Discovering Conceptual Metaphors Using Source Domain Spaces

Samira Shaikh[1], Tomek Strzalkowski[1], Kit Cho[1], Ting Liu[1], George Aaron Broadwell[1], Laurie Feldman[1], Sarah Taylor[2], Boris Yamrom[1], Ching-Sheng Lin[1], Ning Sa[1], Ignacio Cases[1], Yuliya Peshkova[1] and Kyle Elliot[3]

[1]State University of New York – University at Albany	[2]Sarah M. Taylor Consulting LLC	[3]Plessas Experts Network

samirashaikh@gmail.com

Abstract

This article makes two contributions towards the use of lexical resources and corpora; specifically making use of them for gaining access to and using word associations. The direct application of our approach is for detecting linguistic and conceptual metaphors automatically in text. We describe our method of building *conceptual spaces*, that is, defining the vocabulary that characterizes a Source Domain (e.g., Disease) of a conceptual metaphor (e.g., Poverty is a Disease). We also describe how these conceptual spaces are used to group linguistic metaphors into conceptual metaphors. Our method works in multiple languages, including English, Spanish, Russian and Farsi. We provide details of how our method can be evaluated and evaluation results that show satisfactory performance across all languages.

1 Introduction

Metaphors are communicative devices that are pervasive in discourse. When understood in a cultural context, they provide insights into how a culture views certain salient concepts, typically broad, abstract concepts such as poverty or democracy. In our research, we are focusing on metaphors on targets of governance, economic inequality and democracy, although our approach works for metaphors on any target. Suppose it is found in a culture that its people use metaphors when speaking of poverty; for example, they may talk about "symptom of poverty" or that "poverty infects areas of the city". These expressions are linguistic metaphors that are instances of a broader conceptual metaphor: Poverty is a Disease. Similarly, if it is found that common linguistic metaphors about poverty for peoples of a culture include "deep hole of poverty" and "fall into poverty", it would lead to the conceptual metaphor: Poverty is an Abyss. A communicator wishing to speak of ways to deal with poverty would use metaphors such as "treat poverty" and "cure poverty" to make their framing consistent with the conceptual metaphor of Disease, whereas she would use metaphors such as "lift out of poverty" when speaking to people who are attuned to the Abyss conceptual metaphor. Here Disease and Abyss are source domains, and poverty is the target domain. Relations, like "symptom of", "infect" and "fall into" from the respective source domains are mapped onto the target domain of poverty.

In order to discover conceptual metaphors and group linguistic metaphors together, we make use of corpora to define the conceptual space that characterizes a source domain. We wish to discover the set of relations that are used literally for a given source domain, and would create metaphors if applied to some other target domain. That is, we wish to *automatically* discover that relations such as "symptom", "infect", "treat" and "cure" characterize the source domain of Disease, for example. To create the conceptual spaces, we employ a fully automated method in which we search a balanced corpus using specific search patterns. Search patterns are so created as to look for co-occurence of

relations with members of a given source domain. Relations could be nouns, verbs, verb phrases and adjectives that are frequently used literally within a source domain. In addition, we calculate the frequency with which relations occur in a given source domain, or Relation Frequency. We then calculate the Inverse Domain Frequency (IDF), a variant of the inverse document frequency measure quite commonly used in field of information retrieval; the IDF captures the degree of distribution of relations across all source domains under consideration. Using these two measures, the relation frequency and inverse domain frequency, we are able to rank relations within a source domain. This ranked list of relations are then used to group linguistic metaphors belonging to the same source domain together. A group of linguistic metaphors so formed is a conceptual metaphor.

2   Related Research

Most current research on metaphor falls into three groups: (1) theoretical linguistic approaches (as defined by Lakoff & Johnson, 1980; and their followers) that generally look at metaphors as abstract language constructs with complex semantic properties; (2) quantitative linguistic approaches (e.g., Charteris-Black, 2002; O'Halloran, 2007) that attempt to correlate metaphor semantics with their usage in naturally occurring text but generally lack robust tools to do so; and (3) social science approaches, particularly in psychology and anthropology that seek to explain how people deploy and understand metaphors in interaction, but which lack the necessary computational tools to work with anything other than relatively isolated examples.

