

# Generating Conceptual Metaphors from Proposition Stores

Ekaterina Ovchinnikova\*, Vladimir Zaytsev\*, Suzanne Wertheim<sup>+</sup>, Ross Israel\*

\* USC ISI, 4676 Admiralty Way, CA 90292, USA

{katya, vzaytsev, israel}@isi.edu

<sup>+</sup> Worthwhile Research & Consulting, 430 1/2 N Genesee Av., Los Angeles, CA 90036, USA

worthwhileresearch@gmail.com

## Abstract

Contemporary research on computational processing of linguistic metaphors is divided into two main branches: metaphor recognition and metaphor interpretation. We take a different line of research and present an automated method for generating conceptual metaphors from linguistic data. Given the generated conceptual metaphors, we find corresponding linguistic metaphors in corpora. In this paper, we describe our approach and its evaluation using English and Russian data.

## 1 Introduction

The term *conceptual metaphor* refers to the understanding of one concept or conceptual domain in terms of the properties of another (Lakoff and Johnson, 1980; Lakoff, 1987). For example, development can be understood as movement (e.g., *the economy moves forward*, *the engine of the economy*). In other words, a conceptual metaphor consists in mapping *source* conceptual domain (e.g., vehicle) properties to a target domain (e.g., economy). For example, an economy develops like a vehicle moves. In language, conceptual metaphors are expressed by *linguistic metaphors*, i.e. natural language phrases expressing the implied mapping of two domains.

Contemporary research on computational processing of linguistic metaphors is divided into two main branches: 1) metaphor recognition, i.e. distinguishing between literal and metaphorical use of language, and 2) metaphor interpretation, inferring the intended literal meaning of a metaphorical expression; see (Shutova, 2010b) for an overview.

In this paper, we take a different line of research and present an automated method for generating conceptual metaphors from linguistic data. Given a *proposition store*, i.e. a collection of weighted

tuples of words that have a determined pattern of syntactic relations among them, for each lexically expressed concept we collect a weighted list of propositions that represent properties of this concept. For each target concept, we generate a weighted list of possible sources having properties similar to the properties of the target. Given a generated conceptual metaphor, we automatically find linguistic metaphors that are realizations of this conceptual metaphor in a corpus of texts. The proposed method does not crucially rely on manually coded lexical-semantic resources and does not require an annotated training set. It provides a mechanism that can be used for both metaphor recognition and interpretation.

We test our approach using corpora in two languages: English and Russian. We select target concepts and generate potential sources for them. For top-ranked sources, we automatically find corresponding linguistic metaphors. These linguistic metaphors are then validated by expert linguists. In addition, we evaluate proposed conceptual metaphors generated for English against the Master Metaphor List (Lakoff et al., 1991)

## 2 Related Work

Contemporary research on computational processing of linguistic metaphors (LMs) mainly focuses on metaphor recognition and metaphor interpretation. Automatic metaphor recognition is performed by checking selectional preference violation (Fass, 1991; Krishnakumaran and Zhu, 2007; Baumer et al., 2010), relying on existing lexical-semantic resources (Peters and Peters, 2000; Wilks et al., 2013), mapping phrases to predefined target and source concepts constituting a conceptual metaphor (Heintz et al., 2013; Mohler et al., 2013). mapping linguistic phrases to predefined concepts constituting a conceptual metaphor (Heintz et al., 2013; Mohler et al., 2013), and employing a classifier (Birke and Sarkar, 2006;

Gedigian et al., 2006; Baumer et al., 2010; Hovy et al., 2013; Mohler et al., 2013). Approaches based on selectional preferences are known to overgeneralize and miss conventional metaphors (Shutova, 2010b). Approaches relying on hand-coded knowledge and manually annotated training sets have coverage limitations.

Automatic metaphor interpretation is performed using two principal approaches: 1) reasoning with manually coded knowledge of source domains (Narayanan, 1999; Barnden and Lee, 2002; Agerri et al., 2007; Ovchinnikova et al., 2014) and 2) deriving literal paraphrases for metaphorical expressions from corpora (Shutova, 2010a). The first approach has yielded only limited scale, while the second approach does not explain the target-source mapping.

There has been considerably less work on generating conceptual metaphors (CMs). The *CorMet* system (Mason, 2004) generates CMs given predefined target and source domains. Given two domain-specific corpora and two sets of characteristic keywords for each domain, *CorMet* finds domain-characteristic verbs with the highest relative frequency that have the keywords as their arguments. For each verb, the system learns domain-specific selectional preferences represented by WordNet synsets expressing domain concepts. All possible mappings between the top target and source domain concepts are then scored according to polarity (i.e., structure transfer between domains), the number of verbs instantiating the mapping, and the systematic co-occurrence of that mapping with different mappings. For example, given the LAB and FINANCE domains, *CorMet* finds a CM "Money is a Liquid". In contrast to *CorMet*, our method does not require target and source domains to be predefined and does not rely on selectional preferences.

Gandy et al. (2013) propose another method for generating CMs. Given predefined target nouns, the method classifies verb and adjective phrases (called *facets*) containing these nouns as being either metaphorical or not based on the assumption that a metaphor usually involves a mapping from a concrete to an abstract domain. For generating source candidates, the method finds nouns that are non-metaphorically associated with the same facets in the corpus. Source and target nouns are clustered into concepts represented by WordNet synsets. The obtained target-source concept pairs

are filtered and scored according to several heuristics. This method is restricted to metaphorical mappings from concrete to abstract domains, e.g. "Power is a Building", whereas our method does not imply such limitation.

