MetaNet: Deep semantic automatic metaphor analysis

Ellen Dodge and Jisup Hong
International Computer Science Institute
1947 Center Street Ste. 600
Berkeley, CA 94704
edodge@icsi.berkeley.edu
jhong@icsi.berkeley.edu

Elise Stickles
University of California, Berkeley
1203 Dwinelle Hall
Berkeley, CA 94720-2650
elstickles@berkeley.edu

Abstract

This paper describes a system that makes use of a repository of formalized frames and metaphors to automatically detect, categorize, and analyze expressions of metaphor in corpora. The output of this system can be used as a basis for making further refinements to the system, as well as supporting deep semantic analysis of metaphor expressions in corpora. This in turn provides a way to ground and test empirical conceptual metaphor theory, as well as serving as a means to gain insights into the ways conceptual metaphors are expressed in language.

1 Introduction

Recognition of the ubiquity of metaphor in language has led to increased interest in automatic identification of metaphoric expressions in language. Typical approaches to metaphor analysis in linguistics comprise (a) theory-driven introspective top-down methods, and (b) bottom-up corpus approaches with an emphasis on analyzing how metaphors are used in discourse. Computational approaches tend to focus on the task of metaphor detection (i.e. determining whether a particular expression metaphoric or not) rather than attempting to identify and analyze which conceptual metaphors are being expressed.

The MetaNet approach described here bridges the two linguistic methodologies above by providing (a) a linguist-friendly interface for formally representing conceptual metaphor theoretic analyses and principles, and (b) an automatic metaphor detection system that applies those analyses and principles to identify metaphoric expressions with-

in large-scale corpora. What results is an integrated system that connects the output of the metaphor detection process to rich information that enables further semantic analysis. This serves as a means for advancing and refining conceptual metaphor theory, and increasing our understanding of how metaphors are used in language.

1.1 Related work

Our work addresses two important criticisms that have been directed toward much previous linguistic work in conceptual metaphor analysis. One issue is that such analyses are often idiosyncratic, with methods of analysis and representations of metaphor varying from analyst to analyst; to address this, metaphor study needs rigorous methodological analyses that can be replicated (Pragglejaz 2007, Kövecses 2011). Another criticism is that metaphor theorists often take a top-down approach that relies on analysis of data gathered from introspection; this tends to limit discovery of new metaphors, and focus analysis on those metaphors the analyst has already identified or vetted from the literature. This contrasts with a bottom-up, corpus-based approach espoused by Stefanowitsch (2006), Deignan (2005), Martin (2006), and others, who argue that identifying as many metaphors as possible in a corpus leads to a clearer picture of the full inventory of metaphoric expressions, as well as providing a measure of their relative frequency of use. Furthermore, such a method can serve to verify theories based on previously-identified metaphors, as well as aiding the discovery of previously-unidentified metaphors.

Various computational approaches have been applied to the task of metaphor detection. Among the first systems, Fass (1991) used selectional pref-

erence violations as a cue for nonliteralness, and then relied on comparisons to a knowledge base for further disambiguation. Gedigian et al. (2006)'s system achieved high accuracy at classifying verbs in PropBank annotated texts, though only from a limited domain for a small range of source domain frames, using features consisting of the verb plus its argument filler types expressed as WordNet synsets. In a larger-scale system, Shutova et al. (2010) used unsupervised clustering methods to create noun and verb clusters that represent target and source concepts, respectively. Mappings between them, established by metaphoric seed expressions, were then used to generate novel target-source expressions. Similarly, Mohler et al. (2013)'s system builds semantic signatures that map text to areas in a multidimensional conceptual space and represent associations between concepts. These are compared to known metaphoric ones to detect novel metaphoric expressions. Other systems, such as Turney et al. (2011) and Tsvetkov et al. (2014) determine metaphoricity based on lexical features such as abstractness/concreteness, imageability, and supersenses derived from WordNet.

Our approach to metaphor detection differs from previous approaches in its deliberate dependence on formalization of a particular theory of metaphor and the correctness and completeness of a conceptual metaphor repository expressed in that formalism. By design, we expect the system to succeed at identifying metaphoric expressions to the extent that the formalism and the repository are consistent and correct. The approach thus integrates top-down linguistic and bottom-up computational approaches to metaphor identification and annotation, combining the strengths of each. A significant outcome is that in addition to detecting metaphors in text, our system also yields semantic information about each of these expressions, including identification of source and target domains and links to underlying conceptual metaphors.