Metaphor study in yet other disciplines has included cognitive psychologists (e.g., Allbritton, McKoon & Gerrig, 1995) who have focused on the way metaphors may signify structures in human memory and human language processing. Cultural anthropologists, such as Malkki in her work on refugees (1992), see metaphor as a tool to help outsiders interpret the feelings and mindsets of the groups they study, an approach also reflective of available metaphor case studies, often with a Political Science underpinning (Musolff, 2008; Lakoff, 2001).

In computational investigations of metaphor, knowledge-based approaches include MetaBank (Martin, 1994), a large knowledge base of metaphors empirically collected. Krishnakumaran and Zhu (2007) use WordNet (Felbaum, 1998) knowledge to differentiate between metaphors and literal usage. Such approaches entail the existence of lexical resources that may not always be present or satisfactorily robust in different languages. Gedigan et al (2006) identify a system that can recognize metaphor. However their approach is only shown to work in a narrow domain (Wall Street Journal, for example).

Computational approaches to metaphor (largely AI research) to date have yielded only limited scale, often hand designed systems (Wilks, 1975; Fass, 1991; Martin, 1994; Carbonell, 1980; Feldman & Narayan, 2004; Shutova & Teufel, 2010; inter alia, also Shutova, 2010b for an overview). Baumer et al (2010) used semantic role labels and typed dependency parsing in an attempt towards computational metaphor identification. However, they self-report their work to be an initial exploration and hence, inconclusive. Shutova et al (2010a) employ an unsupervised method of metaphor identification using nouns and verb clustering to automatically impute metaphoricity in a large corpus using an annotated training corpus of metaphors as seeds. Their method relies on annotated training data, which is difficult to produce in large quantities and may not be easily generated in different languages.

More recently, several important approaches to metaphor extraction have emerged from the IARPA Metaphor program, including Broadwell et al (2013), Strzalkowski et al. (2014), Wilks et al (2013), Hovy et al (2013) inter alia. These papers concentrate on the algorithms for detection and classification of individual linguistic metaphors in text rather than formation of conceptual metaphors in a broader cultural context. Taylor et al (2014) outlines the rationale why conceptual level metaphors may provide important insights into cross-cultural contrasts. Our work described here is a first attempt at automatic discovery of conceptual metaphors operating within a culture directly from the linguistic evidence in language.

3   Our Approach

The process of discovering conceptual metaphors is necessarily divided into two phases: (1) collecting evidence about potential source domains that may be invoked when metaphorical expressions are used; and (2) building a conceptual space for each sufficiently evidenced source domain so that linguistic metaphors can be accurately classified as instances of appropriate conceptual metaphors. In

this paper, we concentrate on the second phase only. Strzalkowski et al (2013) in their work have described a data-driven linguistic metaphor extraction method and our approach builds upon their work.

During the source domain evidencing phase, we established a set of 50 source domains that operate frequently with the target concepts we are focusing on (government, bureaucracy, poverty, wealth, taxation, democracy and elections). These domains were a joint effort of several teams participating in the Metaphor program and we are taking this set as a starting point. These are shown in Table 1.

A_GOD	CONFINEMENT	GAME	MONSTER	PLANT
A_RIGHT	CRIME	GAP	MORAL_DUTY	PORTAL
ABYSS	CROP	GEOGRAPHIC_FEATURE	MOVEMENT	POSITION AND CHANGE OF POSITION ON A SCALE
ADDICTION	DARKNESS	GREED	NATURAL_PHYSICAL_FORCE	RACE
ANIMAL	DESTROYER	HUMAN_BODY	OBESITY	RESOURCE
BATTLE	DISEASE	IMPURITY	PARASITE	STAGE
BLOOD_STREAM	ENERGY	LIGHT	PATHWAY	STRUGGLE
BODY_OF_WATER	ENSLAVEMENT	MACHINE	PHYSICAL_BURDEN	THEFT
BUILDING	FOOD	MAZE	PHYSICAL_HARM	VISION
COMPETITION	FORCEFUL_EXTRACTION	MEDICINE	PHYSICAL_LOCATION	WAR

Table 1. Set of 50 source domains that operate frequently with target concepts being investigated. Only English names are shown for ease of presentation, equivalent sets in Spanish, Russian and Farsi have been created.