Our method is similar to (Mason, 2004) and (Gandy et al., 2013) in that we rely on characteristic predicate, but we use general propositions instead of verb and adjective phrases, see Sec. 3.

### 3 Building Proposition Stores

We build upon the idea that by parsing a sentence and abstracting from its syntactic structure (e.g., dropping modifiers) we can obtain common sense knowledge, see (Schubert, 2002; Clark and Harrison, 2009). For example, the sentence *Powerful summer storms left extensive damage in California* reveals common sense knowledge about storms possibly leaving damage and being powerful. This knowledge can be captured by *propositions*, i.e. tuples of words that have a determined pattern of syntactic relations among them (Clark and Harrison, 2009; Peñas and Hovy, 2010; Tsao and Wible, 2013). While many of such tuples can be erroneous due to parse errors, statistically higher frequency tuples can be considered more reliable.

We generated propositions from parsed English and Russian corpora. We parsed English Gigaword (Parker et al., 2011) with *Boxer* (Bos et al., 2004). As one of the possible formats, *Boxer* outputs logical forms of sentences in the style of (Hobbs, 1985), generalizes over some syntactic constructions (e.g., passive/active), and performs binding of arguments. For example, *Boxer* represents the sentence *John decided to go to school* as:

$$\text{John}(e_1, x_1) \wedge \text{decide}(e_2, x_1, e_3) \wedge \text{go}(e_3, x_1) \\ \wedge \text{to}(e_4, e_3, x_2) \wedge \text{school}(e_5, x_2)$$

The following propositions can be extracted from this output:

(NV *John decide*)  
 (NV *John go*)  
 (NVV *John decide go*)  
 (VPN *go to school*)  
 (NVPN *John go to school*)  
 (NVVPN *John decide go to school*)

A *proposition store* is a collection of such tuples such that each tuple is assigned its frequency in a corpus. For Russian, we used the ruWac corpus (Sharoff and Nivre, 2011) parsed with the *Malt* dependency parser (Nivre et al., 2006). We then con-

vert these dependency graphs into logical forms in the style of (Hobbs, 1985).

#### 4 Generalizing over Propositions

A significant amount of the propositions can be further generalized if we abstract from named entities, synonyms, and sister terms. Consider the following tuples:

(*Guardian publish interview with Stevens*)  
 (*newspaper publish interview with John*)  
 (*journal publish interview with Dr. Crick*)

The first two propositions above provide evidence for generating the proposition (*newspaper publish interview with person*). All three propositions above can be generalized into (*periodical publish interview with person*). Such generalizations help to refine frequencies assigned to propositions containing abstract nouns, infer new propositions, and cluster propositions.

In order to obtain the generalizations, we first map nouns contained in the propositions into WordNet and Wikipedia semantic nodes using the YAGO ontology (Suchanek et al., 2007). YAGO links lexical items to WordNet and Wikipedia for all the languages presented in Wikipedia including English and Russian. Given a single noun  $n$  being an argument in a tuple, the mapping procedure works as follows.

1. Find semantic nodes  $N$  corresponding to  $n$ 
  - If  $n$  is a given name or a surname, then  $N$  is  $\langle \text{wordnet\_person} \rangle$ .
  - If  $n$  is equal to a YAGO lexical item, then  $N$  is the corresponding YAGO semantic node.
  - If  $n$  is a substring of several YAGO lexical items, then  $N$  is a union of all corresponding YAGO semantic nodes.
2. Filter semantic nodes  $N$ 
  - If  $N$  contains class nodes only, the class nodes are returned. We do not perform disambiguation and map ambiguous nouns to all possible semantic nodes.
  - If  $N$  contains both class and instance nodes, the class nodes are returned (e.g., *nirvana* can be mapped to the class  $\langle \text{wordnet\_nirvana} \rangle$  and to the instance  $\langle \text{Nirvana\_band} \rangle$ ).
  - If  $N$  contains instance nodes only, then the instance nodes are mapped to the corresponding classes using the YAGO hierarchy.

Noun compounds require special treatment. We map the longest sequence of lexemes in each

compound to semantic nodes. For example, for *the New York Times* we obtain the node  $\langle \text{The\_New\_York\_Times} \rangle$  instead of the nodes  $\langle \text{New\_York} \rangle$  and  $\langle \text{wordnet\_time} \rangle$ . If different parts of the compound are mapped to different nodes, we prefer class nodes over instances. For example, parts of the compound *musician Peter Gabriel* are mapped to  $\langle \text{wordnet\_musician} \rangle$  as well as to  $\langle \text{Peter\_Gabriel} \rangle$ . We prefer the node  $\langle \text{wordnet\_musician} \rangle$ , because it refers to a class. If several classes can be produced for a noun compound, the mapping procedure returns all of them. Given the obtained mappings to YAGO classes, we merge identical propositions and sum their frequencies. For example, given two propositions (*live in city*, 10) and (*live in New York*, 5), we obtain (*live in*  $\langle \text{wordnet\_city} \rangle$ , 15).

