

Marianna Bolognesi* and Laura Aina

Similarity is closeness: Using distributional semantic spaces to model similarity in visual and linguistic metaphors

DOI 10.1515/cllt-2016-0061

Abstract: The semantic similarity that characterizes two terms aligned in a metaphor is here analysed through a corpus-based distributional semantic space. We compare and contrast two samples of metaphors, representative of visual and linguistic modality of expressions respectively. Popular theories of metaphor claim that metaphors transcend their modality to influence conceptual structures, thus suggesting that different modalities of expression would typically express the same conceptual metaphors. However, we show substantial differences in the degree of similarity captured by the distributional semantic space with regard to the modality of expression (higher similarity for linguistic metaphors than for visual ones). We argue that this is due to two possible variables: Conventionality (linguistic metaphors are typically conventional, while visual are not) and Complexity (visual metaphors have modality-specific inner complexities that penalize the degree of similarity between metaphor terms captured by a language-based model). Finally, we compare the similarity scores of our original formulations with those obtained from different possible verbalizations of the same metaphors (acquired by replacing the metaphor terms with their semantic neighbours). We show that while this operation does not affect the average similarity between metaphor terms for visual metaphors, the similarity changes significantly in linguistic metaphors. These results are discussed here.

Keywords: distributional semantics, metaphor analysis, multimodality

***Corresponding author: Marianna Bolognesi**, Argumentation Theory and Rethorics, Universiteit van Amsterdam, Spuistraat 134, Amsterdam 1012VB, Netherlands,
E-mail: M.M.Bolognesi2@uva.nl

Laura Aina, Institute for Logic, Language and Computation, Universiteit van Amsterdam, Amsterdam, Netherlands

1 Introduction

Nowadays, metaphor is recognized as a complex cognitive phenomenon that affects language use and, possibly under certain conditions, (see for example Steen 2013), conceptual processing as well (Lakoff and Johnson 1980).

Through metaphors we understand one entity in terms of another, the latter generally being perceived as easier or more concrete. For example, if a journalist refers to the phenomenon of immigration by defining it as a *tsunami*, then properties of the latter concept (such as its unavoidable destructive power) are meant to be transferred onto the concept *immigration*. The cognitive operations that have been proposed to account for human metaphor understanding include interaction views (e.g. Black 1979), in which the two terms of the metaphor select each other's relevant portions of meaning to generate a metaphorical comparison; blendings (e.g. Turner and Fauconnier 2002) in which new conceptual structures arise from the blend of the compared conceptual domains; and semantic disambiguation (e.g. Giora 2008), in which the literal and the metaphorical meaning of a word are selected or inhibited, depending on the context in which the word appears. All of these models revolve around the idea that a metaphor, in order for it to be meaningful and understandable, requires that the two aligned terms have a degree of similarity to legitimize their metaphorical comparison. However, the nature of such similarity has not yet been investigated.

The current project aims to analyse the semantic similarity between metaphor terms by means of distributional semantics. According to the *distributional hypothesis*, the closer the meanings of two words are, the more these words will tend to occur in the same linguistic contexts (Harris 1954; Firth 1957; Miller and Charles 1991). As a consequence, in the distributional semantics framework, the semantic similarity of two words is computed as a function of the similarity of the linguistic contexts in which the two words are used (Lenci 2008). In particular, starting from matrices of words co-occurrences derived from text corpora, distributional semantic models (DSMs) exploit the geometric distance between two word vectors to weight the semantic similarity between the two words (Deerwester et al. 1990; Turney and Pantel 2010). In this sense, semantic similarity is indicated by geometrical proximity.

We hereby explore to what extent the similarity between two concepts involved in a metaphor matches the distributional similarity of their corresponding words, captured by a distributional semantic model based on text corpora.

A growing body of scientific literature has previously tackled aspects of metaphor comprehension by means of distributional semantics: for example, a

pioneering study conducted by Kintsch (2000) showed how Latent Semantic Analysis (Landauer and Dumais 1997) can be used to model metaphor comprehension in a qualitative fashion. In a more recent and extensive project (Utsumi 2011) categorization and comparison processes involved in metaphor comprehension were compared and modelled through distributional semantics.¹ Moreover, within the computational linguistics and machine learning community, the interest towards statistical modelling of metaphor has recently been growing (Shutova 2015; Veale et al. 2016). These studies typically aim at modelling metaphor structure (rather than the cognitive processes that lead to metaphor comprehension), and metaphor detection in text corpora. Many of these studies indeed make use of distributional semantics methods with the goal of modelling metaphors aspects like relational analogy (Turney 2006), non-compositionality (Vecchi et al. 2011), and verb metaphoricity (Del Tredici and Bel 2016).