Some of the domains are self explanatory, while others require a further specification since the labels are sometimes ambiguous. For example, PLANT represents things that grow in the soil, not factories; similarly, BUILDING represents artifacts such as houses or edifices, but not the act of constructing something; RACE refers to a running competition, not skin color, etc.

Consequently, each of these domains need to be seeded with the prototypical representative elements to make the meaning completely clear. This seeding occurs during the first phase of the process when a linguistic expression, such as "cure poverty" is classified as a linguistic metaphor. This process of classifying "cure poverty" as metaphorical is described in detail in Strzalkowski et al. (2013). Part of the seeding process is to establish that a source domain different than the target domain (here: poverty) is invoked by the relation (here: cure). To find the source domain where "cure" is typically used literally, we form a linguistic pattern [cure [OBJ: X/nn]] (derived automatically from the parsed metaphoric expression) which is subsequently run through a balanced language corpus. Arguments matching the variable X are then clustered into semantic categories, using lexical resources such as Wordnet (Felbaum, 1998) and the most frequent and concrete category is selected as a possible source domain (proto-source domain). From the balanced language corpus, it is possible to compute the frequency with which the arguments resulting from search appear with relation ("cure"). We determine concreteness by looking up concreteness score in MRC psycholinguistic database (Coltheart 1981, Wilson 1988). As may be expected, the initial elements of the proto-source obtained from the above patterns will include: *disease, cancer, plague*, etc. These become the seeds of the source domain DISEASE in our list. The same process was performed for each of the 50 domains listed here, for each of the 4 languages under consideration. Additional Source Domains are continuously generated bottom-up fashion by this phase 1 process elaborated above. In Table 2, we show seeds so obtained for a few source domains.

DISEASE	disease, cancer, plague
ABYSS	abyss, chasm, crevasse
BODY_OF_WATER	ocean, lake river, pond, sea
PLANT	plant, tree, flower, weed, shrub, vegetable
GEOGRAPHIC_FEATURE	land, land form, earth, mountain, plateau, island, valley

Table 2. Example of seeds corresponding to a few source domains

Once such seeds are obtained, we perform another search through a balanced corpus in the corresponding language to discover relations that characterize the source domains. The purpose of source domain spaces in our research is two-fold: a) to provide a sufficiently complete characterization of a source domain via a list of relations ; and b) such a list of relations should sufficiently distinguish *between* different source domains. Creating these spaces is phase 2 of the conceptual metaphor discovery process.

We search for nouns, verbs and verb phrases, and adjectives that co-occur with seeds of given source domain with sufficiently high frequency and sufficiently high mutual information. Our goal with this process is to approximate normal usage patterns of relations within source domains. The results of balanced corpora search form our *conceptual spaces*. The balanced corpora we use are English: Corpus of Contemporary American English (Davies, 2008), Spanish: Corpus del Español Actual (Davies, 2002), Russian: Russian National Corpus[2] and Farsi: Bijankhan Corpus (Oroumchian et al., 2006). In addition to retrieving the relations, we retrieve the frequency with which these relations can be found to co-occur with seeds of a source domain, Relation Frequency (RF). We calculate Inverse Domain Frequency (IDF) of all relations across all 50 source domains using a variant of the inverse document frequency measure. The formula for IDF is as given below:

IDF = log (total number of source domains / total number of source domains a relation appears in)

For example, if a relation such as "dive into" is found to appear in two source domains, BODY_OF_WATER and GEOGRAPHIC_FEATURE, then the IDF for "dive into" would be log (50/2). The rank of a relation is computed as the product of RF and IDF. However, computing rank using RF without normalization results in inflated ranks for relations that are quite common across domains even when they do not sufficiently disambiguate between the domains. We assume a normal distribution of frequencies of relations within a source domain and normalize RF by taking its logarithm. We also normalize with respect to seeds within a source domain. If a relation frequency is disproportionately high with a specific seed, we disregard that frequency. For example, one of the seeds for the source domain of BUILDING is "house". A search through balanced corpus for nouns adjacent to "house" revealed a disproportionately large number for "white", which is meant to be the White House, and would be disregarded.