1.2 System overview

There are three key components in our system: (1) a repository of formalized metaphors, frames, metaphor constructions, and metaphoric relational patterns; (2) an automated metaphor extraction system that utilizes information from the repository to identify expressions of metaphor in text and annotate them for additional semantic information; and

(3) computational tools to evaluate, analyze, and visualize the extracted metaphor data. Together, these are used to form a 'cycle' of analysis, in which analysis of extracted data serves as a means to refine and expand the repository, which in turn improves metaphor extraction results. The system is currently in development for analysis of American English, Mexican Spanish, and Russian.

2 Improvements to Extraction Based on Formalization of Metaphor Theory

Since Lakoff and Johnson's first book on conceptual metaphor theory (1980), the field has come to recognize the hierarchical and taxonomic nature of metaphors and the concepts that comprise their source and target domains. For example, consider the phrases *poverty infects society* and *crime is plaguing the nation*, which instantiate the specific metaphors POVERTY IS A DISEASE and CRIME IS A DISEASE, respectively. However, they inherit much of their semantics from a more general metaphor, SOCIAL PROBLEMS ARE AFFLICTIONS; this in turn inherits from a yet more general metaphor, NEGATIVELY EVALUED CONDITIONS ARE PHYSICALLY HARMFUL, as shown in Figure 1.
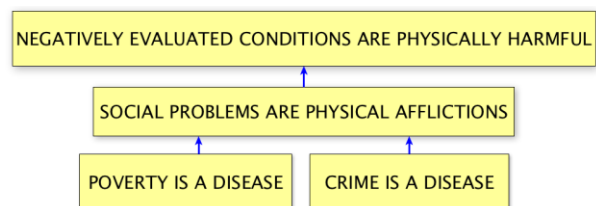

Figure 1. Metaphor inheritance network

It is also clear that the semantic domains of these metaphors are themselves hierarchically related: poverty and crime are social problems, which are negative conditions; meanwhile, disease is a type of physical affliction, which in turn is something that causes physical harm. These domains are represented in our system as semantic frames (Fillmore 1976) similar to those instantiated in FrameNet (Ruppenhofer et al., 2010), which constitute conceptual gestalts that describe particular situations or events along with their participants and other basic conceptual structures. By developing a system that formally represents these structures and relations in an ontology of frames and metaphors, we enable the possibility of a rigorous

system of representation that can be computationally implemented and leveraged for improved metaphor detection.

2.1 Repository of metaphors and frames

The MetaNet project represents an effort to formally represent and categorize metaphors and frames that comprise their source and target domains, and relations between them. Frames are coherent, conceptual gestalts organized in a hierarchical structure. These range from experiential universal structures such as Motion Along a Path and Verticality, to more specific situations such as Physical Restraints and Disease; they also include less physically concrete culturally-based frames like Poverty and Corruption. More-specific frames incorporate the semantics and internal structure of the more-general frames they inherit from, forming a complex network of related concepts. Relations between frames define how elements of a parent frame are incorporated by the child frame. For instance, the 'subcase of' relation indicates that the child fully inherits and elaborates the structure of the parent frame. In addition to traditional ontological relations, we also include relations specific to frame semantics and metaphor theory. A fragment of this network is illustrated in Figure 2.
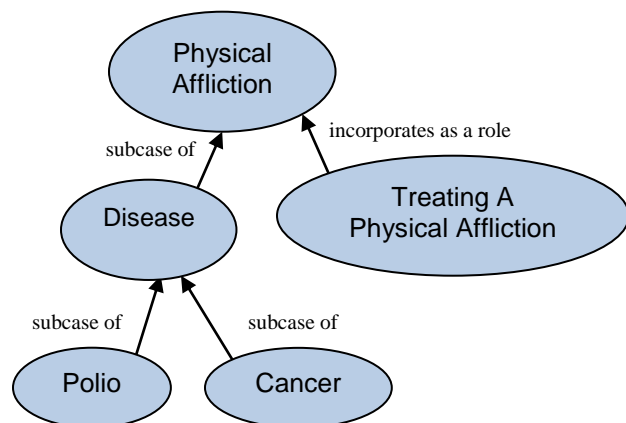

Figure 2. Non-metaphoric frame network pattern

Sub-networks like the group of Physical Affliction frames in Figure 2 are further grouped together to form families of frames, which define collections of broader, but still coherent, conceptual domains.

