A Tiered Approach to the Recognition of Metaphor

David B. Bracewell, Marc T. Tomlinson, Michael Mohler, and Bryan Rink

Language Computer Corporation
Richardson, TX
{david,marc,michael,bryan}@languagecomputer.com

Abstract. We present a tiered-approach to the recognition of metaphor. The first tier is made up of highly precise expert-driven lexico-syntactic patterns which are automatically expended on in the second tier using lexical and dependency transformations. The final tier utilizes an SVM classifier using a variety of syntactic, semantic, and psycholinguistic features to determine if an expression is metaphoric. We focus on the recognition of metaphors in which the target is associated with the concept of "Economic Inequality" and examine the effectiveness of our approach for metaphors expressed in English, Farsi, Russian, and Spanish. Through experimental analysis we show that the proposed approach is capable of achieving 67.4% to 77.8% F-Measure depending on the language.

1 Introduction

Metaphor is a pervasive literary mechanism allowing individuals to view one concept in terms of the properties of another. As recent empirical studies have shown, metaphor is abundant in natural language occurring, as often as every third sentence [1]. Because of its abundance and often unique usage, it is important to build a system that can recognize and understand metaphor in order to aid natural language processing applications, such as authorship identification and semantic interpretation.

Metaphors enrich our conversation and provide a mechanism to map an abstract *target* domain into in a concrete *source* domain [2] allowing for the target to be discussed, understood, and affect assessed through the source. Below lists some examples of metaphor:

1. They cannot escape **poverty's grasp**.
2. The **burden of the tax** is onerous.
3. **Wages have stagnated**.

In the second example above, "tax" is being described as a "burden" transferring the the heavy weight property associated with a burden to tax.In the third example, "wages" are being described as stagnated invoking a mapping to a body of water or a volume of air which has ceased to move.

Automated methodologies for processing metaphor can be broken down into two main categories: recognition and interpretation. Interpretation of metaphor involves determining the intended, or literal, meaning of a metaphor [3, 4]. The recognition of metaphor entails determining if an expression is literal or figurative. Work on automated metaphor recognition dates back to the early 1990's with the work of Fass [5] based on selectional preference violations and more recently with the work of Shutova [6] using hierarchical graph factorization clustering.

In this paper, we propose a tiered approach to the recognition of metaphor. The first tier is made up of highly precise expert-constructed lexico-syntactic patterns which recognize metaphors associated with a predefined set of source and target domains. The second tier builds off the first by automatically expanding the lexico-syntactic patterns with dependency information and a series of lexical and dependency transformations. Finally, the third tier uses a combination of highly precise identification of target elements (spans of text associated with a target domain) with an SVM classifier to determine if a target element and a candidate source span of text linked to the target through a dependency chain is metaphoric.

In this paper, we limit our focus to recognition of metaphors in which the target is associated with the abstract domain of Economic Inequality. In particular, we focus on the following sub-domains of Economic Inequality: *Poverty*, *Wealth and Social Class*, and *Taxation*. We examine the effectiveness of our approach in multiple languages: English, Farsi, Russian, and Spanish.

The paper will proceed as follows. In section 2, we will present related work in metaphor processing. Then in section 3 will layout our tiered-approach to recognizing metaphor. Next, in section 4 will give experimental results of our approach for English, Farsi, Russian, and Spanish, Finally, in section 5 we will present concluding remarks.

2 Related Work

Metaphor has been studied by researchers in many fields, including, psychology linguistics, sociology, anthropology, and computational linguistics. A number of theories of metaphor have been proposed including the Contemporary Theory of Metaphor [2], the Conceptual Mapping model [7], the Structure Mapping Model [8], and the Attribute Categorization Hypothesis [9]. Based on these theories, databases of metaphors, such as the Master Metaphor List (MML) [10] for English and the Hamburg Metaphor Database (HMD) [11] for French and German, have been constructed. The MML provides links between domains (source and target) and their conceptual mappings. The HMD fuses EuroWordNet synsets with the MML source and target domains.

