relevance s_v of node v.

In the simplest case, one could simply have a single input word w, and constrain the target relevance t_{v_w} for the corresponding node v_w for w to be 1. To do this, one would provide $C_{\mathrm{T}} = \{(v_w, 0, 1)\}$, $C_{\mathrm{S}} = \emptyset$ as inputs to the algorithm.

Given the graph and the inputs, we seek to find a set of values for t_v, s_v based on the following constrained objective.

$$
\begin{aligned}
&\text{minimize} \\
&\qquad \sum_{v\in V}(t_v + s_v) + c \sum_{(u,v)\in A}(\sigma_{uv} + \tau_{uv}) \\
&\text{subject to} \\
&\quad t_v + \tau_{uv} + \epsilon \geq w_{uv} t_u && \forall\,(u,v) \in A \setminus A_{\mathrm{m}\to\mathrm{s}} \quad (1)\\
&\quad s_v + \sigma_{uv} + \epsilon \geq w_{uv} s_u && \forall\,(u,v) \in A \setminus A_{\mathrm{t}\to\mathrm{m}} \quad (2)\\
&\quad s_v \geq t_v && \forall\, v \in V_{\mathrm{m}} \quad (3)\\
&\quad t_v + s_v \leq 1 && \forall\, v \in V \setminus V_{\mathrm{m}} \quad (4)\\
&\quad t_v \in [t_{\min}, t_{\max}] && \forall\,(v, t_{\min}, t_{\max}) \in C_{\mathrm{T}} \quad (5)\\
&\quad s_v \in [s_{\min}, s_{\max}] && \forall\,(v, s_{\min}, s_{\max}) \in C_{\mathrm{S}} \quad (6)\\
&\quad t_v \in [0,1] && \forall\, v \in V \quad (7)\\
&\quad s_v \in [0,1] && \forall\, v \in V \quad (8)\\
&\quad \tau_{uv} \geq 0 && \forall\,(u,v) \in A \quad (9)\\
&\quad \sigma_{uv} \geq 0 && \forall\,(u,v) \in A \quad (10)
\end{aligned}
$$

These inequalities have a natural interpretation. Constraints (1) and (2) consider arcs (u, v) from nodes u to nodes v and their corresponding arc weights w_{uv}. Any arc (u, v) that is not a metaphor-to-source one indicates to what degree v should acquire target domain relevance from u. Additionally, any arc (u, v) that is not a target-to-metaphor one indicates to what degree v should acquire source domain relevance from u. A small constant ϵ, which is set to 0.1, determines an extra loss of relevance that occurs at every hop along an arc in the graph. Slack variables τ_{uv} and σ_{uv} ensure that there is a feasible solution, but are highly discouraged from becoming non-zero, as the constant c is fixed to a very high value $> 2|V|$ in the objective function.

Conceptual metaphor nodes have a special function in this graph. At any conceptual metaphor node $v \in V_{\mathrm{m}}$, source domain relevance is positively tied to target domain relevance. Thus, it is only here that target domain relevance may be converted into source domain relevance.

At all other nodes, t_v and s_v constrain each other such that $t_v + s_v \leq 1$. The intuition here is that words from the target domain, which are used literally, are undesirable as outputs of the algorithm. The algorithm's output should instead consist of metaphorically relevant words from the source domain. In light of this, the algorithm constrains the two variables t_v, s_v for any v with respect to each other. If a word is fully in the target domain, it is not considered as being in the source domain, and vice versa.

The minimization objective ensures that the relevance and slack variables do not grow arbitrarily for no reason. A number of techniques may be used for this optimization process. Fortunately, in our setting, the number of variables is just $O(|A|)$, as nodes without arcs are irrelevant. This is an important difference from linear programming algorithms that need to keep track of pairwise connections between all nodes (de Melo, 2013). Hence, the current implementation relies on graph pruning and barrier optimization using CPLEX.