In addition to relations between frames, structure within frames is also represented in the repository. This includes such elements as participant roles, aspectual and causal structures, relationships between roles, and lexical units that evoke the frame. Figure 3 illustrates partial frame representations of the Poverty and Disease frames. Internal frame structure not only enables improved analysis by requiring the analyst to consider the details of each frame, but also provides additional information in metaphor detection. As the detection system identifies the frames that contribute to the identified metaphor, the detailed semantics of those concepts can be accessed via these frame entries.

Metaphors are essentially representations of mappings between frames. The structure of the source domain frame maps onto the structure of the target domain frame (Figure 3); hence, in POVERTY IS A DISEASE, the *impoverished people* of the Poverty frame are understood as the *patient* experiencing the disease in the Disease frame. Specifically, the roles of the Disease frame map onto their counterparts of the Poverty frame.
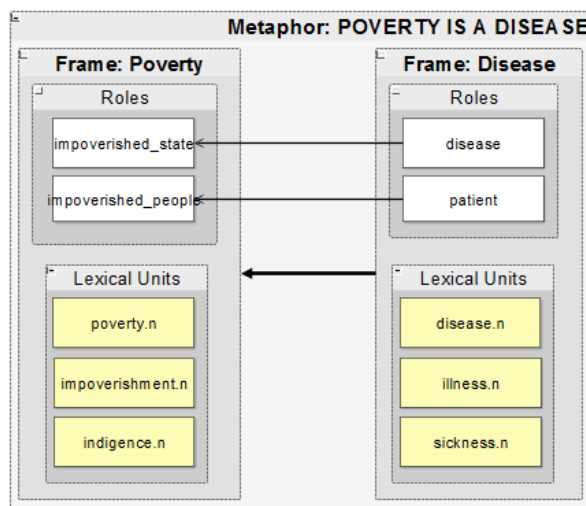

Figure 3. Metaphor structure

Furthermore, just as frame-frame relations define how one frame incorporates the semantics of another, metaphor-metaphor relations define the hierarchy of metaphors (Figure 1). The result is a complex, lattice-like inheritance network of concepts and metaphors.

The computational architecture used for this purpose is a Wiki, based on the Semantic Mediawiki format (Krötzsch et al. 2006). Linguists trained in conceptual metaphor theory create frame and metaphor entries as individual pages, specify-

ing for each metaphor its source and target domain frames, role-to-role mappings, and relations between that metaphor and others in a network. Initially the repository was seeded from metaphors previously identified in the past 30 years of metaphor literature; including comprehensive analysis of primary metaphors provides broad coverage of conceptual metaphors that are applicable to many target domains. For example, the metaphor MORE IS UP can be found in such varied expressions as *prices skyrocketed, she had high hopes,* and *studying boosted his GPA.* Following this initial stage, additional metaphors are added as analysts find them via focused study of particular target domains; however, the system can identify metaphoric language even in the absence of specific metaphors by utilizing the frame network to find more general-level metaphors, as will be shown in section 3.2.

2.2 Metaphor constructions

Previous research has demonstrated that metaphors tend to be expressed in certain regular constructional patterns (Croft 2002; Sullivan 2007, 2013). For example, the noun phrase *poverty trap* has the source domain lexeme *trap* modified by the target domain lexeme *poverty*; these noun-noun metaphor constructions consistently appear in this dependency relation. In contrast, the reverse construction with the source modifying the target is not observed in corpus data (Sullivan 2013). Building on this research, our project has defined a set of grammatical constructions that represent several different types of frequently occurring patterns. In each, the construction specifies which constructional element evokes the source domain, and which evokes the target domain (Table 1). The source term fills different slots in these various patterns; the source may appear not only as a verb, but also as a noun or adjective. As described in section 3.1, these constructions enable us to link a broad range of potentially metaphoric expressions (e.g. lexical units, expressed in some constructional pattern) to the frames and conceptual metaphors in our repository. However, while this process eliminates much data that is *not* metaphoric, it does not positively identify expressions that *are* metaphoric: for example, the same noun-noun pattern that identifies *poverty trap* also returns the literal

expression *bear trap*. Hence, disambiguating between these two types of expressions requires a second step of metaphoricity evaluation.