An active area of research in computational linguistics has been on the recognition of figurative language [12, 15]. The recognition of metaphoric expressions [5, 18], one part of the more general figurative language recognition task, has in particular seen a number of advances in recent years. Much of the early work on the recognition of metaphor used hand-crafted world knowledge. The met* [5] system determined if an expression was literal or figurative using selectional preferences. Figurative expressions were then determined to be metonymic using hand-crafted patterns or metaphoric using a manually constructed database of analogies. The CorMet [18] system determined the source and target concepts of a metaphoric expression using domain-specific selectional preferences mined from Internet resources.

Exemplar-based approaches to metaphor recognition have shown good results although are often limited in the metaphoric expressions they can identify. The Metaphor Interpretation, Denotation, and Acquisition System (MIDAS) [19] employed a database of conventional metaphors that could be searched to find a match for a metaphor discovered in text. In [20] semantic signatures were utilized to expand the metaphoric expressions producing a more durable and robust system for linking into the metaphor example repository.

Clustering-based approaches have also been prominent in the recognition of metaphor. In [17] noun-verb clustering starting from a small seed of one-word metaphors were used to generate clusters representing target and source concepts connected via a metaphoric relation. The clusters were then used to annotate the metaphoricity of text. Extending upon the noun-verb clustering work, [6] examined the use of hierarchical graph factorization clustering of nouns in a fully unsupervised approach to metaphor recognition. In [21] imageability and topicality were coupled with proto source induction for the recognition of metaphor.

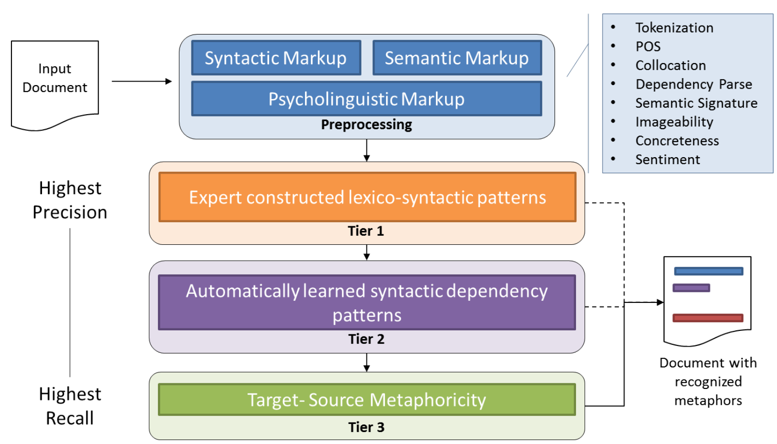

Fig. 1. Architecture of the proposed three-tiered metaphor recognition system.

3 Methodology

We propose a supervised approach to the recognition of metaphor that melds human knowledge and machine learning into a single three-tiered architecture.

The first tier, discussed in detail in section 3.1, is made up of high precision expert-constructed lexico-syntactic patterns for a predefined set of source and target domains. The second tier, discussed in detail in section 3.2, consists of syntactic dependency patterns which are automatically derived from the first-tier patterns. The final tier,

discussed in detail in section 3.3, uses a combination of highly accurate target domain identification using semantic signatures [3] with an SVM classifier that utilizes a variety of syntactic, semantic, and psycholinguistic features to determine if an expression is metaphoric. The overall architecture is illustrated in Figure 1.

3.1 Tier 1: Expert Constructed Lexico-Syntactic Patterns

The first tier in the proposed metaphor recognition system is made up of high precision lexico-syntactic patterns, examples are shown in Figure 2. The patterns define a source domain, target domain, and any lexical relation needed to exist for the two to be considered metaphoric. For example, In the English example in Figure 2, the pattern consists of a placeholder for a noun phrase from the source domain "BODY OF WATER" and noun phrase from the target domain of "POVERTY" and in order for the two to be metaphoric their must exist the word "of" between the source and target.

English: [BODY_OF_WATER:NOUN] of [POVERTY:NOUN] - sea of poverty - river of poverty - ocean of poverty
Farsi: [POVERTY:NOUN] [ABYSS:NOUN] - ورطه فقر "abyss of poverty"
Russian: [BUILDING:VERB] [POVERTY:NOUN] - Они построили нищету и ничего более. "They have built poverty and nothing else."
Spanish: [HUMAN:NOUN] de la [POVERTY:NOUN] - Los rostros de la pobreza en Mexico "The faces of poverty in Mexico"

Fig. 2. Examples of high precision lexico-syntactic patterns used in the first tier of the proposed metaphor recognition system.

