Multi-dimensional abstractness in cross-domain mappings

Jonathan Dunn
Department of Computer Science / Illinois Institute of Technology
`jonathan.edwin.dunn@gmail.com`

Abstract

Metaphor is a cognitive process that shapes abstract target concepts by mapping them to concrete source concepts. Thus, many computational approaches to metaphor make reference, directly or indirectly, to the abstractness of words and concepts. The property of abstractness, however, remains theoretically and empirically unexplored. This paper implements a multi-dimensional definition of abstractness and tests the usefulness of each dimension for detecting cross-domain mappings.

1 Introduction

The idea of metaphor as cross-domain mapping goes back, at least, to Black (1954), who made explicit an earlier implicit view that linguistic metaphors depend upon non-linguistic (i.e., conceptual) connections between networks of concepts. Black's premises were later employed to represent groups of related linguistic metaphoric expressions using non-linguistic conceptual metaphors (for example, Reddy, 1979, and Lakoff & Johnson, 1980). Inherent in this approach to representing metaphor is the idea that metaphor is, at its core, a matter of cross-domain mapping (e.g., Lakoff, 1993); in other words, metaphor is a cognitive process that builds or maps connections between networks of concepts. The study of cognitive metaphor processes has largely focused on content-specific representations of such mappings within a number of content domains, such as TIME and IDEAS. Thus, a cross-domain mapping may be represented as something like ARGUMENT IS WAR.

Computational approaches to metaphor, however, have represented cross-domain mappings using higher-level properties like abstractness (Gandy, et al., 2013; Assaf, et al., 2013; Tsvetkov, et al., 2013; Turney, et al., 2011), semantic similarity (Li & Sporleder, 2010; Sporleder & Li, 2010), domain membership (Dunn, 2013a, 2013b), word clusters that represent semantic similarity (Shutova, et al. 2013; Shutova & Sun, 2013), and selectional preferences (Wilks, 1978; Mason, 2004). Most of these approaches rely on some concept of abstractness, whether directly (e.g., in terms of abstractness ratings) or indirectly (e.g., in terms of clusters containing abstract words). Further, these approaches have viewed abstractness as a one-dimensional scale between abstract and concrete concepts, with metaphor creating mappings from concrete source concepts to abstract target concepts.

Although both theoretical and computational treatments of metaphor depend upon the concept of abstractness, little has been done to either define or operationalize the notion. To fill this gap, this paper puts forward a multi-dimensional definition of abstractness and implements it in order to test the usefulness of the dimensions of abstractness for detecting cross-domain mappings.

2 Multi-dimensional abstractness

This approach recognizes four dimensions of abstractness: Domain of the Referent, Domain of the Sense, Fact-Status, and Function-Status, each of which has a range of values from more abstract to less abstract, as shown in Table 1. Domain refers to top-level categories in a hierarchical ontology as in, for example, ontological semantics (Nirenburg & Raskin, 2004), which uses four top-level domains: PHYSICAL, MENTAL, SOCIAL, ABSTRACT. Each concept belongs within a certain domain so that, at the highest level, cross-domain mappings can be represented as mappings between, for example, a PHYSICAL concept and an ABSTRACT concept. This dimension corresponds most with the traditional one-dimensional

approach to abstractness.

Here we divide domain membership into two types: (i) Domain of the Sense and (ii) Domain of the Referent. The idea is that a concept may refer to an object in one domain but define properties of that concept relative to another domain. For example, the concept *teacher* refers to a PHYS- ICAL object, a human who has physical properties. At the same time, the concept *teacher* is defined or distinguished from other humans in terms of SOCIAL properties, such as being focused on the education of students. Thus, the referent of the concept is within the PHYSICAL domain but its sense is within the SOCIAL domain. This is also true, for example, of countries (e.g., Mexico) which refer to a PHYSICAL location but also to a SOCIAL entity, the government and people who reside in that physical location. It is important to distinguish sense and reference when searching for cross-domain mappings because many concepts inherently map between different domains in this way (and yet are not considered metaphoric). Within both types of Domain, ABSTRACT is the category with the highest abstractness and PHYSI- CAL with the least abstractness.