Upon obtaining an optimal solution, the set of all words with non-zero s_v form the overall output. If V_{t} is the set of all word nodes in V, a categorical distribution with

$p_v = \frac{s_v}{\sum\limits_{v \in V_t} s_v}$ can be used to draw words from this output at random.

4.3. Graph Construction

The graph is constructed using a number of lexical resources. While the MetaNet repository already covers many of the most fundamental metaphors that shape our thinking, it is only extremely sparsely populated in the number of words attached to schemas and lacks rich knowledge about semantic relationships or associations between words. We shall thus draw on several additional sources to construct the graph.

Many of the nodes represent words, which are here taken to include multi-word expressions. Word nodes are identified by their string form and their language. The latter is necessary to distinguish words from different languages with the same string form.

Note that many of the links between nodes will be symmetric bidirectional ones, reflecting, for instance, synonymy or translation relationships. Such links result in two arcs (u, v), (v, u) that are inverses of each other and share the same arc weight $w_{uv} = w_{vu}$. In this case, the algorithm will aim at obtaining $t_v + \epsilon \geq w_{uv} t_u \geq w_{uv}^2 t_v - w_{uv}\epsilon$ to the extent possible. In other words, it will try to keep the two nodes close to each other, aiming at placing t_u in the interval $\left[w_{uv} t_v - \epsilon, \frac{1}{w_{uv}}(t_v + \epsilon)\right]$ and similarly for s_u, s_v.

From the MetaNet repository itself, we can adopt conceptual metaphors along with their links to source and target schemas and their links to related and subcase metaphors. This includes schemas as well as any links to words (lexical units). We can also incorporate links to FrameNet frames, which some schemas include as part of their metadata. All of this is imported in four different languages (English, Spanish, Russian, Farsi).

FrameNet (Baker et al., 1998; Ruppenhofer et al., 2006) is a lexical resource based on the theory of frame semantics (Fillmore, 1985), and thus highly compatible with conceptual structure of the MetaNet repository. From FrameNet, we can incorporate nodes for frames and lexical units and the corresponding links between them. Additionally, we can include the frame hierarchy, i.e., inheritance relationships between frames in both directions. FrameNet frames are useful as conceptual structures even beyond linguistic annotation. The FrameBase project, for instance, relies on them for knowledge representation (Rouces et al., 2015a; Rouces et al., 2016; Rouces et al., 2017).

From WordNet (Fellbaum, 1998), we can obtain nodes for words, word senses, and sense relationships. Links between words and word senses are incorporated, and hyponym/hypernym, similarity, and derivation links between word senses are taken as well, with manually specified edge type-specific weights. In order to increase the coverage of the graph, schemas without lexical units are automatically linked to WordNet synsets, using the first sense heuristic for disambiguation. Previous work has found that WordNet-like resources include certain types of commonsense knowledge that is highly relevant for capturing entailments in the target and source domains, although the coverage still tends to be limited (Lönneker, 2003).

From VerbNet (Schuler, 2005), we can include verb entries and their links to WordNet senses. The associated SemLink resource (Palmer et al., 2014) provides mappings between verb entries and FrameNet frames and lexical units, which are included as well.

Overall, this process yields a rich graph with numerous connections between words, schemas, and other entities. For example, the graph connects schemas from the MetaNet repository with corresponding FrameNet frames in the following ways: 1) by means of explicitly provided links from the repository, 2) by means of indirect connections through other resources (WordNet, VerbNet), 3) by means of indirect connections via shared terms, and 4) various hybrid forms of indirect connections, often based on semantic relations between words.

4.4. Randomization

While the aforementioned process has used static arc weights in the graph, in some settings, it may be advantageous to integrate a measure of chance into the processing. For this, we can treat each original arc weight $\hat{w}_{uv}$ as a mere hyperparameter and draw the actual arc weight wuv using one of the following two schemes.