In the present study we focus on an unexplored variable: the *modality* in which metaphors are expressed. *Linguistic* and *visual* metaphors are analysed and contrasted, in order to observe potentially different behaviours due to the diverse features and constraints of these two modalities of metaphor expression. The distributional similarity that characterizes two terms aligned in visual vs. linguistic metaphors is used here as a variable to investigate different possible behaviours in which the visual and the linguistic modalities construct and express metaphorical comparisons.

According to Conceptual Metaphor Theory (CMT, Lakoff and Johnson 1980), metaphors *transcend* their modality of expression to influence deeper conceptual structures. It follows that the distributional similarity between metaphor terms should be independent from the modality in which a metaphor is expressed, and possibly the same conceptual metaphors that are typically expressed in language should be found also in images.

In the past decades the research into different modalities of metaphor expression, and in particular the visual modality, has attracted different scientific communities, ranging from scholars in marketing theory (e. g. Phillips and McQuarrie 2004; Ng and Koller 2013) to cognitive psychology (e. g. van weelden et al. 2012) and from linguistics (e. g. Perez-Sobrino 2016) to media studies (e. g. Forceville 2005). Although slightly different working definitions of visual

¹ As described in Bowdle and Gentner (2005) metaphor comprehension might function in two cognitively distinct ways: as an online cross domain comparison, in which the addressee needs to mentally activate both the metaphor terms in order to unravel the metaphor, or as a categorization process, in which only the target domain is activated in the addressee's mind, and perceived as polysemous, whereas the metaphorical and the literal meaning constitute two different senses; the correct sense is then selected in context, through lexical disambiguation.

metaphor have been provided by scholars from different fields, an encompassing working definition can be formulated as follows: prototypical visual metaphors are highly structured images that present perceptually-based incongruities that stimulate the viewers to construct cross-domain mappings to unravel the intended message. Such images are often used in different genres, such as advertising, political cartoons, and art. An example of visual metaphor is given in Figure 1, where a globe (representing the Earth) is fused with a scoop of melting ice-cream. The image is a WWF social campaign about the effects of global warming.



Figure 1: A social campaign issued by WWF about global warming. Author and copyright owner: BBDO Belgium.

The present study is embedded in a larger project² that aims at modelling and contrasting the semantic similarity between metaphor terms in visual vs. linguistic metaphors, relying on the *distributional hypothesis*. Three different types of distributional similarity are investigated through extensive analyses based on corpora that encode different streams of semantic information, and that can therefore model three different types of distributional similarity: semantic

² EU Marie Curie Intra European Fellowship, awarded to dr. Marianna Bolognesi (COGVIM n° 629076 – Project Acronym: COGVIM; Call identifier FP7-PEOPLE-2013-IEF).

features retrieved from mental simulations, elicited in a property generation task; tag co-occurrence across Flickr annotated images; and syntactic patterns plus semantic collocates in text corpora. The present study reports the latter analysis, and tackles the following research questions:

- RQ1: How does distributional similarity based on text corpora capture the similarity between two terms compared in a metaphor?
- RQ2: Does the modality in which a metaphor is expressed influence how distributional similarity models metaphor similarity?
- RQ3: Does the conventionality of some linguistic metaphors affect the degree of similarity between two terms aligned in linguistic metaphors vs. visual metaphors?
- RQ4: Does the inner complexity of visual metaphors affect the degree of similarity between two terms aligned in linguistic metaphors vs. visual metaphors?
- RQ5: What type of semantic information is shared between metaphorical terms and captured by DSMs?
- RQ6: Is the metaphor formulation still valid when substituting terms with a neighbour in the semantic space? In that case, does the semantic similarity between the terms remain consistent with the trends registered with the original formulation?