Fact-Status is an ontological property as opposed to a domain within a hierarchical ontology. It represents the metaphysical property of a concept's dependence on human consciousness (Searle, 1995). In other words, PHYSICAL-FACTS are those, like *rocks* and *trees*, which exist in the external world independent of human perceptions. NON-INTENTIONAL facts are involuntary human perceptions such as pain or fear. INTENTIONAL facts are voluntary products of individual human consciousness such as *ideas* and *opinions*. COL- LECTIVE facts are products of the consciousness of groups of humans, such as *laws* and *govern- ments*. Thus, all categories except for PHYSICAL- FACTS are dependent on human consciousness. NON-INTENTIONAL and INTENTIONAL facts depend only on individuals, and in this sense are less abstract than COLLECTIVE facts, which exist only if a group of humans agrees to recognize their existence. This dimension of abstractness measures how dependent on human consciousness and how socially-constructed a concept is, with COLLEC- TIVE facts being more socially-constructed (and thus more society-dependent) than the others.

The final dimension of abstractness is Function- Status, which reflects how embedded function in-

Property	Value
Domain of the Referent	Abstract Mental Social Physical
Domain of the Sense	Abstract Mental Social Physical
Fact-Status	Collective Intentional Non-Intentional Physical
Function	Institutional Physical-Use Non-Agentive None
Event-Status	Object State Process
Animacy	Human Animate Inanimate Undefined

Table 1: Concept properties and values.

formation is in the sense of a concept. Function information is human-dependent, being present only as assumed by humans; thus, this dimension is also related to how human-centric a particular concept is. Many concepts have no function information embedded in them, for example *rock* or *tree*, and these are the least human-dependent. Some concepts have NON-AGENTIVE functions, sometimes called NATURAL functions; for example, the function of a *heart* is to pump blood. Some concepts have PHYSICAL-USE functions, in which the embedded function is a reflection of how humans use a physical object; for example, the function of a *hammer* is to drive nails. Finally, many concepts have embedded within them INSTITU- TIONAL functions, those which perform a social function only insofar as a group of individuals agree that the social function is performed. For example, a group of individuals may declare that certain taxes will be collected on income; but if others do not consent to the performance of that function then it is not performed (e.g., if the group had no legal authority to do so). Thus, INSTITU- TIONAL functions have the highest abstractness.

In addition to these dimensions of abstractness, two properties are added in order to test how they interact with these dimensions of abstractness: Event-Status, distinguishing OBJECTS from STATES and PROCESSES, and Animacy, distinguishing HUMANS from ANIMATE non-humans and INANIMATE objects.

3 Implementation

The system has two main steps: first, the input text is mapped to concepts in the Suggested Upper Merged Ontology (Niles & Pease, 2001); second, features based on the ontological properties of these concepts are used to represent the input sentences as a feature vector. The text is processed using Apache OpenNLP for tokenization, named entity recognition, and part of speech tagging. Morpha (Minnen, et al., 2001) is used for lemmatization. At this point word sense disambiguation is performed using SenseRelate (Pedersen & Kolhatkar, 2009), mapping the lexical words to the corresponding WordNet senses. These WordNet senses are first mapped to SynSets and then to concepts in the SUMO ontology, using existing mappings (Niles & Pease, 2003). Thus, the input to the second part of the system is the set of SUMO concepts which are pointed to by the input text. The properties of these concepts are contained in a separate knowledge-base developed for this system and available from the author. Each concept in SUMO has a value for each of the concept properties. This value is fixed and is the same across all instances of that concept. Thus, SenseRelate disambiguates input text into WordNet synsets which are mapped onto SUMO concepts, at which point the mapping from concepts to concept properties is fixed.

Feature Type	Number
Relative value frequency	23
Main value / concepts	6
Other values / concepts	6
Number of value types	6
Total	41

Table 2: Concept properties and values.

The concept properties discussed above are used to create a total of 41 features as shown in Table 2: First, 23 features contain the total number of instances of each possible value for the properties in each sentence relative to the number of concepts

present. Second, 6 features contain the relative frequency of the most common values of a property (the "main" value) and 6 features the relative frequency of all the other values (the "other" value). Third, 6 features contain the number of types of property values present in a sentence relative to the number of possible types.

4 Evaluation of the Features

We evaluated these features in a binary classification task using the VU Amsterdam Metaphor Corpus (Steen, et al., 2010), which consists of 200,000 words from the British National Corpus divided into four genres (academic, news, fiction, and spoken; the spoken genre was not evaluated) and annotated by five linguists. Metaphorically used prepositions have been untagged, as have ambiguously metaphoric sentences. Non-sentence fragments have been removed (e.g., titles and by-lines), along with very short sentences (e.g., "He said.").