Option 1 Draw w_{uv} from $\mathcal{U}(0, \hat{w}_{uv})$ for $(u, v) \in A$: Drawing arc weights from a uniform distribution ensures that the ranking of source relevance scores is perturbed. Hence, one can repeatedly obtain different highest-rated source domain words among those that the original unperturbed graph would provide.

Option 2 Draw $w_{uv} from \mathcal{N}(\hat{w}_{uv}, \sigma^2)$ for arbitrary $(u, v) \in A = V \times V$ and accept if larger than some threshold $w_{\min} \geq 0$: In this alternative scheme, one instead draws arcs using a normal distribution, allowing even previously non-existent arcs to be created with a non-zero probability. Note that the threshold $w_{\min}$ is needed to ensure the non-negativity of arc weights and in practice also to avoid a large quadratic number of arcs.

5. Results

This section describes initial experiments and statistics about the system.

5.1. Graph Creation

The input graph construction was based on June 30, 2013 dumps of the MetaNet repository and of FrameNet, additionally relying on WordNet 3.0, VerbNet 3.2, and SemLink 1.2.2c.

Metaphors	1,125
— English	613
— Spanish	373
— Farsi	82
— Russian	57
Schemas	1,372
— Source	656
— Target	571

Table 1: Input MetaNet Repository

Table 1 provides statistics about the input repository. Note that being a source or target schema is not an inherent property of a schema, but just refers to its involvement in metaphors described in the repository.

Table 2 shows the number of words immediately connected to schemas. It first lists the overall number, and then provides the breakdown by schemas that serve as target schema for some metaphor and schemas that serve as source schema for some metaphor. Clearly, only few terms are activated if one relies on the MetaNet repository in its original form. Once links to FrameNet frames are included, the coverage increases greatly. Adding VerbNet and SemLink does not increase the coverage in a meaningful way, but the main benefit of including these resources is that they serve as segue to WordNet synsets due to the incorporated links. The additional automatically predicted WordNet mappings lead to significant further increases in coverage. Table 2 provides the counts for the overall graph.

These numbers, however, only reflect those words that are somewhat unambiguously connected to a schema via mappings. The power of the approach of this paper lies in the fact that a dense network of semantic or commonsense links in the input graph allows the algorithm to cast a much wider net of possible words. For this, we can rely on the links between conceptual metaphors, links between schemas, links between frames, and WordNet's semantic relations, which overall result in a graph with over 700,000 directed arcs.

Table 3 shows the number of unique words that are connected to a schema at different maximum depth levels (max. number of hops in graph). Thus, with the extended graph, the system is able to select from a much larger pool of candidate words when making suggestions.

5.2. Algorithm Outputs

For instance, running the system using just the original repository for the target word *anger* (with target relevance score 1.0), the algorithm does not find any relevant source domain words. However, using the final graph, it can find several relevant metaphors, including ANGER IS FIRE, ANGER IS HEAT, and ANGER IS INSANITY. The top-ranked output words are *mad*, *steam*, *crazy*, *simmer*, *blow off*, *boil*, *warm*, *hot*, *stew*, which for the most part can quite well be used to describe anger metaphorically. In addition, many hundreds of other candidate words are returned. In lower ranks, one finds words such as *bake*, *kick*, *hammer*, *poison*, *smash*, *microwave*.

At the input side, instead of *anger*, we can also enter alternative target domain words such as *enrage*, *fury*, and so on, and the algorithm still finds relevant source domain words.

From an accuracy perspective, since the output results from graph links that overwhelmingly have been manually created, the presented words are clearly connected to the relevant domains. Still, the method simply provides relevant source domain words in a rather open-ended manner, and some will obviously be suitable, while in other cases it may still be challenging for writers to come up with a suitable way of employing the suggested terms in a sentence such that the proposed metaphorical interpretation comes to bear. Many suggestions may thus prove unsuitable. To address this, among the output words, one could apply additional filtering to select words that highlight one particular property of the target domain using the technique suggested by Veale and Hao (2007). This involves retaining from the set of proposed source domain words only those that satisfy the constraint of having significant corpus or Web frequency occurrence counts for patterns such as for instance *as* ⟨*property*⟩ *a* ⟨*word*⟩ (e.g., *as innocent as a child*) or ⟨*property*⟩ ⟨*word*⟩ (e.g., *innocent child*).