We have defined a set of 51 source domains which either frequently occur with the target domain of Economic Inequality or one of the three sub-domains we are focusing on ("Poverty", "Wealth and Social Class", and "Taxation") in metaphoric expressions or are generic concepts often used in metaphors, e.g. "Movement" and "Vertical Scale". For each of the source domains we have defined lexical items for Nouns, Verbs, and Adjectives, which are strong exemplars of the given domain. The target domain lexical items are defined using semantic signatures of the sub-domain using the method described in [3] and enhanced using language specific lexicons for concepts not captured by the semantic signatures.

The pattern matching process allows for a gap of up to two words between any element in a pattern. For example, the English pattern listed in Figure 2 would also match "ocean of malnourished poor" and "river of unwanted and poor". The patterns in the first tier are checked against the training data and those with a precision less than 98% are discarded. This process left us with 450 patterns in English, 258 in Farsi, 65 in Russian, and 325 in Spanish covering the three sub-domains of Economic Inequality.

3.2 Tier 2: Automatically Learned Syntactic Dependency Patterns

The lexico-syntactic patterns used in the first tier are high precision, but low recall. To overcome this limitation the second tier constructs a set of automatically learned syntactic dependency patterns. The second tier leverages the first tier patterns as seed patterns. Dependency transformations are performed atop these seed patterns to discover new candidate patterns.

Following from the work of [25] we use two types of transformations. The first transformation replaces single lexical restrictions in the pattern with a wildcard. As an example, it would replace "of" in the English example shown in Figure 2 with a wildcard ("$T*$") creating the pattern:

$$\text{[BODY\_OF\_WATER:NOUN]}\ \mathbf{T*}\ \text{[POVERTY:NOUN]}$$

The second transformation works over expert defined syntactic dependency relations by replacing each dependency relation with a wildcard. Using the same example from Figure 2 with dependency information:

$$\text{[BODY\_OF\_WATER:NOUN]} \xrightarrow{prep} \text{of} \xrightarrow{pobj} \text{[POVERTY:NOUN])}$$

the pattern would be transformed into:

$$\text{[BODY\_OF\_WATER:NOUN]} \xrightarrow{T*} \text{of} \xrightarrow{pobj} \text{[POVERTY:NOUN]}$$
$$\text{[BODY\_OF\_WATER:NOUN]} \xrightarrow{prep} \text{of} \xrightarrow{T*} \text{[POVERTY:NOUN]}$$

where the relations "prep" and "pobj" get replaced by relational wildcards, meaning that other syntactic dependencies will also be considered. Using the expanded set of patterns containing wildcards, we search our training data to find matches. The matches with associated wildcards filled in are then assigned a confidence score based on the number of metaphors it matched and the number of non-metaphoric expressions it matched. Patterns matching less three times and patterns with a confidence less than 95% are discarded.

3.3 Tier 3: Target-Source Metaphoricity

The final tier of the linguistic metaphor identification system utilizes a variety of syntactic, semantic, and psycholinguistic features in an SVM classifier with highly accurate identification of target domains using semantic signatures (discussed briefly in Section 3.6 and in detail in [26]). The processing for recognizing metaphors in the third tier is as follows. First, candidate target elements (spans of text related to a target domain) are identified using semantic signatures. Semantic signatures [3] are constructed using the semantic knowledge in Wikipedia and WordNet. They have already been shown to be highly effective in identifying target domain in text and in identifying potential conceptual metaphors related to a linguistic metaphor.

Second, candidate source elements (spans of text which may be related to a source domain, known or not) are selected as all phrases within a predefined distance of the target element on a collapsed dependency tree. A collapsed dependency tree is one in which multiple dependency nodes have been merged based on a given criteria. We use Malt parser [27] in all four languages and collapse based on named entities, WordNet (and its foreign language equivalents) collocations, and noun/verb + preposition.