The first step was to evaluate the features individually for their usefulness in detecting metaphoric language, allowing us to ask theoretical questions about which dimensions of abstractness are most related to metaphor. The Classifier-SubSetEval feature in Weka (Hall, et al., 2009) was used with the logistic regression classifier on the full corpus with 10 fold cross-validation. Three different search algorithms were used to ensure that the best possible combination of variables was found: the Greedy Stepwise, Linear Forward Selection, and Rank Search algorithms. The final feature rating was computed by taking the reverse ranking given by the GreedyStepwise search (e.g., the top ranked feature out of 41 is given a 41) and adding the number of folds for which that feature was selected by the other two algorithms. Table 3 below shows the top variables, arranged by score.

An interesting finding from this selection process is that each of the concept properties made the list of the top 16 features in the form of the *Property: Other* feature. In other words, the number of minority values for each property is useful for detecting cross-domain mappings. Next, each of the values for the Function property was a top feature, while only two of the Domain-Sense and one of the Domain-Referent properties made the list. The properties of Animacy and Fact are represented by the number of types present in the utterance, and

Property	Feature	Score
Function	Other Values	45.5
Fact-Status	Other Values	41
Animacy	Types	39.1
Fact-Status	Collective	37.8
Event-Status	Other Values	31
Function	Non-Agentive	30.6
Animacy	Other Values	29.8
Function	Physical-Use	29.8
Fact-Status	Types	28.3
Domain-Sense	Abstract	27.1
Domain-Sense	Other Values	25.1
Domain-Sense	Mental	22.1
Domain-Referent	Social	21.8
Function	None	20.5
Function	Institutional	17.1
Domain-Referent	Other Values	12.8

Table 3: Top features.

Fact is also significant for the number of concepts with the Collective value. These are interesting, and unexpected, findings, because the most important properties for detecting metaphor are not the traditional Domain-defined notions of abstractness, either Sense or Referent, but rather those notions of abstractness which are tied to a concept's degree of dependence on human consciousness and degree of being socially-constructed.

Using these top 16 variables, a binary classification task was performed on the entire VU Amsterdam Corpus, prepared as described above, using the logistic regression algorithm with 10 fold cross-validation, giving the results shown below in Table 4. These results show that while the full set of 41 features performs slightly better than the select set of the top 16, the performance gain is fairly small. For example, the F-measure on the full corpus raises from 0.629, using only the top 16 variables, to 0.649 using the full set of 41 variables. Thus, a similar performance is achieved much more efficiently (at least, in terms of the evaluation of the feature vectors; the top 16 variables still require many of the other variables in order to be computed). More importantly, this shows that the different dimensions of abstractness can be used to detect cross-domain mappings, licensing the inference that each of these operationalizations of abstractness represents an important and independent property of cross-domain mappings.

Var.	Corpus	Prec.	Recall	F1
Select	Full	0.655	0.629	0.629
All	Full	0.672	0.691	0.649
Select	Academic	0.655	0.682	0.600
All	Academic	0.639	0.676	0.626
Select	Fiction	0.595	0.597	0.592
All	Fiction	0.642	0.642	0.642
Select	News	0.749	0.813	0.743
All	News	0.738	0.808	0.746

Table 4: Results of evaluation.

5 Relation between the dimensions of abstractness

In order to determine the relationship between these dimensions of abstractness, to be sure that they are not measuring only a single scale, principal components analysis was used to determine how many distinct groups are formed by the properties and their values. The written subset of the American and Canadian portions of the International Corpus of English, consisting of 44,189 sentences, was used for this task. The corpus was not annotated for metaphor; rather, the purpose is to find the relation between the features across both metaphor and non-metaphor, using the direct oblimin rotation method.

#	Main Features	CL.	Vari.
1	Domain-Sense: Types	.834	18.7%
	Domain-Ref.: Types	.816	
	Event-Status: Types	.808	
2	Fact-Status: Main	.778	14.2%
	Fact-Status: Physical	.774	
3	Domain-Sense: Physical	.509	11.1%
	Domain-Ref.: Physical	.548	
	Event-Status: Object	.451	
4	Fact-Status: Intentional	.990	10.6%
	Fact-Status: Collective	.990	
	Fact-Status: Other	.913	
5	Domain-Sense: Abstract	.997	6.6%
	Domain-Ref.: Abstract	.997	
6	Domain-Sense: Main	.851	5.8%
	Domain-Ref.: Main	.773	
7	Function: Physical-Use	.876	4.4%
8	Event-Status: Process	.574	3.6%
9	Animacy: Main	.800	2.9%
10	Function: Non-Agentive	.958	2.4%

Table 5: Grouped features.