In Table 3, we show a few top ranked relations for the source domains DISEASE and BODY_OF_WATER. In columns 1 and 2, we show the source domain and the relation. Column 3 shows the relation frequency and column 4 shows the part of speech of relation (V=verb or verb phrase, N=noun, ADJ=adjective). An RF score of 800 for row 1 indicates that the relation "diagnose with" appears 800 times with one or more of the seeds we search for source domain DISEASE ("diagnose with cancer", "diagnose with disease" and so on. In column 5, we show the position where the relation is commonly found to co-occur with the source domain. For example, "afflict" in row 2 has a position "after" which means it appears after DISEASE: "DISEASE afflict(s)"; whereas row 3 would be read as "affict with DISEASE" since it appears "before". In column 6, we show the normalized RF*IDF score. The highest RF*IDF score for a relation across our spaces is 2.165. From Table 3, we can see that even if frequency for some relations may be relatively low, their rank would be high if they are strongly associated with a single source domain.

	1. Source Domain	2. Relation	3. RF	4. Type	5. Position	6. Norm RF*IDF
1	DISEASE	diagnose with	800	V	before	1.94
2	DISEASE	afflict	85	V	after	1.67
3	DISEASE	afflict with	33	V	before	1.52
4	DISEASE	cure of	29	N	before	1.46
5	BODY_OF_WATER	dive into	49	V	before	2.01
6	BODY_OF_WATER	wade through	44	V	before	1.88
7	BODY_OF_WATER	wade into	42	V	before	1.84
8	BODY_OF_WATER	rinse in	41	V	before	1.80

Table 3. A few top ranking relations for the source domains DISEASE and BODY_OF_WATER. Relations are ranked by their normalized RF*IDF score.

---

[2] http://ruscorpora.ru/en/

With the conceptual spaces defined in this manner, we can now use them to group linguistic metaphors together. Shaikh et al (2014) have created a repository of thousands of automatically extracted lingusitic metaphors in all four languages, which we are using to create conceptual metaphors. To discover which conceptual metaphors exist within such large sets of linguistic metaphors would be quite challenging, if not impossible, for a human expert. We automatically assign each linguistic metaphor to ranked list of source domains.

Consider the linguistic metaphor "plunge into poverty", where the relation is "plunge into". We search through our conceptual spaces and retrieve a list of source domains where the relation "plunge into" may appear. From this list, only the domains that have this relation RF*IDF score higher than a threshold are considered. This threshold is currently assigned to be 0.40, although it is subject to further experimentation. The source domain where the RF*IDF score of "plunge into" is the highest is chosen as the source domain, along with the next source domains only if the difference in scores is 5% or lower. Tables 4 and 5 depicts this part of algorithm for two relations, "plunge into" and "explorar" (from Spanish – "explore"). The relation "plunge into" is thus assigned to BODY_OF_WATER source domain. "explorar" is assigned to GEOGRAPHIC_FEATURE and BODY_OF_WATER since difference in RF*IDF scores is less than 5%.

Relation	Source Domains	RF*IDF
plunge into	BODY_OF_WATER	1.82
	DARKNESS	1.28
	ABYSS	0.68
	WAR	0.57
	GEOGRAPHIC_FEATURE	0.48

Relation	Source Domains	RF*IDF
explorar	GEOGRAPHIC_FEATURE	0.77
	BODY_OF_WATER	0.76
	PHYSICAL_LOCATION	0.56
	PATHWAY	0.56
	BUILDING	0.41

Table 4 and Table 5. Assigning relations of linguistic metaphor to source domains. "plunge into" is assigned to BODY_OF_WATER; "explorar" is assigned to GEOGRAPHIC_FEATURE and BODY_OF_WATER

Once this process of assigning linguistic metaphors to source domains is accomplished for all linguistic metaphors in our repository, we validate the resulting conceptual metaphors. A small percentage of metaphors cannot be assigned to any of the 50 Source Domains. We explain the validation process in Section 4. In Tables 6 and 7, we show sample conceptual metaphors in English and Spanish. Our validation process revealed an interesting insight regarding forming conceptual metaphor, wherein they should contain relations that are anchors for that given source domain that we shall describe next.