The described simple mapping procedure has obvious limitations, especially concerning the preference of classes over instances. However, the procedure is computationally cheap and does not require a training dataset. In future work we plan to compare its performance with the performance of the *Wikifier* system (Ratinov et al., 2011) that identifies important entities and concepts in text, disambiguates them and links them to Wikipedia.

#### 5 Generating Salient Concept Properties

In a metaphor, properties of a source domain are mapped to a target domain. In natural language, concepts and properties are represented by words and phrases. There is a long-standing tradition for considering computational models derived from word co-occurrence statistics as being capable of producing reasonable property-based descriptions of concepts; see (Baroni and Lenci, 2008) for an overview. We use the proposition stores to derive salient properties of concepts that can be potentially mapped in a metaphor.

Given a seed lexeme  $l$ , we extract all tuples from the proposition store that contain that lexeme. For each tuple  $t = \langle x_1, \dots, x_n \rangle$  there is a set of patterns of the form  $p_i(t) = \langle x_1, \dots, x_{i-1}, \_, x_{i+1}, \dots, x_n \rangle$ , i.e. a pattern  $p_i$  of tuple  $t$  is obtained by replacing the  $i$ th argument of  $t$  with a blank value. The weight of tuple  $t$  relative to lexeme  $l$  occurring at position  $i$  in  $t$  is computed as follows:

$$\text{weight}_l(t) = \frac{\text{freq}(t)}{\sum_{t' \in T: p_i(t) = p_i(t')} \text{freq}(t')}, \quad (1)$$

where  $T$  is a set of all tuples, and  $freq(t)$  is a frequency of tuple  $t$ . For example, the following are the top-ranked tuples containing the English lexeme *poverty*:

(NVPN majority live in poverty)  
(NVAdv poverty affect increasingly)  
(NPN lift out of poverty)  
(VN fight poverty)  
(VN eliminate poverty)  
(NN poverty eradication)  
(AdvPN deep in poverty)

## 6 Generating Source Lexemes

Some property tuples presented in Sec. 5 already suggest conceptual metaphors: "Poverty is a Location" (*live in poverty*), "Poverty is an Enemy" (*fight poverty*), "Poverty is an Abyss" (*deep in poverty*). We generate potential source lexemes for a seed target lexeme  $l$  in three steps:

1. Find all tuples  $T_l$  containing  $l$ .
2. Find all potential source lexemes  $S$  such that for each  $s \in S$  there are tuples  $t, t'$  in the proposition store such that  $l$  occurs at position  $i$  in  $t$  and  $s$  occurs at position  $i$  in  $t'$ . The set of tuples containing  $l$  and  $s$  at the same positions is denoted by  $T_{l,s}$ .
3. Weight potential source lexemes  $s \in S$  using the following equation:

$$weight_l(s) = \sum_{t \in T_{l,s}} weight_l(t), \quad (2)$$

We generated potential source lexemes for the target domains of Poverty, Wealth, and Friendship. For each target domain, we selected two seed target lexemes with different parts of speech. In English, we selected *poverty* and *poor*, *wealth* and *buy*, *friendship* and *friendly*. In Russian, we selected бедность (poverty) and бедный (poor), богатство (wealth) and покупать (buy), дружба (friendship) and дружеский (friendly). The top-ranked source lexemes for each seed target are presented in Tables 1 and 2.

Nouns seem to be better seeds as compared to adjectives and verbs. Looking at the source lexemes for nouns, we find a) semantically related words (e.g., *poverty*: *emission*, *recession*, 2) abstract supercategories (e.g., *poverty*: *situation*), or 3) real potential sources (e.g., *poverty*: *disease*).

The list of source lexemes for verbs and adjectives contains many random frequent words (e.g., *buy*: *set*, *poor*: *first*). One possible explanation is that the selected seed adjectives and verbs occur in more contexts and combine more freely than selected nouns. A larger study over a wide range of targets is needed to draw further conclusions.

Obviously, many of the found source lexemes are semantically related to the target lexemes. In order to find sources that share patterns with the target, but are not closely semantically related, we need to compute "anti-relatedness". For doing so, we use the Latent Dirichlet allocation (LDA) model (Blei et al., 2003).<sup>1</sup> LDA is a probabilistic topic model, in which documents are viewed as distributions over topics and topics are viewed as distributions over words. We generated English and Russian LDA models using the Gensim toolkit<sup>2</sup> applied to lemmatized English Gigaword and ruWac corpora with stop words removed; the number of topics was equal to 50.

Following (Rus et al., 2013), we define the LDA-based relatedness of words as follows:

$$rel(w_1, w_2) = \sum_{t=1}^T \phi_t(w_1)\phi_t(w_2), \quad (3)$$

where  $T$  is the number of topics and  $\phi_t(w)$  is the probability of word  $w$  in topic  $t$ . We filter out source lexemes such that their relatedness to the target lexeme is above a threshold. The value of the threshold defines the number of final source lexemes. The results in Tables 3, 4, 5, 6, 7 were obtained with a threshold of 0.04.

The results in the tables show that some of the original source lexemes (Tables 1 and 2) that are semantically close to the target were filtered out, e.g., *poverty*: *corruption*, *recession*, *situation*, *problem*. At the same time, some of the filtered out sources that are semantically similar to the target still have potential to constitute a CM. For example, "Poverty is a Crime" has realizations in the corpus, e.g., *a powerful indictment of the iniquities of racial discrimination and the crime of poverty*.