2 Method

We formalized a sample of representative visual metaphors and a sample of representative linguistic metaphors into A-is-B correspondences by applying state-of-the-art procedures, as described in Section 2.2.1. Then we computed the distributional similarity between each AB pair by comparing the vectors of A and B, containing the linguistic contexts in which each A and B typically occurs, as they are retrieved by the adopted distributional method, Distributional Memory (Baroni and Lenci 2010), described in Section 2.1.

2.1 Distributional memory

For our analyses, we used Distributional Memory³ (DM; Baroni and Lenci 2010), a multi-purpose structured DSM where distributional information of English is

³ <http://clic.cimec.unitn.it/dm/>

stored in the form of a third-order tensor. Each entry in the model consists in fact of a weighted corpus-derived triple $t_v = \langle \langle w_1, l, w_2 \rangle, v_t \rangle$ where l is the lexico-syntactic dependency link that associates the words w_1 and w_2 , and v_t is the weight assigned to the triple via a measure of association: the Local Mutual Information (LMI).

Example of t_v triples in DM are:

$\langle \langle \text{teacher} - n, \text{use}, \text{book} - n \rangle, 53.6 \rangle$

$\langle \langle \text{food} - n, \text{sbj.intr}, \text{cost} - v \rangle, 79.8 \rangle$

The distributional information stored in DM was extracted from the concatenation of (1) the ukWac corpus⁴; (2) a mid-2009 dump of the English Wikipedia⁵; and (3) the British National Corpus.⁶ The resulting corpus (2.83 billion tokens) was tokenized, POS-tagged, lemmatized, and dependency-parsed. Among the different versions of DM available, we use the TypeDM version, where the set of links consists both of dependency paths and lexicalized links, and the scoring function, which assigns the weight to the triple, computes the number of distinct types displayed by a link when it co-occurs with the relevant words.

As suggested by its authors, DM then relies on fine-grained features when computing distributional similarity, and, by making use of a three-way structure, makes it possible to flexibly model various semantic phenomena, by transforming the tensor into different semantic spaces. In particular, we are interested in the transformation from the original tensor to the word-context matrix ($w_1 \times w_2$, word by link-word matrix), which is used as the co-occurrence matrix for our semantic similarity computations. This means that we make use of vectors labelled with words w_1 , whose dimensions are labelled with tuples of type $\langle l, w_2 \rangle$.

For example, in the case of the example triples shown above, we will consider *teacher*–*n* and *food*–*n* vectors labels, and $\langle \text{use}, \text{book} - n \rangle$ and $\langle \text{sbj.intr}, \text{cost} \rangle$ context features.

In our experiments, following this framework, we model the distributional similarity of the pairs of words involved in a metaphor by computing the cosine between their word vectors.

⁴ <http://wacky.sslmit.unibo.it/>

⁵ http://en.wikipedia.org/wiki/Wikipedia:Database_download

⁶ <http://www.natcorp.ox.ac.uk/>

2.2 Materials

The sets of visual and linguistic metaphors used for the distributional analyses were randomly extracted from the VU Amsterdam Metaphor Corpus⁷ of linguistic metaphors and the VisMet Corpus⁸ of visual metaphors. These corpora are balanced and are representative of the two modalities, and therefore they have modality-specific inherent variability.

A sample of 50 visual metaphors and a sample of 50 linguistic metaphors were randomly selected. However, as often happens when dealing with real-world data (see discussion in Goodall et al. 2013), in order to be suitable for the present investigation, the metaphors had to meet a number of criteria⁹: they had to cover the different genres in which visual and linguistic metaphors are typically expressed,¹⁰ and different types of realization,¹¹ and only visual metaphors without meaningful linguistic anchors¹² were selected, such as the example in Figure 1.

2.2.1 Metaphor identification procedure

In the field of metaphor studies, fundamental criticisms have been raised against the arbitrary approach employed by metaphor experts to identify and formalize conceptual metaphors within the CMT framework. Such formulations are, in fact, based on the verbalizations of experts who look at sets of metaphoric expressions in language use and discuss the best wordings to express a conceptual metaphor that would encompass and explain such linguistic metaphoric expressions (see Steen et al. 2010 for a review). To tackle this problem, a reliable and objective procedure has been developed to identify metaphors starting from language use (MIPVU, Steen et al. 2010).