This procedure identified 10 components with eigenvalues above 1 containing unique highest value features, accounting for a cumulative 83.2% of the variance. These components are shown in Table 5 along with the component loadings of the main features for each component and the amount of variance which the component explains. All features with component loadings within 0.100 of the top feature are shown.

These components show two important results: First, the division of the Domain property into Sense and Referent is not necessary because the two are always contained in the same components; in other words, these really constitute a single-dimension of abstractness. Second, Domain, Function, and Fact-Status are not contained in the same components, but rather remain distinguishable dimensions of abstractness.

The important point of this analysis of the relations between features is that, even for those systems which do not represent abstractness in this way (e.g., systems which use numeric scales instead of nominal attributes), the dimensions of abstractness used here do represent independent factors. In other words, there is more than one dimension of abstractness. Domain membership, which corresponds most closely to the traditional one-dimensional view, refers essentially to how concrete or physical a concept is. Thus, *love* is more abstract than *grass*, but no distinction is possible between *love* and *war*. Fact-Status refers to how dependent on human consciousness a concept is. PHYSICAL concepts do not depend upon humans in order to exist. Thus, PHYSICAL concepts will be represented with the same degree of abstractness by both the Domain and Fact-Status properties. However, Fact-Status adds distinctions between abstract concepts. For example, *ideas* are not physical, but *laws* are both non-physical and depend upon complex social agreements. Function-Status refers to how much of the definition of a concept is dependent upon Function information which is, ultimately, only present in human understandings of the concept. This dimension adds distinctions between even physical concepts. For example, *canes* are just as physical as *sticks*, but *cane* embeds function information, that the object is used to help a human to walk, and this function information is dependent upon human consciousness. These two additional and distinguishable dimensions of abstractness, then, operationalize how dependent a concept is on human consciousness and how socially-constructed it is.

Using the traditional one-dimensional approach to abstractness, not all metaphors have abstract target concepts. For example, in the metaphoric expressions "My car drinks gasoline" and "My surgeon is a butcher," the concepts CAR and SURGEON are both PHYSICAL concepts in terms of Domain, and thus not abstract. And yet these concepts are the targets of metaphors. However, the concept DRINKING, according to this system, has an INTENTIONAL Fact-Status, because it is an action which is performed purposefully, and thus is an action which only sentient beings can perform. It is more abstract, then, than a verb like *uses*, which would not be metaphoric. The second example, however, cannot be explained in this way, as both SURGEON and BUTCHER would have the same concept properties (they are not included in the knowledge-base; both map to HUMAN). This phrase occurs only twice in the 450+ million word Corpus of Contemporary American English, however, and represents a rare exception to the rule.

6 Conclusions

This paper has examined the notion of abstractness, an essential component of many theoretical and computational approaches to the cross-domain mappings which create metaphoric language. There are two important findings: First, of the four posited dimensions of abstractness, three were shown to be both (1) members of separate components and (2) useful for detecting metaphoric mappings. These three dimensions, Domain Membership, Fact-Status, and Function-Status, are different and distinguishable ways of defining and operationalizing the key notion of abstractness. Second, and perhaps more importantly, the Fact-Status and Function-Status dimensions of abstractness, which are not directly present in the traditional one-dimensional view of abstractness, were shown to be the most useful for detecting metaphoric mappings. Although more evidence is needed, this suggests that cross-domain mappings are mappings from less socially-constructed source concepts to more socially-constructed target concepts and from less consciousness-dependent source concepts to more consciousness-dependent target concepts. This multi-dimensional approach thus provides a more precise definition of abstractness.

References

Assaf, D., Neuman, Y., Cohen, Y., Argamon, S., Howard, N., Last, M., Koppel, M. 2013. Why "dark thoughts" aren't really dark: A novel algorithm for metaphor identification. *2013 IEEE Symposium on Computational Intelligence, Cognitive Algorithms, Mind, and Brain*: 60–65. Institute of Electrical and Electronics Engineers.