## 7 Generating Conceptual Metaphors

A conceptual metaphor is a triple  $\langle C_t, C_s, P \rangle$  consisting of a target concept  $C_t$ , a source concept  $C_s$ ,

<sup>1</sup>We also experimented with the Latent Semantic Analysis model (Dumais, 2004), but the LDA model proved to provide more relevant results.

<sup>2</sup><http://radimrehurek.com/gensim/>

<b>poverty</b>	<b>poor</b>	<b>wealth</b>	<b>buy</b>	<b>friendship</b>	<b>friendly</b>
terrorism	economic	resource	hold	cooperation	bilateral
area	local	power	win	tie	first
disease	low	security	provide	country	political
corruption	strong	lot	rise	peace	international
problem	public	support	produce	development	positive
violence	first	gas	build	nation	military
situation	international	hundred	wear	trade	european
emission	bad	information	include	agreement	public
recession	national	percent	set	understanding	close
crime	political	level	import	deal	official

Table 1: Source lexemes for Poverty, Wealth, and Friendship generated from English Gigaword.

бедность (poverty)	бедный (poor)	богатство (money)	покупать (buy)	дружба (friendship)	дружеский (friendly)
отсутствие (absence)	маленький (small)	сила/энергия (power/energy)	купить (buy)	отношение (relationship)	последний (last)
проблема (problem)	молодой (young)	количество/сумма/ объем (amount/sum)	использовать (use)	встреча (encounter)	личный (personal)
боль/страдание (pain/suffering)	русский (Russian)	опыт (experience)	получить (get)	война (war)	взаимный (mutual)
чувство (feeling)	дорогой (expensive)	возможность (possibility)	найти (find)	любовь (love)	особый (special)
состояние (state)	местный (local)	средство (means)	сделать (make)	связь (tie)	деловой (business)
факт (fact)	бывший (former)	мир (world)	брать (take)	общение (communication)	человеческий (human)
мир (world)	старый (old)	работа (work)	приобрести (purchase)	работа (work)	подобный (similar)
нарушение (violation)	простой (simple)	любовь (love)	любить (love)	граница (border)	небольшой (small)
страна/город (country/city)	настоящий (real)	информация (information)	продавать (sell)	интервью (interview)	легкий (light)
болезнь/заболе- вание(disease)	уважаемый (respected)	знание (knowledge)	работать (work)	знакомство (acquaintance)	общий (common)

Table 2: Source lexemes for Бедность, Богатство, and Дружба generated from Russian ruWac.

and a set  $P$  of properties transferred from  $C_s$  to  $C_t$ . Each concept is a set of words and phrases. In order to obtain potential source concepts, we cluster generated source lexemes.

Both (Mason, 2004) and (Gandy et al., 2013) employ the WordNet hierarchy for clustering. In this paper, we take a similar approach and use the YAGO hierarchy based on WordNet and Wikipedia.<sup>3</sup> Source lexemes belong to the same concept if they are all hyponyms of the same YAGO node in the hierarchy and share  $k$  or more patterns. The value of  $k$  defines the number of final source concepts. The results in Tables 3, 4, 5, 6, 7 were obtained with  $k = 5$ . The weight of the potential source concept is equal to the sum of the weights of the corresponding source lexemes. For example, *area*, *room*, *apartment*, *city*, *region* were clustered together based on the patterns (*live in X*), (*reside in X*), etc.

<sup>3</sup>Note that a preliminary clustering of synonyms was done at the proposition generalization step (Sec. 4).

Tables 3 and 4 contain detailed results for the targets Wealth and Бедность (Poverty) showing some of the patterns common for the targets and the sources that can be used to explain the conceptual metaphor. For English, "Wealth is Blood", because one can donate and pump both money and blood. For Russian, "Бедность это Враг" ("Poverty is an Enemy"), is based on the idea that one can fight against poverty, defeat it, etc.

The usage of the YAGO hierarchy is just one out of many options. In the future, we will investigate different word clustering algorithms (including distributional models, n-gram-based language models, etc.) and their effect on the obtained sources.

## 8 Finding Linguistic Metaphors

For each potential conceptual metaphor, we look for supporting linguistic metaphors in corpora. A large number of LMs supporting a particular CM suggests that this CM might be cognitively plausi-

ble. However, it should be noted that if a CM is not supported by any LMs it does not mean that this CM is wrong. The target-source mapping can be still cognitively relevant, but not yet conventional enough to be represented linguistically.

We use a simple method for finding LMs. If a target lexeme and a source lexeme are connected by a dependency relation in a sentence, then we assume that this dependency structure contains a LM. For example, in the phrases *medicine against poverty*, *chronic poverty*, бедность – это болезнь, которую надо лечить ("poverty is a disease that should be treated"), болезнь хронической бедности ("disease of chronic poverty") target words (poverty, бедность) are related by a dependency with source words (medicine, chronic, болезнь, хронический). Heintz et al. (2013) present a similar approach mapping sentences to LDA topic models for target and source domains. Our method allows us to exploit dependency links and output only sentences containing target words being modified by source words or vice versa.