The final stage is to determine if each of the target element - source element pairs is metaphoric. We utilize an SVM classifier with the following features to make this judgment: (1) Imageability, (2) Concreteness, (3) Degree of dependency violation, (4) Selectional Strength, (5) Topicality, (6) Semantic Categories, and (7) Family Resemblance. Each of these features will be described in the following subsections. We optimized the C parameter of the SVM classifier by performing a grid search over the training data utilizing 10 fold cross-validation.

Imageability and Concreteness Imageability and Concreteness are concepts from the field of psychology relating to how well an object represented by a word can be imagined (Imageability) or linked to a sensory experience (Concreteness). Imageability has been found to be a strong indicator in the recognition of metaphors [28, 29].

We obtain imageability and concreteness scores by expanding the MRC [30] to full coverage of all WordNet senses across all parts-of-speech. Our methodology then moves beyond WordNet, allowing us to estimate psycholinguistics ratings in cases where such language resources are unavailable. The full details of the expansion can be found in [31].

Figure 3 lists examples of high and low imageability source elements in metaphoric and non-metaphoric phrases. The highly imageable element is “sting” and appears in the metaphoric expression “taxes will sting”. The low imageability element is “outperform” and appears in the non-metaphoric expression “the rich now out perform”. In this simple case, the use of an imageability feature would help to accept “taxes will sting” as being metaphoric while helping to reject “the rich now outperform”.

High Imageability:	If youre lucky and have itemized deductions, you will get a refund, but **taxes still sting**, especially for the middle class that pays more than its fair share.
Low Imageability:	But **the rich now outperform** the middle class by as much as the middle class outperform the poor.

Fig. 3. Example of high and low imageability source terms in metaphoric and non-metaphoric phrases.

3.4 Degree of dependency violation and Selectional Strength

The degree of dependency violation feature calculates how unexpected a source and target element are to share a given dependency relation. High degrees of violation are indicative of metaphor. For example, given the following sentence with metaphor highlighted in bold:

> OK, our friends on the left have one narrow statistic that says **wealth inequality is soaring**, but to be fair this does not capture the distribution either.

"wealth inequality" (target element) and "soaring" (source element) are unlikely to share the dependency relation of subject making them appearing in this relation a high degree of violation. In contrast, in the following example:

> Obama told Joe the Plumber in 2008 that its fair to **tax any income** over $250,000 at 39 percent and that when you "spread the wealth around; it's good for everybody."

"tax" (target element) and "any income" (source element) are likely to be seen with the dependency relation direct object making the combination a low dependency violation.

In a similar vein to the degree of dependency violation feature is the selectional strength feature. Selectional strength is a measure of how variable an element is in its selectional preference (here meaning between dependency relations). In simpler words, it relates to how many distinct classes of words share a given dependency relation with the source element. Source elements with low selectional strength, e.g. "is" or "think", are less likely to be metaphoric.

An example of a high selectional strength, i.e. has few semantic classes occurring in the given relation), is "tilts the field" in "But our **tax system tilts the field**". In contrast, the word "kill" has a low selectional strength in English as we tend to kill many types of things. This can be problematic as kill is often used metaphorically as is the case in "Many profitable employers argue that **taxes will kill jobs** and diminish our states competitive edge".

Both of these features are calculated using large corpora for each of the four language. Dependency relations are determine using Malt Parser and then combined as described earlier.

3.5 Topicality

Topicality is a measure of how topically related a word or phrase is to its context, i.e. hammer would topically related to a context discussing home improvement and not topically related to a context discussing legislation. Topically unrelated words are highly likely to be metaphoric.

The topical relatedness of a candidate source element is calculated by constructing a graph $G = (V, E)$ where, the vertices are the lemmatized version of the words in the candidate passage (sentence containing the candidate source and its context of two sentences before and after) and weighted edges exist between vertices whose similarity is greater than the average of all pairs. Similarity is determined using the cosine similarity between the row vectors constructed using Latent Semantic Analysis (LSA) over a large corpus.