⁷ <http://www.vismet.org/metcor/search/showPage.php?page=start>

⁸ <http://www.vismet.org/VisMet/>

⁹ A more detailed description of how the materials is reported in Bolognesi (2016).

¹⁰ For linguistic metaphors: academic texts, news, fiction and conversations (see Metaphor Corpus documentation); for visual metaphors: advertising, political cartoons, artworks (see Vismet Corpus documentation).

¹¹ For linguistic metaphors indirect and direct realizations; for visual metaphors juxtapositions, fusions and replacements (as described in the respective corpora documentation).

¹² As described in the article mentioned above, the images were manipulated so that the linguistic anchors (slogans and other words presented in the picture) were blurred. Then the images were shown to the participants and only those that were correctly interpreted without relying on the information encoded in the linguistic anchors were kept as instances of representative (monomodal) visual metaphors.

This procedure relies on the idea that the majority of metaphors found in language are not direct comparisons expressed in a direct way through words (such as, for example, “my lawyer is a shark”), but are instead words used in a metaphorical way in a given context. In this sense, the majority of linguistic metaphors are expressed *indirectly*, and they are based on a contrast between the *contextual* meaning of the word (which is metaphorical) and its *basic* meaning (which is literal). According to this procedure, given a text with a potentially metaphorical word, the contextual meaning and the basic meaning of that word are very different, and the contrast between the two is the basis on which the metaphor is created. For example, in the (highly conventionalized and frequently used) sentence “I *see* what you mean”, the contextual meaning of *see* is *understand*, while the basic meaning refers to the physical ability of sight. The two meanings (understanding and physically seeing) belong to different conceptual domains and it follows that the word *see* is to be considered metaphorical in the linguistic context mentioned above. In this sense, the linguistic metaphors identified through MIPVU are expressed at a different level compared to the conceptual metaphors formulated in accordance to CMT (Lakoff and Johnson 1980). The latter ones are formulated on the basis of the analysis of *sets* of metaphoric linguistic expressions, rather than on one specific instance of text.¹³ It follows that the metaphors formulated in accordance with MIPVU do not (necessarily) overlap with the formulations of the conceptual metaphors suggested within the CMT framework, because the two dimensions of meaning that they tackle are different. While MIPVU tackles the linguistic dimension of metaphor, and is based on context-specific meaning, CMT tackles a higher (and less language-specific) conceptual dimension. In the case of *understand-is-see*, the linguistic metaphor identified through MIPVU overlaps with its conceptual formulation UNDERSTAND-IS-SEE, but this is not necessarily always the case. For example, consider the following statement, extracted from the Metaphor corpus:

- (1) *The other side of press management depending on informal contacts is less easy to assess.*

In this example, the word *side* is marked for metaphoricity in the Metaphor Corpus. According to MIPVU it can be argued that the contextual meaning of

¹³ In addition, the choice of the exact wordings used to express a conceptual metaphor in CMT are quite arbitrary, as described above.

side is *aspect*,¹⁴ while its basic meaning is *surface*.¹⁵ This leads to the following formalization of the linguistic metaphor: *aspect-is-surface*. This kind of formalization (or similar ones based on synonyms) cannot be explicitly found in state-of-the-art repositories of conceptual metaphors, even though the metaphoric expression is quite conventional. However, one could argue that this linguistic metaphor relates to and provides an extension to the underlying conceptual metaphor IDEAS-ARE-OBJECTS (and therefore ideas have -metaphorically-tangible surfaces, as objects do). As the reader can see, in this case the linguistic and the conceptual formalizations, obtained through two different methods (MIPVU and CMT), do not fully overlap, although they relate to the same phenomenon.

For the identification of the metaphor terms involved in the sample of visual metaphors, the VISMIP procedure was applied (Šorm and Steen under review). This procedure relies on the idea that visual metaphors typically present (different types of) perceptually incongruous elements that violate the expected scenario and need to be mentally replaced with other elements, whose function it is to restore the visual feasibility (i. e. perceptual congruency) of the scenario. Detecting such elements (step 3 of the VisMip procedure) and replacing them with elements that would help restore the expected scenario (step 4) is one of the main cognitive operations that needs to be performed to unravel the metaphor. In this sense, the perceptual incongruities and their replacements constitute the metaphor terms (or part of them), or they cue to the abstract concepts that constitute the actual conceptual domains of the metaphor, by means of metonymies. For example, if an anti-tobacco social campaign shows an open package of cigarettes with bullets inside, instead of the (expected) cigarettes, then the bullets need to be mentally replaced by the viewer with cigarettes in order to restore the expected scenario (i. e. a package of cigarettes with cigarettes inside). In this metaphor, the cigarettes are therefore compared to bullets, and this comparison triggers features of the base term *bullet*, such as *dangerous*, *pain-producing*, *bad*, *deadly*, etc., which are mapped onto the target term *cigarette*.