This method has limitations with respect to both precision and recall. First, not all LMs are accommodated in a dependency structure. For example, in the text fragment *There is no "magic bullet" for poverty, no cure-all* the target word *poverty* is not related to the source word *cure-all* by a dependency link. Second, this method over-generates. In the previous section, there was an example of the conceptual metaphor "Poverty is a Location". While *poor* is a target lexeme and *country* is a source lexeme for this CM, the phrase *poor country* is not metaphorical. One more problem concerns ambiguity. The phrase *friendship is a two-way deal* instantiating the "trade/business" meaning of *deal* seems to contain a LM, whereas *personal friendship is a big deal* instantiating the meaning "important" is not metaphorical.

Tables 3 and 4 contain examples of the extracted sentences potentially containing LMs for the CMs generated for the targets Wealth and Бедность (Poverty) along with the patterns explaining the target-source mapping.

## 9 Evaluation

In this section, we describe an evaluation of the proposed approach. First, we present a validation of the generated CMs. As mentioned in Sec. 7, nouns proved to be the best seeds for generating potential sources. For two of the selected

noun targets per language (*poverty*, *wealth*, бедность (poverty), богатство (wealth)) we generate source lexemes, select 100 top-ranked lexemes, cluster them into source concepts, and select 10 top-ranked CM proposals (Tables 5 and 6).

In order to obtain a lexically richer representation of the domains, we expand the sets of these target and source lexemes with semantically related lexemes using English and Russian ConceptNet resources (Speer and Havasi, 2013) and top ranked patterns from the proposition stores.<sup>4</sup> For example, the expansion of the lexeme *disease* results in {*disease*, *symptom*, *syndrome*, *illness*, *unwellness*, *sickness*, *sick*, *medicine*, *treatment*, *treat*, *cure*, *doctor*, ... }.

Given parsed English Gigaword and ruWac corpora, we extracted sentences that contain dependency structures relating target and source lexemes. For each language, we randomly selected at most 10 sentences per target-source pair. For some pairs, less than 10 sentences were retrieved. In total, we obtained 197 sentences for English and 186 for Russian. Each sentence was validated by three linguist experts. The experts were asked if the sentence contains a metaphor mapping indicated target and source domains. The Fleiss' kappa (Fleiss, 1971) is 0.69 for English and 0.68 for Russian.

Tables 5 and 6 show potential sources with numbers of the corresponding sentences containing a LM. Column ALL contains the number of sentences per a proposed CM such that all three experts agreed that there is a metaphor in these sentences. Column TWO contains the number of sentences such that any two experts agreed that there is a metaphor in them. Column ONE contains the number of sentences such that only one expert thought there is a metaphor.

We consider a CM to be approved if at least one of the associated LMs was positively validated by all three experts. According to the validation results in Tables 5 and 6, 18 out of 20 top-ranked conceptual metaphors proved to be promising for English (90%). For Russian, the experts approved 15 of 20 proposed CMs (75%).

In several cases, experts validated given sentences as containing LMs, but disagreed with the label of the corresponding source domain. For ex-

<sup>4</sup>ConceptNet combines several resources developed manually (e.g., WordNet, Wiktionary) and thus provides high quality semantic relations. The usage of it is optional. ConceptNet semantic relations can be replaced by semantically similar words provided by a distributional model.

ample, the experts complained about "Poverty is a Slump" as being a label for the examples like *Getting a college degree does not assure one will lift out of poverty*. The source lexeme *slump* generates many relevant patterns, e.g., *lift out of X*, *deep X*, *pull out of X*, etc., but Abyss seems to be a better label for this domain.

In some cases two concepts constituting a CM are mapped to the same source domain. For example, both Poverty and Terrorism are mapped to Enemy. Our method produces the CM "Poverty is Terrorism" and many "enemy"-patterns as overlapping properties (*fight against X*, *war on X*). Does it mean that the method overgenerates and "Poverty is Terrorism" is not a valid CM? In our experimental study we find that sentences like *Poverty is a form of terrorism, causing its victims to live in fear* are validated by the experts as metaphorical. Thus, even if  $D_1$  and  $D_2$  are mapped to the same domain,  $\langle D_1, D_2, P \rangle$  might still be a valid CM.

In order to compare our system results to the *CorMet* system (Mason, 2004), we present an evaluation against the Master Metaphor List (Lakoff et al., 1991). The Master Metaphor List (MML) is composed of manually verified metaphors common in English. We restrict our evaluation to the elements of MML used for evaluating *CorMet* (Table 7). For each target and source domain in Table 7, we select English seed lexemes and expand the sets using ConceptNet and top-ranked patterns from the proposition store. For target nouns, we generate potential sources. For each target lexeme, we take a set of 100 top-ranked potential sources and check if the set contains MML source lexemes. For example, for the CM "Fighting a War is Treating Illness", we obtain the lexeme sets  $T = \{\textit{war, fight, combat, battle, attack, ...}\}$  and  $S = \{\textit{disease, treatment, medicine, doctor, ...}\}$ . Mappings found between  $T$  and  $S$  are shown in Table 7. For example, *treatment* was mapped to *attack* with weight 0.51.<sup>5</sup> The table also contains the original *CorMet* mappings and scores. Similar to *CorMet*, our system found reasonable CMs in 10 of 13 cases (77%).<sup>6</sup>

We used LMs associated with CMs in the Master Metaphor List for error analysis. We found that

<sup>5</sup>Weights were scaled between 0 and 1.