Each lemma, l_i, is assigned a topicality score by:

$$score(l_i) = \frac{\sum_{l_j \in P(l_i)} count(l_j)}{\sum_{l_j} count(l_j)} \tag{1}$$

Where $P(l_i)$ is the set of lemmas for which a path exists l_i from l_j and $count(l_j)$ is the number of times l_j occurred in the candidate passage.

3.6 Semantic Categories

Following the work of [20] and [3], we incorporate the semantic features made available through the semantic signatures. Semantic signatures are a set of highly related and interlinked (WordNet) senses corresponding to a particular domain with statistical reliability. To generate the semantic signature we build off: (1) The semantic network encoded in WordNet; (2) The semantic structure underlying Wikipedia; and (3) Collocation statistics provided by statistical analysis of large corpora. We use the Princeton WordNet [32] for English, FarsNet [33] for Farsi, RussNet [34] for Russian, and the Multilingual Central Repository [35] for Spanish as our underlying WordNets.

We create target - source pairs of possible semantic categories using the output of the semantic signature and Wikipedia categories. A source and target element with a semantic mismatch, i.e. relating to semantic categories with little to no similarity, are more indicative of metaphor than those with no mismatch. For example, the target element “tax” and candidate source element “cow” in the expression “tax cow” represent a semantic mismatch as the corresponding semantic categories (taken from Wikipedia) “Taxation” and “Domesticated Animals” are semantically unrelated.

3.7 Family Resemblance

Family resemblance [36] based theories of categorization suggest that an item is classified based on its resemblance to the prototypical items in the category. Conceptual categories are arranged in a hierarchy in which the higher an item is in the tree the more generic it is and the lower it is in the tree the more specific it is, e.g. The concept "beagle" would be lower in the hierarchy than "canine". Each conceptual category has a set of culturally salient prototypical examples which are the items that most come to mind when imagining the given category. A prototypical example of furniture for an American would likely be "chair" whereas for a Japanese person it is likely to be "futon".

We approximate the source element's likelihood of being a prototypical example using a combination of TF-IDF and distributional semantics. We first find semantically similar concepts to the candidate source element using its associated semantic class as induced through distributional semantics. The items in the semantic class make up the examples (prototypical and not) for the candidate source element's category. We then use the TF-IDF values of the concepts in the semantic class to calculate a z-score for the candidate source element. The lower the candidate source element's z-score the less likely it is a prototype for the associated category.

4 Experimentation

For experimentation, we used a training set of roughly 1,000 metaphors per language over a wide variety of genres of data, including blog posts, forum posts, news articles, and transcripts. For testing we had a set of approximately 100 metaphors per language from the same genres as the training set. Spanish and English documents came from ClueWeb09[1], Farsi documents were gathered from web crawls, and Russian documents came from Ruwac[2].

The training and testing set were both annotated by multiple annotators. We did this not to determine inter-annotator agreement, but because we found a single annotator's recall in recognizing metaphor was poor. This problem arises based on the annotators' background and to how standard the metaphoric expression has become, e.g. "tax system" is a metaphor, but has become so standard that most people will not recognize it as one.

The results of the experimentation are shown in Table 1. We gave the system credit in recognizing a metaphor when the source and target elements it chose overlapped with the source and target elements chose by an annotator. This was because even two annotators would rarely agree on the exact same spans of text for the target and source elements.

As can be seen from Table 1, the tiered approach is able to precisely recognize metaphoric expressions in all four languages. Analyzing the errors, we found that rare words, errors in part of speech, and errors in the dependency parse caused a majority of the problems, especially in the non-English languages. In other cases the selectional strength, imageability, or concreteness of a term was too low causing a valid metaphor to be identified as non-metaphoric.

Table 2 shows the results when only the first two tiers of the system, expert-construct lexical patterns and automatically learned syntactic dependency patterns, were used. As one would expect the first two tiers resulted in high precision

[1] http://lemurproject.org/clueweb09/index.php
[2] http://corpus.leeds.ac.uk/ruscorpora.html

	Precision	Recall	F-Measure
English	73.8%	82.3%	77.8%
Farsi	60.0%	83.0%	69.6%
Russian	76.9%	51.9%	61.4%
Spanish	54.3%	88.7%	67.4%

Table 1. Results of the tiered system for recognizing metaphors.