Figure 2 shows another instance of a visual metaphor used for analysis belonging to the genre of political cartoons. Here, the cartoonist represents a bar-code as if it were a jail, exploiting the visual resemblance between the vertical lines of the barcode and the bars of a jail cell. The representational meaning of this metaphorical image can be formalized as *barcode-is-jail*.

14 Third listed sense of the entry “side” in the Macmillan online dictionary.

15 Second listed sense of the entry “side” in the Macmillan online dictionary.




Figure 2: A political cartoon used for analysis, retrieved from the VisMet corpus of visual metaphors. Credits and copyright owner: Vangelis Pavlidis.

Finally, Figure 3 shows an instance of a visual metaphor belonging to the genre of advertising, in which a (Listerine) mouthwash bottle is ‘enriched’ with the handle of a grenade attached to the bottle cap, to suggest the incredible power of this product in killing bacteria (comparable to the power of a grenade).



Figure 3: An advertisement used for analysis, retrieved from the VisMet corpus of visual metaphors. Credits and copyright owner: Ferdi Rizkiyanto.

As a result of the application of the MIPVU and VISMIP procedures, we obtained a list of 50 representative linguistic metaphors and 50 representative visual

metaphors, formalized in A-is-B expressions, where A and B are two words representing the two compared entities.

We acknowledge the fact that the distributional analyses reported in Sections 3.1–3.5 are based on the exact A and B terms that we obtained and agreed upon, from the application of MIPVU and VISMIP. Arguably, slightly different verbalizations of the same concepts (e. g. “prison” instead of “jail” in Figure 3) might result in slightly different distributional similarities between AB pairs. In reply to this possible objection, we bring two arguments. First, DM has been evaluated by the authors in a synonym detection task (Baroni and Lenci 2010: 693–694), which shows that this method can reliably detect two synonyms as opposed to two non-synonyms, by retrieving a significantly higher cosine similarity between synonyms, compared to non-synonyms.¹⁶ It follows that synonyms are distributionally similar in DM, i. e. they tend to appear in the same linguistic patterns. Thus, it can be argued that the distributional similarity between A and B (for example *barcode* and *jail*), and the distributional similarity between A and B1 (for example *barcode* and *prison*), where B and B1 are synonyms, should not be substantially different within DM. Our second argument is the actual analysis reported in Section 3.6, where for each AB pair retrieved through MIPVU and VISMIP we compared the cosine similarities between A and B to the cosine similarities between the 10 closest terms related to A and the 10 closest terms related to B. By doing to, we covered different possible formalizations of the same metaphors.

2.3 Procedure

As a first step (in Section 3.1) we computed the average semantic similarity between all the AB metaphor pairs in our list of metaphors ($N=100$). We then interpreted this result by comparing it with two samples of *non-metaphors* within DM. More specifically, we compared the metaphor similarity with:

- The average similarity between randomly paired words, extracted from a sample of 50K randomized pairs of words belonging to the same part of speech (*Random pairs*), in order to understand whether the similarity that is registered for metaphors is substantially above pure chance.

¹⁶ However, antonyms (e. g. *dead – alive*) are often as close as synonyms in the semantic space. Such counter-intuitive phenomenon is well-known in the DS community, and is ascribed to the fact that indeed antonyms tend to have common semantic properties, though their opposite polarity, as they belong to the same semantic domain. As a consequence, they tend to appear in very similar linguistic contexts.

- The average similarity between each word in DM and its 10 closest words (*Word-neighbour pairs*), in order to understand how metaphor similarity compares to the similarity of highly semantically related words.