<sup>6</sup>Note that Mason (2004) evaluates the correspondences between the *CorMet* mappings and the MML mappings by hand which introduces subjectivity, whereas our evaluation is done automatically.

target	source	ALL	TWO	ONE
wealth	blood	10	10	10
	water	9	10	10
	drug	9	10	10
	food	9	9	10
	power	8	9	10
	security	7	9	10
	resource	7	7	9
	victory	2	3	5
	support	0	1	1
	troops	0	0	0
poverty	war	10	10	10
	slump	10	10	10
	violence	9	9	10
	price/cost	8	9	9
	house/area/ country/...	7	9	9
	disease	7	7	7
	crop	3	7	9
	terrorism	3	3	5
	tension	0	9	10
	crisis	0	9	10
Total (percentage)		118 (.6)	149 (.76)	163 (.83)

Table 5: Validation of English linguistic metaphors found for potential sources.

for two missing CMs, our system produced mappings, but they were assigned low weights because of the low frequencies of the corresponding propositions in the corpus. Therefore they were not included into 100 top-ranked source proposals. For example, for the CM "Investments as Containers for Money" (*The bottom of the economy dropped out, I'm down to my bottom dollar*), the system found the pattern (*NPN bottom of X*) for both *economy* and *container*. For the CM "People as Containers for Emotions" (*I was filled with rage, She could hardly contain her anger*), it found the propositions (*VPN fill with emotion/feeling*), (*NVN people contain feeling*) and the patterns (*VPN fill with X*) and (*NVN X contain Y*) for both *people* and *container*. For "People as Machines" (*He had a breakdown, what makes him tick, Fuel up with a good breakfast*), no mappings were found.

The presented evaluation shows how many of the proposed top-ranked CMs are approved by experts and how many of the MML CMs are found by the system, but we do not learn much about the quality of the CM ranking measure. In the future we plan to evaluate how well our ranking of CMs correlates with human ranking.

## 10 Conclusion and Future Work

Contemporary research on computational processing of linguistic metaphors is divided into two main branches: metaphor recognition and

T	source	patterns	example
wealth	blood water	people donate X, pump X, frozen X pool of X, X flow, drain X	Money is the blood of the state and must circulate Of course , their rivers of money mean they can offer far more than a single vote
	drug	injection of X, stash X, distribute X	Money is a drug, and a drug addiction can make any of us lower our standards
	food	taste of X, eat up X , taste X	We would go visit him at his mansions and have a taste of wealth and privilege
	power security	temptation of X, transfer of X, misuse X ensure X, social X, control of X	But its enormous wealth rules the world today Rich people, however, imagine that their wealth protects them like high, strong walls round a city
	resource victory support	plunder X, abundant X, sharing of X campaign with X, symbol of X, ensure X government need X, lend X, X for program	Money is a resource, not a reward Consequently, the attainment of wealth is often a hollow victory Money is a support system, and indeed can be seen as a form of trust
	troops	withdraw X, shell X, contribute X	Many of the treasures were later auctioned to raise money for the troops

Table 3: Examples of English LMs found for potential sources for the target Wealth.

T	source	patterns	example
бедность (poverty)	пропасть (abyss) враг (enemy)	упасть в X (fall into X), глубокий X (deep X), скатиться в X (slide into X) победа над X (victory over X), борьба против X (fight against X), война с X (war against X)	Огромное количество людей скатывается в бедность (Many people slide into poverty) Боритесь с бедностью, бедность – это враг всех людей (Fight against poverty, poverty is the enemy of all people)
	болезнь/ заболевание (disease)	хронический X (chronic X), страдать от X (suffer from X), переносить X (stand X)	путь преодоления "болезни хронической бедности"(a way to overcome the "disease of chronic poverty")
	власть (power)	господство X (domination of X), X править (X rule), X господствовать (X dominate)	Ввиду страшно высоких налогов везде господствовала ужасающая нищета. (Because of the high taxes terrible poverty ruled everywhere.)
	страдание/боль (suffering/pain)	жаловаться на X (complain about X), X измучить (X frazzle), избавление от X (deliverance from X)	Бедность - это стыд, боль, паника, отчаяние (Poverty is a shame, pain, panic, despair)
	смерть (death)	обречь на X (doom to X), преодолеть X (overcome X), христос победить X (Christ defeat X)	в деревне смертельная нищета была очевидна (the deadly poverty in the village was evident)
	нагота (nakedness)	стыдиться X (be ashamed of X), скрывать X (hide X), увидеть X (see X)	нагота бедности кое-где прикрыта зеленью (the nudity of poverty is covered in some places)
	страна/город/район/... (country/city/region/...)	большинство жить в X (majority live in X), родиться в X (be born in X), вырасти в X (grow up in X)	Он читает рабочим проповеди из Библии, продолжая жить в полной нищете (He reads the Bible to the workers and continues to live in abject poverty)
	нарушение (violation)	вопиющий X (egregious X), решение устранить X (decision to eliminate X), протест против X (protest against X)	бедность – нарушение прав человека (poverty is a violation of human rights)
чувство (feeling)	унизительный X (humiliating X), стыдиться X (be ashamed of X), ужасный X (terrible X)	неспособность любить, бедность чувств (inability to love, poverty of feelings)	

Table 4: Examples of Russian LMs found for potential sources for the target Бедность (Poverty).