5.3. Cross-Lingual Applicability

In the above example, despite the use of the English input word *anger*, the algorithm also finds pertinent words in other languages such as the Spanish *calentar* and *arrojar*. This is possible because the MetaNet data includes cross-lingual connections of frames across languages. It is trivial to extend this even further by incorporating further multilingual terms for FrameNet frames (Čulo and de Melo, 2012) and WordNet synsets (de Melo and Weikum, 2009; de Melo and Weikum, 2014), or on multilingual word vectors (de Melo, 2015).

Obviously, there are also important caveats here. The MetaNet database currently described conventional metaphor in four languages, and any cross-lingual connection that the algorithm emits will need to be made via conventional metaphor links in one of those languages. In the literature, there have been studies about how metaphors compare across language boundaries (Gutiérrez et al., 2016). Clearly, many metaphors are highly language-specific, and thus when the algorithm is applied cross-lingually, some of the emitted output candidates are likely to be unsuitable.

However, in many cases, they turn out to be appropriate. For one, this may stem from similarities in metaphorical language use across related languages. For instance, in many Western languages, the word *transparency*, which in its original sense refers to the property of allowing the transmission of light through an object, is also used to refer to public evaluability and accountability. Additionally, this may also stem from a broadly shared experiential basis. For example, the connection between anger and heat derives from human biology. Empirically, Tsvetkov et al. (2014) found that they were able to apply models trained to detect English linguistic metaphors also to the task of detecting linguistic metaphors in other languages, some of which are not phylogenetically close (specifically, they considered Spanish, Farsi, and Russian).

6. Related Work

6.1. Metaphor Detection

Numerous papers have studied the task of automatically identifying metaphoric expressions in text. Many systems aim at achieving this by detecting violations of selectional preference restrictions (Fass, 1991; Shutova et al., 2010). For instance, the verb *to kill* usually applies to living beings, so when it is found in contexts such as *my process got killed*, it is quite likely that the word is being used metaphorically. Some approaches have additionally relied on lexical resources such as FrameNet (Gedigian et al., 2006) and HowNet (Tang et al., 2010) to increase the qual-

	Schemas	Target Schemas	Source Schemas
Original Repository	803	258	513
+ Direct FrameNet Links	5,349	2,283	2,776
+ VerbNet/SemLink	5,368	2,290	2,783
+ WordNet	6,302	2,626	2,849
+ WordNet mappings	6,789	3,304	3,497

Table 2: Word–Schema Connections, where subsequent rows show the counts as additional resources are added (including all previously mentioned ones).

Depth	1	2	3	4	5	6	7	8
Original Repository	742	742	742	742	742	742	742	742
Final Graph	742	1,202	9,024	29,246	62,411	94,065	117,372	125,038

Table 3: Words with Indirect Schema Connections.

ity. However, many such systems are brittle as they rely on rather small amounts of manually provided data.

The approach by Shutova et al. (2015) draws on a massive collection of images and videos with associated tags to draw inferences about which predicate-argument pairs are more likely to be concrete. The underlying assumption is that such multimodal data is more closely grounded in the real world and hence linguistic descriptions are more likely to be concrete and literal (e.g., *cutting hair*) as opposed to more metaphorical (e.g., *cutting costs*), which is often predominant in newswire text.

Only few systems have attempted to go beyond identifying individual linguistic metaphors towards recognizing more general conventionalized conceptual metaphors. The CorMet system (Mason, 2004) attempts to automatically infer metaphor mappings from a corpus by studying systematic differences in verb selectional preferences between domains. Such techniques could be used to extend the number of metaphors in our graph.