Constructional pattern	Examples
T-subj_S-verb	*poverty infects*
T-subj_S-verb-conj	*poverty infects and maims*
T-subj-conj_S-verb	*homelessness and poverty infect*
S-verb_T-dobj	*escape poverty*
S-verb_T-dobj-conj	*escape despair and poverty*
S-verb_Prep_T-noun	*slide into poverty / pull up out of poverty*
S-noun_of_T-noun	*trap of poverty*
T-noun_poss_S-noun	*poverty's undertow*
S-noun_prep_T-noun	*path to poverty*
T-noun_mod_S-noun	*poverty trap*
S-adj_mod_T-noun	*burdensome poverty*
T-noun_cop_S-noun-adj	*poverty is a disease / poverty is burdensome*

Table 1. Constructional Patterns

3 Metaphor Extraction and Identification

Our automatic metaphor identification system divides into two main phases. In the first, we use a set of manually defined metaphoric constructional patterns to identify candidate expressions with explicitly realized potential target and source elements. In the second, we identify the frames that are evoked by these elements, and use our conceptual network of frames and metaphors, along with a set of patterns of relationships between nodes in the network, to determine the likelihood of a candidate expression being metaphoric. These phases are presented in detail below.

3.1 Matching constructional patterns

The first step in the process is to identify potentially metaphoric expressions in the corpus; the system can search for metaphors for a particular target domain family, metaphors that make use of a particular source domain family, or simply all the metaphoric expressions in the data. This search is performed by making use of the metaphoric constructional patterns as described in section 2.2. They are represented as SPARQL queries that

43

specify document structural constraints, including grammatical constraints. To search texts for constructional matches, we construct Resource Description Framework (RDF) models of each sentence in terms of an ontology defined in the Web Ontology Language (OWL). The ontology defines the classes Document, Sentence, and Word, and properties, some of which are shown in Table 2 and with their domain and range.

	Domain	Range
inDocument	Sentence	Document
inSentence	Word	Sentence
follows	Word	Word
precedes	Word	Word
dep	Word	Word
hasIdx	Sentence,Word	integer
hasForm	Word	string
hasLemma	Word	string
hasPOS	Word	string

Table 2. Document properties

The resulting RDF representation of the input text constitutes a graph structure in which words, sentences, and documents are nodes in the graph. Properties serve to characterize nodes in terms of string or integer information, such as form, lemma, part of speech (POS), or position, as well as in terms of a node's relation to other nodes. Such relations, with respect to Word nodes, include ordering relations and grammatical dependency relations. While Table 2 shows only the root of the dependency relations, dep, the ontology includes a grammatical relations hierarchy that represents a merger of the hierarchies used by the Stanford (De Marneffe et al., 2006), RASP (Briscoe et al., 2006), and Spanish Freeling (Lloberes et al., 2010) dependency parsers.

Generating this representation requires that NLP tools such as lemmatizers, POS taggers, and dependency parsers be available for the language in question. Because dependency parsing is the most computationally expensive step in this process, in cases where the metaphor extraction is being run only for certain target or source domains, a preprocessing step identifies sentences of interest based on the presence of a word from those domains.

In order to search large volumes of text using SPARQL constructional pattern queries, documents are converted to RDF and uploaded to an OpenRDF Sesame triplestore. Constructional pattern matching queries are run in succession over each document, with queries written so that each match result includes a sentence index, as well as the lemma and word index of the potentially metaphoric lexical elements. Documents are processed in parallel to the extent possible given hardware limitations. With six compute servers each providing 16 cores and running a local triplestore, we were able to run metaphor detection on a preprocessed 500 million word subset of the English Gigaword corpus (Graff & Cieri 2003) in 6 hours.

3.2 Evaluating metaphoricity

The preceding phase of the metaphor extractor returns pairs of words that are related to each other by a constructional pattern where one word may be the source domain of a metaphor, and the other word may be the target domain of that metaphor. While the constructional patterns represent a necessary constraint on metaphoric expression, they are not sufficient to guarantee metaphoricity. Hence, the second phase of metaphor detection makes use of the network of frames and metaphors instantiated in the metaphor repository in order to disambiguate between metaphoric and non-metaphoric expressions in the pool of candidates.