Master Metaphor List mapping	CorMet mapping	Proposed system mapping
Fortifications→Theories	none	base→theory (.49), cement→theory (.48)
Fluid→Emotion	liquid-1→feeling-1 (.25)	channel→feeling (.57), water→emotion (0.69)
Containers for Emotions → People	container-1→person-1 (.13)	none
War→Love	military unit-1→feeling-1 (.34)	kill→love (.21), fight/combat/battle→feeling (.57), fight→emotion (.99)
Injuries→Effects of Humor	weapon-1 →joke-1 (.18)	crack→joke (.5), cut→joke (.5), wound→joke (.47)
Fighting a War→Treating Illness	military action-1 → medical care-1 (.4)	attack→treatment (.51), defence→doctor(.5), fight→disease (1.0), defeat→disease (.24)
Journey→Love	travel-1→feeling-1 (.17)	move→feeling (1.57), train→feeling (.56), transfer→emotion (.99)
Physical Injury→Economic Harm	harm-1→loss-3 (.2)	cut→loss (.47), suffer→loss (.18), strike→loss (.13), hit→recession (.64), hurt→recession (.58)
Machines → People	none	none
Liquid → Money	liquid-1 → income-1 (.56)	water→money (.99), channel→money (.33), flow→dollar (.18)
Containers for Money → Investments	container-1 → institution-1 (.35)	none
Buildings → Bodies	none	house→body (.41)
Body → Society	body part-1 → organization-1 (.14)	head→organization (.54), arm→organization (.34), face→organization (.34), face→nation (.76)

Table 7: Evaluation against Master Metaphor List and *CorMet* results.

T	source	ALL	TWO	ONE
богатство (wealth)	энергия/сила (energy/force)	10	10	10
	вода (water)	10	10	10
	власть (power)	9	10	10
	бог (god)	9	10	10
	игра (game)	8	10	10
	материал (material)	3	4	4
	право (right)	1	1	1
	ребенок (child)	1	1	1
	страна/город/район/... (country/city/region/...)	0	10	10
	информация (information)	0	0	0
бедность (poverty)	пропасть (abyss)	10	10	10
	враг (enemy)	9	10	10
	болезнь/заболевание (disease)	9	9	9
	власть (power)	8	10	10
	страдание/боль (suffering/pain)	5	10	10
	смерть (death)	3	5	6
	нагота (nakedness)	1	3	4
	страна/город/район/... (country/city/region/...)	0	10	10
	нарушение (violation)	0	4	4
	чувство (feeling)	0	0	4
Total (percentage)	96 (.52)	135 (0.73)	141 (0.76)	

Table 6: Validation of Russian linguistic metaphors found for potential sources.

metaphor interpretation. In this paper, we take a different line of research and present an automated method for generating conceptual metaphors from linguistic data. The proposed approach does not crucially rely on manually coded resources and does not require an annotated training set. It provides a mechanism that can be used for metaphor recognition given predefined target and source domains (see Sec. 8). It also enables metaphor interpretation through patterns shared by target and source lexemes (cf. Tables 3 and 4). All developed tools, generated resources, and validation data are freely available for the community as an open source project at <http://ovchinnikova.me/proj/metaphor.html>.

In the future, we will investigate different word clustering algorithms and their effect on the obtained CMs. We also aim at providing a more solid evaluation of the CM ranking score and compute how well it correlates with human ranking. Furthermore, we will study how well salient properties shared by targets and sources help to explain conceptual metaphors. For doing so, we will need to create a gold standard by asking human subjects to explain given conceptual metaphors and provide properties mapped from source to target domains.