SD	Target	Source	Sentence	RF*IDF
BODY_OF_WATER	POVERTY	rise from	The United States has always had a culture with a high regard for those able to **rise from** poverty to riches .	1.78
		slide into	Government aid kept some Americans from **sliding into** poverty last year , some analysts say .	0.81
		shallow	By 2004 , the government was actually giving more , each month , to families in **shallow** poverty than to families in deep poverty .	1.13
		shallow	Those in what could be considered **shallow** poverty , between 50 and 100 percent of the poverty line , received $ 448 .	1.13
		tumble into	About 46 million more people are expected to **tumble into** poverty this year amid the largest decline in global trade in 80 years, according to the World Bank.	1.50
		slip into	Mountains of research tell us that children reared outside of intact marriages are much more likely than other kids to **slip into** poverty, become victims of child abuse, fail at school and drop out, use illegal drugs, launch into premature sexual activity, become unwed teen mothers, divorce, commit suicide and experience other signs of mental illness, become physically ill, and commit crimes and go to jail.	1.35
		deep	The Census Bureau uses a third measure , " **deep** poverty , " which it defines as living on less than half of the amount needed to escape poverty (for a family of three , that means living on less than $ 9,000 a year) .	0.40
		plunge into	With the economic crisis threatening to **plunge** more children **into poverty**	1.82

Table 6. A conceptual metaphor in English: POVERTY is a BODY_OF_WATER

SD	Target	Source	Sentence	RF*IDF
DISEASE	POVERTY	diagnosticar	Más Notas de este Especial Plantea PRD que Hacienda **diagnostique deudas** en estados y municipios… … …	2.04
		heredar	Afirmó que la comuna solventa **deudas heredadas** de la pasada administración que ascienden a 183 millones de pesos.	0.93
		erradicar	Viernes 29 de agosto de 2003 HACIA LA CUMBRE DE CANCUN Ese camino nunca ha **erradicado la pobreza**, dicen organizaciones católicas ……	0.91
		erradicar	Tweet EDUCACIÓN PARA ERRADICAR POBRE El magnate mexicano abogó en la ciudad colombiana de Cartagena de Indias, por **erradicar la pobreza** con educación y empleo para fortalecer las economías.	0.91
		combatir	Envía Banco Mundial expertos a Oaxaca para **combatir pobreza** Viernes, 25 de febrero de 20111 comentario Oaxaca.-	0.86

Table 7. A conceptual metaphor in Spanish: POVERTY is a DISEASE

3.1 Anchor relations in Conceptual Metaphors

When human assessors are presented with a set of linguistic metaphors and the task to assign them into a source domain, some relations will have stronger impact on their decision that others. For example, "cure" would almost invariably be assigned to DISEASE domain, while "dive in" would invoke BODY_OF_WATER domain. Other relations, such as "spread" or "fall into" are less specific, however, when paired with highly evocative relations above are likely to be classified the same way. Thus, there are two types of metaphorical relations in linguistic metaphors: (1) the highly evocative relations that unambiguously point to a specific source domain – we shall call them anchors; and (2) the relations that are compatible with the anchor but are not anchors themselves. We can add another class: (3) the relations that are not compatible with a given anchor. Thus, a set of linguistic metaphors that provides evidence for a conceptual metaphor should contain at least some anchor relations and the balance of the set may be composed of anchor-compatible relations. Our current hypothesis is that there should be at least one anchor for each 7 anchor compatible relations for a group of linguistic metaphors to provide a sufficient evidence for a conceptual metaphor.

As part of our validation process, we conducted a series of experiments with human assessors. One of the tasks was to assign a single linguistic metaphor to one of 50 source domains. As an illustrative example, we show in Table 8, one linguistic metaphor. When presented with this example, a majority of assessors chose ENEMY source domain, while DISEASE was selected second. Additionally, there was greater variance among their selections, only 31% chose the top source domain of ENEMY.