The content of the wiki repository (as described in Section 2.1) is converted to an RDF representation, also in terms of an OWL-defined ontology, and loaded into a triplestore repository. Entries for candidate lexical items in the repository are associated with the frames that they evoke; if the lexical items for English are not already present in the system, FrameNet (https://framenet.icsi.berkeley.edu), WordNet (https://wordnet.princeton.edu), and Wiktionary (https://www.wiktionary.org) data are used to expand the search for the most relevant frame present in the system. After these frames are identified, the system performs searches through the network to determine how the frames are related to one another. If a repository search of the chain of relations that connect the frames includes codified metaphoric mappings, the extractor recognizes the candidate expression as metaphoric.

The likelihood that an expression is metaphoric is determined by attempting to match the relational network between the two frames against a set of pre-defined patterns, which are expressed in

SPARQL and stored in the Semantic MediaWiki, along with the constructional patterns. These patterns fall into two basic types.

The first type are relational configurations that constitute negative evidence for metaphoricity—i.e. they suggest that the expression is not metaphoric. For example, if the potential source and target lexical units evoke the same frame, the system could conclude that the expression is not metaphoric. Similarly, the system can also disregard cases where the frames are too closely related at some point in the network, e.g., if the candidate target lemma evokes a frame that is inherited by the candidate source frame. For example, in the phrases *to cure a disease* and *to cure polio*, *cure* evokes the Treating a Physical Affliction frame, in which one of the roles is the physical affliction being treated. The potential target lemmas *disease* and *polio* evoke the Disease and Polio frames, which inherit from Physical Affliction as shown in Figure 2. The constructional pattern matching phase of the system would identify the expressions as candidates, with *cure* as the source word in both cases, and with *disease* and *polio* as the target words for each phrase. The system, however, is able to exclude these on the basis of a rule that defines a known non-metaphoric network pattern, `TargetIsRoleInSource`, where the frame evoked by the potential target term either directly or recursively inherits from a frame that is incorporated as a role into the frame evoked by the potential source term.

The second type of network relational patterns are a set of rules that constitute positive evidence for metaphoricity. For example, if the two lemmas evoke frames that are defined as target and source frames of a specific conceptual metaphor in the network, then that expression is positively evaluated as a metaphor.

However, it is not necessary that the evoked frames are immediately related to a metaphor entry in the repository. It is not unusual for specific metaphoric mappings not to be present in the conceptual network. This can be due to practical limitations as is often the case with manually created resources, or for principled reasons—for example, in cases where specific metaphors can be predicted from general ones, or for novel extensions that can be interpreted using general metaphors. In those cases, the system is often still able assess metaphoricity on the basis of general map-

pings defined at a higher level. For example, in the phrase *to cure poverty*, *poverty* evokes the Poverty frame and *cure* the Treating a Physical Affliction frame. In the conceptual network, Poverty is defined as a subcase of Social Problem. Furthermore, Treating a Physical Affliction incorporates the Physical Affliction frame as a role within it. In this case, although the more specific metaphor ADDRESSING POVERTY IS TREATING A DISEASE is not present in the repository network, the system can still identify the candidate pair *cure poverty* as metaphoric on the basis of the higher-level metaphoric mapping SOCIAL PROBLEMS ARE PHYSICAL AFFLICTIONS, as illustrated in Figure 4.
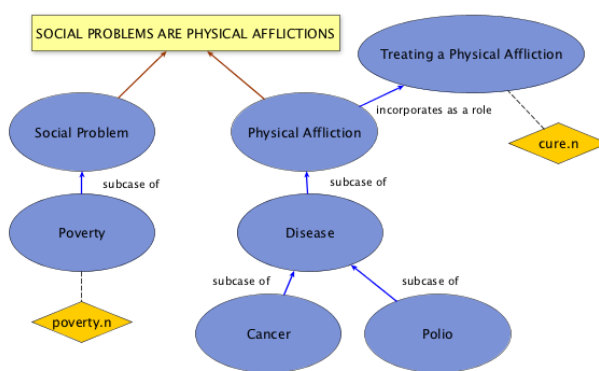

Figure 4. Structures accessed by phrase *cure poverty*.

Consequently, the system is often able to be resilient in the case of specific level gaps in the conceptual network as in the example above.

In addition, the relational patterns are assigned scores that represent the level of confidence that a linguistic expression with a matching frame and metaphor network pattern would actually be metaphoric or non-metaphoric. These scores are used to produce a metaphoricity score for each candidate expression. Although the scores presently assigned to the relational patterns are based on intuition, plans are underway to determine them empirically.