## References

R. Agerri, J.A. Barnden, M.G. Lee, and A.M. Wallington. 2007. Metaphor, inference and domain-

- independent mappings. In *Proc. of RANLP'07*, pages 17–23.
- J. A. Barnden and M. G. Lee. 2002. An artificial intelligence approach to metaphor understanding. *Theoria et Historia Scientiarum*, 6(1):399–412.
- M. Baroni and A. Lenci. 2008. Concepts and properties in word spaces. *Italian Journal of Linguistics*, 20(1):55–88.
- E. PS. Baumer, J. P. White, and B. Tomlinson. 2010. Comparing semantic role labeling with typed dependency parsing in computational metaphor identification. In *Proc. of the NAACL HLT 2010 Second Workshop on Computational Approaches to Linguistic Creativity*, pages 14–22. ACL.
- J. Birke and A. Sarkar. 2006. A clustering approach for nearly unsupervised recognition of nonliteral language. In *Proc. of EACL'06*.
- D. M. Blei, A. Ng, and M. Jordan. 2003. Latent dirichlet allocation. *JMLR*, 3:993–1022.
- J. Bos, S. Clark, M. Steedman, J. R. Curran, and J. Hockenmaier. 2004. Wide-coverage semantic representations from a ccg parser. In *Proc. of COLING'04*, pages 1240–1246.
- P. Clark and P. Harrison. 2009. Large-scale extraction and use of knowledge from text. In *Proc. of the 5th international conference on Knowledge capture*, pages 153–160. ACM.
- S. T. Dumais. 2004. Latent semantic analysis. *ARIST*, 38(1):188–230.
- D. Fass. 1991. met\*: A method for discriminating metonymy and metaphor by computer. *Computational Linguistics*, 17(1):49–90.
- J. L. Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5):378–382.
- L. Gandy, N. Allan, M. Atallah, O. Frieder, N. Howard, S. Kanareykin, M. Koppel, M. Last, Y. Neuman, and S. Argamon. 2013. Automatic Identification of Conceptual Metaphors With Limited Knowledge. In *Proc. of AAAI'13*.
- M. Gedigian, J. Bryant, S. Narayanan, and B. Ciric. 2006. Catching metaphors. In *Proc. of the Third Workshop on Scalable Natural Language Understanding*, pages 41–48. ACL.
- I. Heintz, R. Gabbard, M. Srinivasan, D. Barner, D. S. Black, M. Freedman, and R. Weischedel. 2013. Automatic extraction of linguistic metaphors with LDA topic modeling. In *Proc. of the First Workshop on Metaphor in NLP*, pages 58–66. ACL.
- J. R. Hobbs. 1985. Ontological promiscuity. In *Proceedings of ACL*, pages 61–69, Chicago, Illinois.
- D. Hovy, S. Srivastava, S. K. Jauhar, M. Sachan, K. Goyal, H. Li, W. Sanders, and Eduard Hovy. 2013. Identifying Metaphorical Word Use with Tree Kernels. In *Proc. of the First Workshop on Metaphor in NLP*, pages 52–57. ACL.
- S. Krishnakumaran and X. Zhu. 2007. Hunting elusive metaphors using lexical resources. In *Proc. of the Workshop on Computational approaches to Figurative Language*, pages 13–20. ACL.
- G. Lakoff and M. Johnson. 1980. *Metaphors we Live by*. University of Chicago Press.
- G. Lakoff, J. Espenson, and A. Schwartz. 1991. The Master Metaphor List. Draft 2nd ed. Technical report, University of California at Berkeley.
- G. Lakoff. 1987. *Women, fire, and dangerous things: what categories reveal about the mind*. University of Chicago Press.
- Z. J. Mason. 2004. CorMet: a computational, corpus-based conventional metaphor extraction system. *Computational Linguistics*, 30(1):23–44.
- M. Mohler, D. Bracewell, D. Hinote, and M. Tomlinson. 2013. Semantic Signatures for Example-Based Linguistic Metaphor Detection. In *Proc. of the First Workshop on Metaphor in NLP*, pages 27–35. ACL.
- S. Narayanan. 1999. Moving right along: A computational model of metaphoric reasoning about events. In *Proc. of AAAI/IAAI*, pages 121–127.
- J. Nivre, J. Hall, and J. Nilsson. 2006. Maltparser: A data-driven parser-generator for dependency parsing. In *Proc. of LREC'06*, volume 6, pages 2216–2219.
- E. Ovchinnikova, R. Israel, S. Wertheim, V. Zaytsev, N. Montazeri, and J. Hobbs. 2014. Abductive Inference for Interpretation of Metaphors. In *Proc. of the Second Workshop on Metaphor in NLP*, page to appear. ACL.
- R. Parker, D. Graff, J. Kong, K. Chen, and K. Maeda. 2011. English gigaword fifth edition. *LDC*.
- A. Peñas and E. H. Hovy. 2010. Filling knowledge gaps in text for machine reading. In *Proc. of COLING'10*, pages 979–987.
- W. Peters and I. Peters. 2000. Lexicalised systematic polysemy in wordnet. In *LREC'00*.
- L. Ratinov, D. Roth, D. Downey, and M. Anderson. 2011. Local and global algorithms for disambiguation to wikipedia. In *Proc. of ACL'11*.
- V. Rus, M. Lintean, R. Banjade, N. Niraula, and Stefanescu D. 2013. SEMILAR: The Semantic Similarity Toolkit. In *Proc. of ACL'13*.
- L. Schubert. 2002. Can we derive general world knowledge from texts? In *Proc. of HLT'03*, pages 94–97. Morgan Kaufmann Publishers Inc.

- S. Sharoff and J. Nivre. 2011. The proper place of men and machines in language technology: Processing Russian without any linguistic knowledge. In *Proc. Dialogue 2011, Russian Conference on Computational Linguistics*.
- E. Shutova. 2010a. Automatic metaphor interpretation as a paraphrasing task. In *Proc. of NAACL'10*.
- E. Shutova. 2010b. Models of metaphor in NLP. In *Proc. of ACL'10*, pages 688–697. ACL.
- R. Speer and C. Havasi. 2013. Conceptnet 5: A large semantic network for relational knowledge. In *The People's Web Meets NLP*, pages 161–176. Springer.
- F. M. Suchanek, G. Kasneci, and G. Weikum. 2007. Yago: A Core of Semantic Knowledge. In *Proc. of the 16th international World Wide Web conference*, NY, USA. ACM Press.
- N. Tsao and D. Wible. 2013. Word similarity using constructions as contextual features. In *Proc. JSSP'13*, pages 51–59.
- Y. Wilks, L. Galescu, J. Allen, and Adam Dalton. 2013. Automatic Metaphor Detection using Large-Scale Lexical Resources and Conventional Metaphor Extraction. In *Proc. of the First Workshop on Metaphor in NLP*, pages 36–44. ACL.