Subsequently, human assessors were presented a set of linguistic metaphors where at least one anchor relation was present. In this case, the majority of assessors chose the DISEASE source domain. Even though the "fight against poverty" example was included in the set, the presence of anchors such as "cure poverty" and "treat poverty" lead assessors to choose DISEASE source domain. The variance in selection was also less, a 70% majority choosing DISEASE. We show the conceptual metaphor in Table 9.

The summit has proven that there is a renewed appetite for the ***fight against*** ***poverty****.*
ENEMY: 31%; DISEASE: 17%; ANIMAL, MONSTER,….<10%

Table 8. A single linguistic metaphor was assigned a varied number of source domains by human assessors.

Of course, many government programs aim to ***alleviate*** ***poverty****.*
We seek to stimulate true prosperity rather than simply ***treat*** ***poverty.***
Unless the ***fight against*** ***poverty*** *is honestly addressed by the West, there will be many more Afghanistans.*
Above all, he knows that the only way to ***cure*** ***poverty*** *is to grow the economy.*
DISEASE: 70%; ENEMY: 30%

Table 9. A conceptual metaphor containing anchors. When sample metaphor from Table 8 is included in this set, human assessors still choose the source domain to be DISEASE.
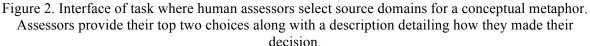
4 Evaluation and Results

A group of human experts who are native speakers and have been substantively trained to achieve high levels of agreement (0.78 Krippendorf's alpha (1970) or higher) form our validation team. In addition, we aim to run crowd-sourced experiments on Amazon Mechanical Turk. In Figure 1, we show a web interface we built to present our human assessors. The task shown here is the assignment of a single linguistic metaphor to one of 50 source domains. Then, we present our validation team with conceptual metaphors we created. Each conceptual metaphor is validated by at least two language experts. This interface is shown in Figure 2. These interfaces are carefully created by our team of social scientists and psychologists, designed to elicit proper responses from native speakers of the language.

Figure 1. Interface of task where human assessors select source domain for a single linguistic metaphor.

Task

1 / 4

Set 1

1) Of course, many government programs aim to alleviate poverty.

2) We seek to stimulate true prosperity rather than simply treat poverty; and that we believe in principle over power.

3) But it also reflects the programs that you really care about, that we've been using to fight poverty, including public housing, tax credits and food stamps.";

4) Unless the fight against poverty is honestly addressed by the West, there will be many more Afghanistans.

5) Above all, he knows that the only way to cure poverty is to grow the economy.

6) However, in his fight against poverty he called on all Americans regardless of their race or ethnicity.

Part A

Question 1: Taken together, which of the concepts below best describes the word POVERTY in the group of sentences above?

First choice   Second choice

[Please select one of the following: ⇕]   [Please select one of the following: ⇕]

Please briefly explain how you arrived at your decision for your first choice:

[   ]

Figure 2. Interface of task where human assessors select source domains for a conceptual metaphor. Assessors provide their top two choices along with a description detailing how they made their decision.

In Table 10, we show the number of conceptual metaphors currently in the repository and the accuracy of our method across four languages, as computed by using validation data. We show the number of conceptual metaphors present in the Governance target domain (metaphors about government and bureaucracy), Economic Inequality (dealing with metaphors of poverty, wealth and taxation) and Democracy (democracy and elections metaphors). These conceptual metaphors on the three target domains of Governace, Economic Inequality and Democracy, when compared across cultures could provide deep insight about peoples' perceptions regarding salient concepts.

We note that Russian and Farsi performance is lower than that in English and Spanish. The size of balanced corpus and accuracy of lexical tools such as stemmers and morphological analyzers affect performance of our algorithm. The Farsi balanced corpus is relatively small when compared to English balanced corpus. The smaller size affects computation of statistics such as Relation Frequency and subsequently the thresholds of RF*IDF scores. One improvement we are currently investigating is that the thresholds may be set specifically for a language.

	ENGLISH	SPANISH	RUSSIAN	FARSI
# of Governance Conceptual Metaphors	27	7	8	7
# of Economic Inequality Conceptual Metaphors	32	26	57	7
# of Democracy Conceptual Metaphors	51	16	18	8
Total # of Conceptual Metaphors	110	49	83	22
Accuracy (%)	85%	76%	67%	62%

Table 10. Number of conceptual metaphors discovered thus far and performance of our approach across four languages.