4 Analysis and Evaluation of Data

The extraction process generates a set of annotated sentences that can be used to both evaluate and refine the system, and to perform various kinds of corpus-based analysis. Annotation information includes the lemma, POS, frame, and frame family for both the source and the target terms, as well as

the name of the conceptual metaphor identified during the metaphor evaluation process. The purpose for which the extraction is being performed will affect which types of input are used (gold standard data vs. corpora).

4.1 System evaluation and improvements using gold standard

To evaluate the accuracy of the metaphor extractor, linguists collected attested sentences and annotated metaphoric expressions for the target domains Government, Bureaucracy, Democracy, Poverty, Taxation, and Wealth; they annotated all in-domain metaphoric expressions in the sentences where both the target and source were explicitly realized. Sentences were manually annotated for source and target word forms, source and target frames, and the constructional pattern used to express the metaphor. The metaphor extractor was run on these collected gold standard sentences, and the output compared to the annotations entered by the linguists. Table 3 shows the number of annotations in the gold standard, the recall (percentage of gold standard examples that were identified), and the precision (percentage of extracted examples that were correct) of the system for three languages.

Lang.	Anno.	Recall	Precision
English	301	0.86 (258/301)	0.85 (258/305)
Spanish	122	0.88 (107/122)	0.86 (107/125)
Russian	148	0.41 (60/148)	0.90 (60/67)

Table 3. Performance over gold standard data

As shown in Table 3, the system exhibits significantly lower recall for Russian than for the other languages. One of the reasons for this is that our instantiation of the conceptual network of frames and metaphors is not as well developed for Russian as for English and Spanish, containing significantly fewer metaphors and frames, as well as lexical units (LUs) which belong to them.[1] For example, Table 4 below shows the number of metaphors, frames, LUs, and the total number of frame-frame relations of the types used for metaphoricity evalu-

---

[1] As linguists continue to work on the repository, these numbers will grow.

ation. These relations include 'incorporates as a role,' 'subcase of,' and 'makes use of.'

	Metaphors	Frames	LUs	Rels
English	787	656	4308	838
Spanish	547	467	3521	506
Russian	127	303	1674	273

Table 4. Summary of repository content

It should be noted, however, that all the systems, including Russian, identified metaphoric expressions with a high degree of precision. Since the functioning of the metaphor detector depends on the correctness of conceptual metaphor theory, of its formalization in our system, and of the metaphor, frame, constructional pattern, and metaphor relational pattern representations in the repository, this result provides positive indication as to the validity in general of these aspects of the system. The metaphor detector thus in some sense implements the predictions of the formalized theory.

This has the added benefit that results contrary to expectation provide invaluable data for refining the system. For example, it is widely accepted that the government is often conceptualized a kind of physical structure, e.g. *foundation of government*, *the government collapsed overnight*, etc. The metaphor detector, based on representations captured in the repository, searching through a large corpus, turned up volumes of expressions such as *government building* and *government house* that are not metaphoric. This becomes a starting point of investigation to correct some aspect of the content of the repository, of the theory, or of its formalization.

4.2 Corpus-based analysis of metaphor

When corpora are used as input to the extraction system, the extraction results can be used to perform various kinds of corpus-based linguistic analyses. Such analyses can help provide an empirical basis for, and suggest refinements and improvements of, Conceptual Metaphor Theory. For instance, instead of relying on intuitions about how a given target domain is metaphorically conceptualized, it is possible to search a corpus and identify which source domain lemmas and frames are used, and with what relative frequency.

The richness of the extracted data and the structural relations identified via our repository enable us to analyze data at varying levels of specificity. For instance, a search of the English Gigaword corpus (Graff & Cieri 2003) for metaphor expressions with the target domain 'poverty' revealed several interesting patterns. Firstly, at the very general level of frame families, we observe that the most frequently occurring source terms were either location/motion related (e.g. being at a location, translational motion, motion impediments) or involved physical harm (e.g. disease, physical combat, or other harmful encounters). Figure 5 shows the relative frequency of these frame families.
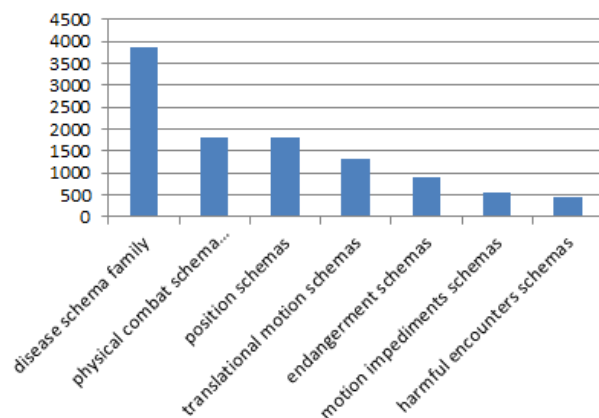
Figure 5. Most frequently occurring source families within poverty data.