6.2. Metaphor Analysis

The Metaphor Magnet system (Veale and Li, 2012) addresses the complementary task of metaphor interpretation, providing a list of attributes that explain what qualities a given source domain shares with a target domain that lend the metaphor its strength. The system takes a target and source domain as input (e.g., LOVE IS A DRUG), and then uses Web corpus frequencies to highlight salient features of the source domain that are shared with the target domain. For instance, for the given example, it highlights *healing*, *satisfying*, and *intoxicating*, providing possible interpretations of how the target concept LOVE is modified when metaphors from the DRUG source domain are invoked.

Shutova et al. (2012) present an unsupervised approach for finding literal paraphrases for a given metaphorical expression. Paraphrases are identified using distributional semantics captured in a vector space model. Selectional preference statistics are then used to select literal expressions among the retrieved paraphrases.

6.3. Metaphor Suggestions

In contrast, the task of automatically suggesting metaphors has received only little attention. Young (1987) uses simple relational database queries to find related words. Abe et al. (2006) take a noun and a set of adjectives as input (e.g., *person* and *young*, *innocent*, *fine*). They then use corpus topic models to find other nouns with these properties. In their experiments, their system mainly emits related words such as *grandchild*, but their intuition is that such a framework could also emit more metaphorical ones such as *puppy*. Veale and Hao (2007) extend this idea to Web-scale knowledge by relying on the Google search engine to find relevant adjectives describing a noun. Terai and Nakagawa (2010) present an alternative model for this same task, based on semantic similarity, which considers both adjectives and verbs as relevant noun properties.

Approaches of this sort excel at finding novel poetic metaphors, e.g., *hope is like a lightbulb* as discussed by Terai and Nakagawa (2010). The system proposed in this paper, in contrast, both in its algorithm and in the kinds of resources we draw on, is biased towards finding variations of fundamental conceptual metaphors that shape human thinking.

7. Conclusion

This paper presents a framework for metaphor suggestion. The approach is centered on the MetaNet semantic repository of conceptual metaphors, which captures various cognitive phenomena, including the connections between target and source domains in common conceptual metaphors.

The approach in this paper involves connecting this information with several other lexical resources in a graph and then optimizing a constrained objective to determine relevance scores for potential source domain words. The algorithm can flexibly integrate additional information sources into its graph and incorporate pre-existing information about relevant source domain words as additional constraints.

The results show that this graph-based approach has a significantly higher coverage than the original repository, to the extent that even words in completely unrelated languages can be processed cross-lingually. Overall, this paves the way for novel applications that use language in a more creative, flexible way.

Acknowledgements

This research was supported in part by the DAAD and by the DARPA SocialSim program.

8. Bibliographical References

Abe, K., Sakamoto, K., and Nakagawa, M. (2006). A computational model of metaphor generation process. In *Proceedings of the 28th Annual Meeting of the Cognitive Science Society*, page 937?942.

Baker, C. F., Fillmore, C. J., and Lowe, J. B. (1998). The Berkeley FrameNet Project. ICCL '98, pages 86–90.

de Melo, G. and Weikum, G. (2009). Towards a universal wordnet by learning from combined evidence. In *Proceedings of the 18th ACM Conference on Information and Knowledge Management (CIKM 2009)*, pages 513–522, New York, NY, USA. ACM.

de Melo, G. and Weikum, G. (2014). Taxonomic data integration from multilingual Wikipedia editions. *Knowledge and Information Systems*, 39(1):1–39, April.

de Melo, G. (2013). Not quite the same: Identity constraints for the Web of Linked Data. In Marie desJardins et al., editors, *Proceedings of the 27th AAAI Conference on Artificial Intelligence (AAAI 2013)*, pages 1092–1098, Menlo Park, CA, USA. AAAI Press.

de Melo, G. (2015). Wiktionary-based word embeddings. In *Proceedings of MT Summit XV*.