5 Conclusion and Future Work

In this article, we presented our approach towards automatic discovery of conceptual metaphors directly from linguistic evidence in a given language. We make use of corpora in two unique ways: the first is to discover prototypical seeds that form the basis of source domains and second is to create conceptual spaces that allow us to characterize the relations that operate within source domains automatically. In addition, our approach also allows us to distinguish *between* source domains as necessary. The validation results show that this is indeed a promising first attempt of tackling a challenging research problem.

We note that the assignment of source domains is limited to the set of 50 in our current prototype. This assumes a closed set of 50 source domains, whereas in reality, there might be many others that operate in the realm of metaphors we are investigating. Although additional source domains are continually being discovered in a bottom-up fashion by the linguistic metaphor extraction process, we cannot account for every source domain that may be relevant. One way of overcoming this limitation would be to define a source domain "OTHER" that would be the all-encompassing domain accounting for any yet undiscovered domains. The details of how it would be represented are still under investigation.

Another potential improvement to our method is to experimentally refine the threshold score of RF*IDF. Through large scale validation experiments, we could learn the optimal thresholds automatically by using machine learning.

6 Acknowledgements

References

David W. Allbritton, Gail McKoon, and Richard J. Gerrig. 1995. Metaphor-based schemas and text Representations: making connections through conceptual metaphors, *Journal of Experimental Psychology*: *Learning, Memory, and Cognition*, 21(3):612-625.

Jonathan, Charteris-Black. 2002. Second language figurative proficiency: A comparative study of Malay and English. *Applied Linguistics* 23(1):104–133.

George Aaron Broadwell, Umit Boz, Ignacio Cases, Tomek Strzalkowski, Laurie Feldman, Sarah Taylor, Samira Shaikh, Ting Liu and Kit Cho. 2013. *Using Imageability and Topic Chaining to Locate Metaphors in Linguistic Corpora*. In Proceedings of The 2013 International Conference on Social Computing, Behavioral-Cultural Modeling, & Prediction (SBP 2013), Washington D.C., USA.

Jaime Carbonell. 1980. Metaphor: A key to extensible semantic analysis. In *Proceedings of the 18th Annual Meeting on Association for Computational Linguistics*.

M. Coltheart. 1981. The MRC Psycholinguistic Database. *Quarterly Journal of Experimental Psychology*, 33A: 497-505.

Davies, Mark. 2008-. *The Corpus of Contemporary American English: 450 million words, 1990-present*. Available online at http://corpus.byu.edu/coca/.

Davies, Mark. 2002-. *Corpus del Español: 100 million words, 1200s-1900s*. Available online at http://www.corpusdelespanol.org.

Dan, Fass. 1991. met*: A Method for Discriminating Metonymy and Metaphor by Computer. *Computational Linguistics*, 17:49-90

Jerome Feldman, and Srinivas Narayanan. 2004. Embodied meaning in a neural theory of language. *Brain and Language*, 89(2):385–392.

Christiane D. Fellbaum. 1998. *WordNet: An electronic lexical database* (1st ed.). MIT Press.

Matt Gedigian, John Bryant, Srini Narayanan and Branimir Ciric. 2006. Catching Metaphors. In *Proceedings of the Third Workshop on Scalable Natural Language Understanding ScaNaLU 2006*, pages 41–48. New York City: NY.

Dirk Hovy, Shashank Shrivastava, Sujay Kumar Jauhar, Mrinmaya Sachan, Kartik Goyal, Huying Li, Whitney Sanders and Eduard Hovy. 2013. Identifying Metaphorical Word Use with Tree Kernels. In the *Proceedings of the First Workshop on Metaphor in NLP, (NAACL)*. Atlanta.

Krippendorff, Klaus. 1970. Estimating the reliability, systematic error, and random error of interval data. *Educational and Psychological Measurement, 30* (1),61-70.

Saisuresh Krishnakumaran and Xiaojin Zhu. 2007. Hunting elusive metaphors using lexical resources. In *Proceedings of the Workshop on Computational Approaches to Figurative Language*, pages 13–20. Rochester, NY.