At a somewhat more specific level, we can examine which specific frames within one of these families are being evoked. Figure 6 looks within the Translational Motion family, and shows the number of extracted metaphor expressions that evoke each of the frames within that family.
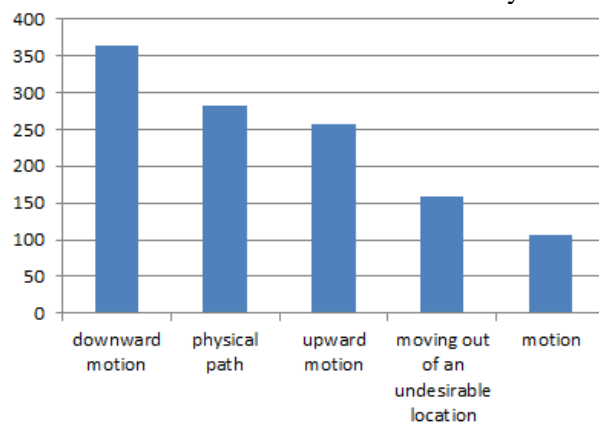

Figure 6. Number of extracted metaphor expressions that evoke various translational motion frames

Looking at a yet more specific level, we can examine which lexical items are used to evoke a given frame. Figure 7, below, shows this data for the Downward Motion frame.
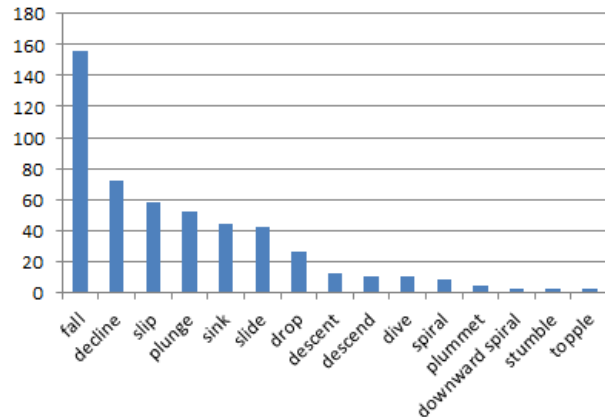

Figure 7. Number of different lexical units in extracted data that evoke the Downward Motion frame.

It is also possible to search for the conceptual metaphor that was discovered during the metaphoricity evaluation phase of the extraction process. For instance, Downward Motion lexemes such as *fall* and *slip* are used in expressions of the conceptual metaphor BECOMING IMPOVERISHED IS MOTION DOWNWARDS (e.g. *the young family fell/slipped into poverty*).

5   Conclusions

Our system moves beyond detection of metaphor, and enables us to perform many kinds of semantic analyses of metaphors in text. This affords the linguistic analyst additional insight into the conceptual structures characteristic of naturally-occurring language. Importantly, the different elements of the system each form part of a cycle, enabling an iterative development process, wherein extracted data informs linguistic analysis, improving the metaphor repository, or the theory, which in turn improves the quality of the extractor output. The resultant MetaNet metaphor repository and the extracted data can serve as valuable resources both for metaphor analysts and for the computational community at large.

Acknowledgments

References

Briscoe, T., Carroll, J., & Watson, R. (2006, July). The second release of the RASP system. In *Proceedings of the COLING/ACL on Interactive presentation sessions* (pp. 77-80). Association for Computational Linguistics.

Deignan, A. (2005). *Metaphor and corpus linguistics*. Amsterdam/Philadelphia: John Benjamins.

De Marneffe, M. C., MacCartney, B., & Manning, C. D. (2006). Generating typed dependency parses from phrase structure parses. In *Proceedings of LREC* (Vol. 6, No. 2006, pp. 449-454).

Feldman, J., Dodge, E. & Bryant, J. (2009). Embodied Construction Grammar. In Heine B., Narrog H., (eds) *The Oxford Handbook of Linguistic Analysis*. Oxford: Oxford U Press, 111–38.