Dodge, E. and Lakoff, G. (2005). Image schemas: From linguistic analysis to neural grounding. In Beate Hampe et al., editors, *From Perception to Meaning: Image Schemas in Cognitive Linguistics*, pages 57–91.

Dodge, E., Hong, J., and Stickles, E. (2015). Metanet: Deep semantic automatic metaphor analysis. In *Proceedings of the Third Workshop on Metaphor in NLP*, pages 40–49, Denver, Colorado, June. Association for Computational Linguistics.

Fass, D. (1991). Met*: A method for discriminating metonymy and metaphor by computer. *Comput. Linguist.*, 17(1):49–90, March.

Fauconnier, G. and Turner, M., (2008). *Rethinking Metaphor*, pages 53–66. Cambridge University Press, September.

Fellbaum, C. (1998). *WordNet: An Electronic Lexical Database*. MIT Press.

Fillmore, C. J. (1985). Frames and the semantics of understanding. *Quaderni di Semantica*, 6(2):222–254.

Gallese, V. and Lakoff, G. (2005). The brain's concepts: The role of the sensory-motor system in conceptual knowledge. *Cognitive Neuropsychology*, 22(3-4):455–479, May.

Gedigian, M., Bryant, J., Narayanan, S., and Ciric, B. (2006). Catching metaphors. In *Proceedings of the Third Workshop on Scalable Natural Language Understanding*, ScaNaLU '06, pages 41–48, Stroudsburg, PA, USA. Association for Computational Linguistics.

Grady, J., Oakley, T., and Coulson, S. (1999). Blending and metaphor. In Raymond W. Gibbs Jr. et al., editors, *Metaphor in Cognitive Linguistics: Selected papers from the 5th International Cognitive Linguistics Conference*, Current Issues in Linguistic Theory 175, page 101ff.

Gutiérrez, E. D., Shutova, E., Lichtenstein, P., de Melo, G., and Gilardi, L. (2016). Detecting cross-cultural differences using a multilingual topic model. *Transactions of the Association for Computational Linguistics (TACL)*, 4:47–60.

Hong, Jisup, E. S. and Dodge, E. (2013). The MetaNet metaphor repository: Formalized representation and analysis of conceptual metaphor networks. In *Proceedings of the 12th International Cognitive Linguistics Conference*.

Lakoff, G. and Johnson, M. (1980). *Metaphors we Live by*. University of Chicago Press, Chicago.

Lakoff, George, J. E. and Schwartz, A. (1991). The master metaphor list (draft 2nd edition).

Lönneker, B. (2003). Is there a way to represent metaphors in wordnets? Insights from the Hamburg metaphor database. In *Proceedings of the ACL 2003 Workshop on the Lexicon and Figurative Language*, pages 18?–27. Association for Computational Linguistics.

Mason, Z. J. (2004). Cormet: A computational, corpus-based conventional metaphor extraction system. *Computational Linguistics*, 30:23–44.

Narayanan, S. S. (1997). *KARMA: Knowledge-based Active Representations for Metaphor and Aspect*. University of California, Berkeley.

Palmer, M., Bonial, C., and McCarthy, D. (2014). Semlink+: Framenet, verbnet and event ontologies. In *Proceedings of Frame Semantics in NLP: A Workshop in Honor of Chuck Fillmore (1929-2014)*, pages 13–17, Baltimore, MD, USA, June. Association for Computational Linguistics.

Rouces, J., de Melo, G., and Hose, K. (2015a). FrameBase: Representing n-ary relations using semantic frames. In *Proceedings of ESWC 2015*, pages 505–521.

Rouces, J., de Melo, G., and Hose, K. (2015b). Representing specialized events with FrameBase. In *Proceedings of the 4th International Workshop on Detection, Representation, and Exploitation of Events in the Semantic Web (DeRiVE 2015) at ESWC 2015*.