George Lakoff, and Mark Johnson. 1980. *Metaphors we live by*. University Of Chicago Press, Chicago, Illinois.

George, Lakoff. 2001. *Moral politics: what conservatives know that liberals don't.* University of Chicago Press, Chicago, Illinois.

Liisa, Malkki. 1992. National geographic: The rooting of people and the territorialization of national identity among scholars and refugees. *Society for Cultural Anthropology*, 7(1):24–44.

James Martin. 1988. A computational theory of metaphor. *Ph.D. Dissertation.*

Musolff, Andreas. 2008. What can critical metaphor analysis add to the understanding of racist ideology? Recent studies of Hitler's anti-semitic metaphors, critical approaches to discourse analysis across disciplines. Critical *Approaches to Discourse Analysis Across Disciplines*, 2(2):1–10.

Kieran, O'Halloran. 2007. Critical discourse analysis and the corpus-informed interpretation of metaphor at the register level. *Oxford University Press*

Farhad Oroumchian, Samira Tasharofi, Hadi Amiri, Hossein Hojjat, Fahime Raja. 2006. Creating a Feasible Corpus for Persian POS Tagging.Technical Report, no. TR3/06, University of Wollongong in Dubai.

Samira Shaikh, Tomek Strzalkowski, Ting Liu, George Aaron Broadwell, Boris Yamrom, Sarah Taylor, Laurie Feldman, Kit Cho, Umit Boz, Ignacio Cases, Yuliya Peshkova and Ching-Sheng Lin. 2014. A Multi-Cultural Repository of Automatically Discovered Linguistic and Conceptual Metaphors. In *Proceedings of the The 9th edition of the Language Resources and Evaluation Conference ,* Reykjavik, Iceland.

Ekaterina Shutova and Simone Teufel. 2010a. Metaphor corpus annotated for source - target domain mappings. *In Proceedings of Language Resources and Evaluation Conference 2010*. Malta.

Ekaterina Shutova. 2010b. Models of metaphor in nlp. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, ACL '10, pages 688–697.

Ekaterina Shutova, Tim Van de Cruys, and Anna Korhonen. 2012. *Unsupervised metaphor paraphrasing using a vector space model* In *Proceedings of COLING 2012*, Mumbai, India

Tomek Strzalkowski, George Aaron Broadwell, Sarah Taylor, Laurie Feldman, Boris Yamrom, Samira Shaikh, Ting Liu, Kit Cho, Umit Boz, Ignacio Cases and Kyle Elliott. 2013. Robust extraction of metaphor from novel data. In *Proceedings of Workshop on Metaphor in NLP, NAACL.* Atlanta.

Tomek Strzalkowski, Samira Shaikh, Kit Cho, George Aaron Broadwell, Laurie Feldman, Sarah Taylor, Boris Yamrom, Ting Liu, Ignacio Cases, Yuliya Peshkova and Kyle Elliot. 2014. Computing Affect in Metaphors. In *Proceedings of the Second Workshop on Metaphor in NLP,* Baltimore Maryland.

Sarah Taylor, Laurie Beth Feldman, Kit Cho, Samira Shaikh, Ignacio Cases,Yuliya Peshkiva, George Aaron Broadwell Ting Liu, Umit Boz, Kyle Elliott. Boris Yamrom, and Tomek Strzalkowski. 2014. Extracting Understanding from automated metaphor identification: Contrasting Concepts of Poverty across Cultures and Languages. *AHFE Conference*, Cracow, Poland.

Yorick, Wilks. 1975. Preference semantics. *Formal Semantics of Natural Language*, E. L. Keenan, Ed. Cambridge University Press, Cambridge, U.K., 329–348.

Yorick Wilks, Lucian Galescu, James Allen, Adam Dalton. 2013. Automatic Metaphor Detection using Large-Scale Lexical Resources and Conventional Metaphor Extraction. In the *Proceedings of the First Workshop on Metaphor in NLP, (NAACL)*. Atlanta.

Wilson, M. D. 1988. The MRC Psycholinguistic Database: Machine Readable Dictionary, Version 2. *Behavioural Research Methods, Instruments and Computers*, 20(1): 6-11.