	Precision	Recall	F-Measure
English	100.0%	7.0%	13.0%
Farsi	100.0%	10.7%	19.4%
Russian	100.0%	15.4%	26.7%
Spanish	100.0%	10.8%	19.5%

Table 2. Results of Tier 1 and 2 at recognizing metaphors.

but low recall. Interestingly, Russian which had the fewest expert constructed patterns had the highest recall.

5 Conclusion

In this paper, we presented a tiered-approach to the recognition of metaphor. The first tier was made up of highly precise expert-constructed lexico-syntactic patterns for a set of 51 source domains and the three predefined sub-domains of Economic Inequality. The second tier was made up of automatically constructed syntactic dependency patterns which were learned by performing lexical and dependency transformations atop the first tier patterns. The final tier used a combination of highly accurate target domain identification using semantic signatures with an SVM classifier using a variety of syntactic, semantic, and psycholinguistic features. We examined the effectiveness of our approach for English, Farsi, Russian, and Spanish. The proposed approach was capable of achieving 67.4% F-Measure in Russian and Spanish, 69.6% F-Measure in Farsi, and 77.8% F-Measure in English.

Acknowledgments

This research is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Defense US Army Research Laboratory contract number W911NF-12-C-0025. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoD/ARL, or the U.S. Government.

References

1. Shutova, E., Teufel, S.: Metaphor corpus annotated for source-target domain mappings. In: Proceedings of LREC. (2010)
2. Lakoff, G., et al.: The contemporary theory of metaphor. Metaphor and thought **2** (1993) 202–251
3. Bracewell, D., Tomlinson, M., Mohler, M.: Determining the conceptual space of metaphoric expressions. In: Computational Linguistics and Intelligent Text Processing. Springer (2013) 487–500
4. Shutova, E.: Models of metaphor in nlp. In: Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics (2010) 688–697
5. Fass, D.: met*: A method for discriminating metonymy and metaphor by computer. Computational Linguistics **17**(1) (1991) 49–90
6. Shutova, E., Sun, L.: Unsupervised metaphor identification using hierarchical graph factorization clustering. In: HLT-NAACL. (2013) 978–988
7. Ahrens, K., Chung, S., Huang, C.: Conceptual metaphors: Ontology-based representation and corpora driven mapping principles. In: Proceedings of the ACL 2003 workshop on Lexicon and figurative language-Volume 14, Association for Computational Linguistics (2003) 36–42
8. Wolff, P., Gentner, D.: Evidence for role-neutral initial processing of metaphors. Journal of Experimental Psychology: Learning, Memory, and Cognition **26**(2) (2000) 529
9. McGlone, M.: Conceptual metaphors and figurative language interpretation: Food for thought? Journal of memory and language **35**(4) (1996) 544–565
10. Lakoff, G.: Master Metaphor List. University of California (1994)
11. Eilts, C., Lönneker, B.: The Hamburg Metaphor Database. (2002)
12. Bogdanova, D.: A framework for figurative language detection based on sense differentiation. In: Proceedings of the ACL 2010 Student Research Workshop. ACLstudent '10 (2010) 67–72
13. Li, L., Sporleder, C.: Using gaussian mixture models to detect figurative language in context. In: Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics. HLT '10 (2010) 297–300
14. Peters, W., Wilks, Y.: Data-driven detection of figurative language use in electronic language resources. Metaphor and Symbol **18**(3) (2003) 161–173
15. Shutova, E.: Computational approaches to figurative language. PhD thesis, University of Cambridge (2011)
16. Shutova, E., Teufel, S., Korhonen, A.: Statistical metaphor processing. Comput. Linguist. **39**(2) (June 2013) 301–353
17. Shutova, E., Sun, L., Korhonen, A.: Metaphor identification using verb and noun clustering. In: Proceedings of the 23rd International Conference on Computational Linguistics. COLING '10, Stroudsburg, PA, USA, Association for Computational Linguistics (2010) 1002–1010
18. Mason, Z.: CorMet: A computational, corpus-based conventional metaphor extraction system. Computational Linguistics **30**(1) (2004) 23–44
19. Martin, J.: A computational model of metaphor interpretation. Academic Press Professional, Inc. (1990)
20. Mohler, M., Bracewell, D., Hinote, D., Tomlinson, M.: Semantic signatures for example-based linguistic metaphor detection. (2013) 27