Rouces, J., de Melo, G., and Hose, K. (2016). Heuristics for connecting heterogeneous knowledge via FrameBase. In *Proceedings of ESWC 2016*, Lecture Notes in Computer Science. Springer.

Rouces, J., de Melo, G., and Hose, K. (2017). FrameBase: Enabling integration of heterogeneous knowledge. *Semantic Web*, 8(6):817–850, August.

Ruppenhofer, J., Ellsworth, M., Petruck, M. R., Johnson, C. R., and Scheffczyk, J. (2006). *FrameNet II: Extended Theory and Practice*. International Computer Science Institute, Berkeley, California. Distributed with the FrameNet data.

Schuler, K. K. (2005). *Verbnet: A Broad-coverage, Comprehensive Verb Lexicon*. Ph.D. thesis, Philadelphia, PA, USA. AAI3179808.

Shutova, E. and Teufel, S. (2010). Metaphor corpus annotated for source - target domain mappings. In Nicoletta Calzolari (Conference Chair), et al., editors, *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*, Valletta, Malta, may. European Language Resources Association (ELRA).

Shutova, E., Sun, L., and Korhonen, A. (2010). Metaphor identification using verb and noun clustering. In *Proceedings of the 23rd International Conference on Computational Linguistics*, COLING '10, pages 1002–1010, Stroudsburg, PA, USA. Association for Computational Linguistics.

Shutova, E., van de Cruys, T., and Korhonen, A. (2012). Unsupervised metaphor paraphrasing using a vector space model. In *Proceedings of COLING 2012: Posters*, pages 1121–1130, Mumbai, India, December. The COLING 2012 Organizing Committee.

Shutova, E., Tandon, N., and de Melo, G. (2015). Perceptually grounded selectional preferences. In *Proceedings of ACL 2015*, pages 950–960.

Steen, G. J., Dorst, A. G., Herrmann, J. B., Kaal, A., Krennmayr, T., and Pasma, T. (2010). *A Method for Linguistic Metaphor Identification*. John Benjamins Publishing.

Tang, X., Qu, W., Chen, X., and Yu, S. (2010). Beyond selectional preference: Metaphor recognition with semantic relation patterns. *Int. J. of Asian Lang. Proc.*, 20(4):141–156.

Terai, A. and Nakagawa, M. (2010). A computational system of metaphor generation with evaluation mechanism. In *Proceedings of the 20th International Conference on Artificial Neural Networks: Part II*, ICANN'10, pages 142–147, Berlin, Heidelberg. Springer-Verlag.

Thibodeau, P. H. and Boroditsky, L. (2013). Natural language metaphors covertly influence reasoning. *PLOS ONE*, 8(1):1–7, 01.

Tsvetkov, Y., Boytsov, L., Gershman, A., Nyberg, E., and Dyer, C. (2014). Metaphor detection with cross-lingual model transfer. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 248–258, Baltimore, Maryland, June. Association for Computational Linguistics.

Čulo, O. and de Melo, G. (2012). Source-Path-Goal: Investigating the cross-linguistic potential of frame-semantic text analysis. *it - Information Technology*, 54.

Veale, T. and Hao, Y. (2007). Comprehending and generating apt metaphors: a web-driven, case- based approach to figurative language. In *Proceedings of the 22nd national conference on Artificial intelligence (AAAI 2007) – Volume 2*, page 1471?14762. AAAI Press.

Veale, T. and Li, G. (2012). Specifying viewpoint and information need with affective metaphors: A system demonstration of the metaphor magnet web app/service. In *Proceedings of the ACL 2012 System Demonstrations*, ACL '12, pages 7–12, Stroudsburg, PA, USA. Association for Computational Linguistics.

Young, L. F. (1987). The metaphor machine: A database method for creativity support. *Decision Support Systems*, 3(4):309–317.