Black, M. 1954. Metaphor. *Proceedings of the Aristotelian Society, New Series*, 55: 273-294.

Dunn, J. 2013a. Evaluating the premises and results of four metaphor identification systems. *Proceedings of the 14th International Conference on Computational Linguistics and Intelligent Text Processing, Volume I*: 471-486. Berlin, Heidelberg: Springer-Verlag.

Dunn, J. 2013b. What metaphor identification systems can tell us about metaphor-in-language. *Proceedings of the First Workshop on Metaphor in NLP*: 1-10. Association for Computational Linguistics.

Gandy, L., Allan, N., Atallah, M., Frieder, O., Howard, N., Kanareykin, S., Argamon, S. 2013. Automatic Identification of Conceptual Metaphors With Limited Knowledge. *Proceedings of the 27th Conference on Artificial Intelligence*: 328–334. Association for the Advancement of Artificial Intelligence.

Hall, M., Frank, E., Holmes, G., Pfahringer, B., Reutemann, P., Witten, I. H. 2009. The WEKA data mining software. *ACM SIGKDD Explorations Newsletter*, 11(1): 10.

Lakoff, G. 1993. The contemporary theory of metaphor. *Metaphor and thought, 2nd edition*: 202-251. Cambridge, UK: Cambridge Univ Press.

Lakoff, G., Johnson, M. 1980. *Metaphors we live by*. Chicago: University Of Chicago Press.

Li, L., Sporleder, C. 2010a. Linguistic Cues for Distinguishing Literal and Non-literal Usages. *Proceedings of the 23rd International Conference on Computational Linguistics: Posters*: 683-691. Association for Computational Linguistics.

Li, L., Sporleder, C. 2010b. Using Gaussian Mixture Models to Detect Figurative Language in Context. *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*: 297–300. Association for Computational Linguistics.

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

Minnen, G., Carroll, J., Pearce, D. 2001. Applied morphological processing of English. *Natural Language Engineering*, 7(3), 207-223.

Niles, I., Pease, A. 2001. Towards a standard upper ontology. *Proceedings of the International Conference on Formal Ontology in Information Systems*: 2-9. Association for Computing Machinery.

Niles, I., Pease, A. 2003. Linking lexicons and ontologies: Mapping WordNet to the Suggested Upper Merged Ontology. *Proceedings of the 2003 International Conference on Information and Knowledge Engineering*: 412-416. World Congress in Computer Science, Computer Engineering, and Applied Computing.

Nirenburg, S., Raskin, V. 2004. *Ontological Semantics*. Cambridge, MA: MIT Press.

Pedersen, T., Kolhatkar, V. 2009. WordNet::SenseRelate::AllWords–A broad coverage word sense tagger that maximimizes semantic relatedness. *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics, Companion Volume: Demonstration Session*: 17-20. Association for Computational Linguistics.

Reddy, M. 1979. The conduit metaphor: A case of frame conflict in our language about language. *Metaphor and Thought, 1st edition*: 284-310. Cambridge, UK: Cambridge Univ Press.

Searle, J. 1995. *The construction of social reality*. New York: The Free Press.

Shutova, E., Sun, L. 2013. Unsupervised Metaphor Identification using Hierarchical Graph Factorization Clustering. *Proceedings of Human Language Technologies: The 2013 Annual Conference of the North American Chapter of the Association for Computational Linguistics*: 978-988. Association for Computational Linguistics.

Shutova, E., Teufel, S., Korhonen, A. 2013. Statistical Metaphor Processing. *Computational Linguistics*, 39(2), 301-353.

Steen, G. J., Dorst, A. G., Herrmann, J. B., Kaal, A. A., Krennmayr, T. 2010. Metaphor in usage. *Cognitive Linguistics*, 21(4), 765-796.

Tsvetkov, Y., Mukomel, E., Gershman, A. 2013. Cross-Lingual Metaphor Detection Using Common Semantic Features. *Proceedings of the First Workshop on Metaphor in NLP*: 45-51. Association for Computational Linguistics.

Turney, P. D., Neuman, Y., Assaf, D., Cohen, Y. 2011. Literal and Metaphorical Sense Identification Through Concrete and Abstract Context. *Proceedings of the Conference on Empirical Methods in Natural Language Processing*: 680-690. Association for Computational Linguistics.

Wilks, Y. 1978. Making preferences more active. *Artificial Intelligence*, 11(3), 197-223.