Fillmore, C. J. 1976. Frame semantics and the nature of language. In Origins and Evolution of Language and Speech, ed. by Stevan R. Harnad, Horst D. Steklis, & Jane Lancaster, 20-32. Annals of the NY Academy of Sciences, Vol. 280.

Graff, D. & Cieri, C. (2003). English Gigaword LDC2003T05. Web Download. Philadelphia: Linguistic Data Consortium.

Kövecses, Z. (2011). Methodological issues in conceptual metaphor theory. In S. Handl & H-J. Schmid (Eds.), *Windows to the Mind: Metaphor, metonymy and conceptual blending* (pp. 23-40). Berlin/New York: Mouton de Gruyter.

Krötzsch, M., Vrandečić, D., & Völkel, M. (2006). Semantic mediawiki. In *The Semantic Web-ISWC 2006* (pp. 935-942). Springer Berlin Heidelberg.

Lakoff, G. (1993). The Contemporary Theory of Metaphor. In A. Ortony (Ed.), *Metaphor and Thought*. Cambridge: Cambridge University Press, 202–251.

Lakoff, G. (1993b). How metaphor structures dreams: The theory of conceptual metaphor applied to dream analysis. *Dreaming*, *3*(2), 77.

Lakoff, G., & Johnson, M. (1980). *Metaphors We Live By*. Chicago: University of Chicago Press.

Lakoff, G., & Johnson, M. (1999). *Philosophy in the flesh: The embodied mind and its challenge to western thought*. Basic books.

Lakoff, G. (2008). The neural theory of metaphor. *The Cambridge handbook of metaphor and thought*, 17-38.

Lakoff, G. (2014). Mapping the brain's metaphor circuitry: metaphorical thought in everyday reason. *Frontiers in Human Neuroscience*, *8*, 958.

Lloberes, M., Castellón, I., & Padró, L. (2010, May). Spanish FreeLing Dependency Grammar. In *LREC* (Vol. 10, pp. 693-699).

Martin J. H.. (2006). "A corpus-based analysis of context effects on metaphor comprehension". In Anatol Stefanowitsch and Stefan Th. Gries (Eds.), *Corpus-Based Approaches to Metaphor and Metonymy* (pp. 214-236). Berlin and New York: Mouton de Gruyter.

Mason, Z. J. (2004). "CorMet: A computational, corpus-based conventional metaphor extraction system." *Computational linguistics*, *30*(1), 23-44.

Mohler, M., Bracewell, D., Hinote, D., & Tomlinson, M. (2013). Semantic signatures for example-based linguistic metaphor detection. *Meta4NLP* 2013, 27.

Pragglejaz Group. (2007). MIP: A method for identifying metaphorically used words in discourse. *Metaphor and Symbol*, *22*(1), 1–39.

Ruppenhofer, J., Ellsworth, M., Petruck, M. R., Johnson, C. R., & Scheffczyk, J. (2010). FrameNet II: Extended theory and practice.

Berkeley, California: International Computer Science Institute

Shutova, E., Sun, L., & Korhonen, A. (2010, August). Metaphor identification using verb and noun clustering. In *Proceedings of the 23rd International Conference on Computational Linguistics* (pp. 1002-1010). Association for Computational Linguistics.

Stefanowitsch, Anatol. (2006). "Words and their metaphors: A corpus-based approach". In Anatol Stefanowitsch and Stefan Th. Gries (Eds.), *Corpus-Based Approaches to Metaphor and Metonymy* (pp. 63-105). Berlin and New York: Mouton de Gruyter.

Sullivan, K. S. (2007). *Grammar in Metaphor: A Construction Grammar Account of Metaphoric Language*. PhD dissertation, University of California Berkeley.

Sullivan, K. S. (2013). *Frames and Constructions in Metaphoric Language*. Amsterdam/Philadelphia: John Benjamins Publishing.

Turney, P. D., Neuman, Y., Assaf, D., & Cohen, Y. (2011). Literal and metaphorical sense identification through concrete and abstract context. In *Proceedings of the 2011 Conference on the Empirical Methods in Natural Language Processing* (pp. 680-690).

Wilks, Y., Galescu, J. A., Allen, J. & Dalton, A. (2013). Automatic Metaphor Detection using Large-Scale Lexical Resources and Conventional Metaphor Extraction. In *Proceedings of the First Workshop on Metaphor in NLP* (pp. 33-44), Atlanta, Georgia, 13 June 2013. Association for Computational Linguistics.