21. Strzalkowski, T., Broadwell, G.A., Taylor, S., Feldman, L., Shaikh, S., Liu, T., Yamrom, B., Cho, K., Boz, U., Cases, I., Elliot, K.: Robust extraction of metaphor from novel data. (2013) 67–76
22. Feldman, J., Narayanan, S.: Embodied meaning in a neural theory of language. Brain and language **89**(2) (2004) 385–392
23. Barnden, J., Glasbey, S., Lee, M., Wallington, A.: Reasoning in metaphor understanding: The att-meta approach and system. In: Proceedings of the 19th international conference on Computational linguistics-Volume 2, Association for Computational Linguistics (2002) 1–5
24. Mohler, M., Bracewell, D.B., Tomlinson, M.T.: Applying textual entailment to the interpretation of metaphor. In: n Proceedings of the 7th IEEE International Conference on Semantic Computing (ICSC 2013). (2013)
25. Rink, B., Roberts, K., Harabagiu, S., Scheuermann, R.H., Toomay, S., Browning, T., Bosler, T., Peshock, R.: Extracting actionable findings of appendicitis from radiology reports using natural language processing. AMIA Summits on Translational Science Proceedings **2013** (2013) 221
26. Bracewell, D.B., Tomlinson, M.T., Mohler, M.: Determining the conceptual space of metaphoric expressions. In: Computational Linguistics and Intelligent Text Processing, Springer Berlin Heidelberg (2013) 487–500
27. Nivre, J., Hall, J., Nilsson, J.: Maltparser: A data-driven parser-generator for dependency parsing. In: In Proc. of LREC-2006. (2006) 2216–2219
28. Broadwell, G.A., Boz, U., Cases, I., Strzalkowski, T., Feldman, L., Taylor, S., Shaikh, S., Liu, T., Cho, K., Webb, N.: Using imageability and topic chaining to locate metaphors in linguistic corpora. In: Social Computing, Behavioral-Cultural Modeling and Prediction. Springer (2013) 102–110
29. Turney, P.D., Neuman, Y., Assaf, D., Cohen, Y.: Literal and metaphorical sense identification through concrete and abstract context. In: Proceedings of the 2011 Conference on the Empirical Methods in Natural Language Processing. (2011) 680–690
30. Coltheart, M.: The MRC psycholinguistic database. The Quarterly Journal of Experimental Psychology **33**(4) (1981) 497–505
31. Mohler, M., Tomlinson, M., Bracewell, D., Rink, B.: Semi-supervised methods for expanding psycholinguistics norms by integrating distributional similarity with the structure of wordnet. In: Proceedings of the 9th Language Resources and Evaluation Conference. (2014)
32. Fellbaum, C.: WordNet, An Electronic Lexical Database. The MIT Press (1998)
33. Shamsfard, M., Hesabi, A., Fadaei, H., Mansoory, N., Famian, A., Bagherbeigi, S., Fekri, E., Monshizadeh, M., Assi, S.: Semi automatic development of farsnet; the persian wordnet. In: Proceedings of 5th Global WordNet Conference, Mumbai, India. (2010)
34. Azarova, I., Mitrofanova, O., Sinopalnikova, A., Yavorskaya, M., Oparin, I.: Russnet: Building a lexical database for the russian language. In: Proceedings of Workshop on Wordnet Structures and Standardisation and How this affect Wordnet Applications and Evaluation. Las Palmas. (2002) 60–64
35. Atserias, J., Villarejo, L., Rigau, G., Agirre, E., Carroll, J., Magnini, B., Vossen, P.: The MEANING Multilingual Central Repository. In: Proceedings of the 2nd Global WordNet Conference (GWC), Brno, Czech Republic (January 2004)
36. Rosch, E., Mervis, C.B.: Family resemblances: Studies in the internal structure of categories. Cognitive psychology **7**(4) (1975) 573–605