Then, we took into account different variables that could predict different similarity values, in order to tackle the other five research questions that we initially formulated. In Section 3.2 we compare how metaphor similarity changes when the modality of expression (visual or linguistic) comes into play. In Section 3.3 we take into account the Conventionality variable, and analyse how this variable influence metaphor similarity. In Section 3.4 we reformulate the visual metaphors at a higher (more abstract) level to investigate whether the inner complexity of visual metaphors (i. e. the fact that they are often combined with metonymies) affects the degree of metaphor similarity captured by DM. In Section 3.5 we provide a qualitative analysis of the type of semantic information that seems to be typically shared between metaphor terms and captured by DM. Finally, in Section 3.6 we analyse whether a new formulation of the metaphor, derived by using neighbours in the semantic space, can preserve the similarity and the semantic features that are relevant for the metaphor in the original formulation.

The overall organization of the analyses is summarized in Figure 4, where the 10 different samples of data are displayed. The two samples of non-

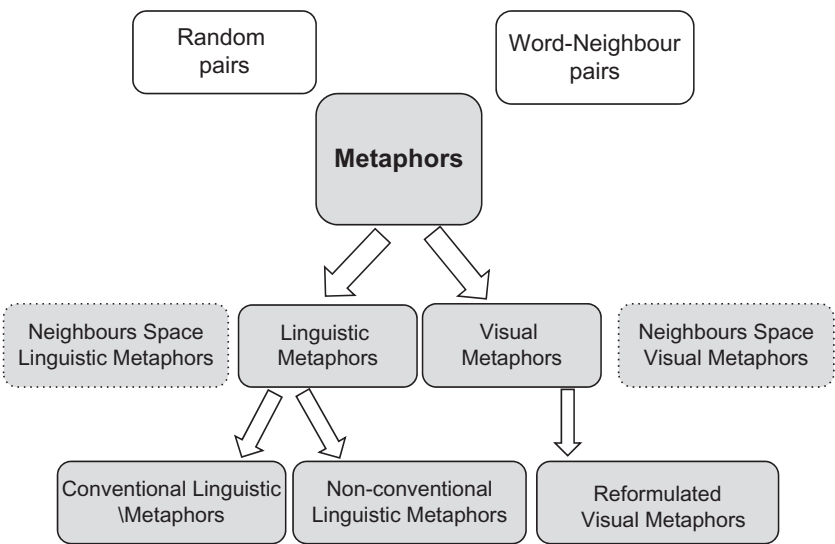


Figure 4: A visual representations of the categories of word pairs analysed in the experiments.

metaphors on the top of the image (Random pairs and Word-neighbour pairs) were only used initially in order to establish the two baselines of semantic similarity within DM, and the general feasibility of the subsequent comparisons. That is, they were only used in the first analysis to establish whether DM is capable of capturing what we call *metaphor similarity*, against similarity by chance and similarity between highly associated words.

While the Random pairs set consists of 50K random pairs of words from the vocabulary of DM, the Word-neighbour pairs set is given by pairing each word in the vocabulary with each of its semantically closest words. For this aim, we used as resource the list of the top 10 nearest neighbours of each word in DM and the relative similarity computations, based on a reduced version of the $W \times LW$ matrix (5K dimensions) derived from DM¹⁷

3 Analysis

3.1 Metaphor similarity in DM

Our first research question was formulated as follows: How does distributional similarity modelled by DM capture the similarity between two terms compared in a metaphor?¹⁸

Table 1 shows that the mean of the similarity values for AB pairs aligned in metaphors is 0.16, a value that is higher than the random pairs average

Table 1: Metaphor similarity in DM, compared to similarity by chance and similarity between highly associated words (μ : arithmetic mean of the sample; σ : standard deviation of the sample).

	<i>N</i>	Semantic similarity	
		μ	σ
Metaphors	100	0.16	0.09
Word-neighbour pairs	306,860	0.36	0.13
Random pairs	50,000	0.05	0.05

¹⁷ Available with DM at <http://clic.cimec.unitn.it/dm/>

¹⁸ For all the data, see the supplementary materials in the following repository: <https://goo.gl/icZC9G>

(CosSim = 0.05), and lower than the top 10 neighbours average (CosSim = 0.36). This result seems to align with the fact that metaphorical mappings rely on a certain degree of common semantic features, though generally not as high as the common semantic features shared by two semantic neighbours, given that the words aligned in a metaphor belong to different semantic domains. In fact, the relation which holds between the source and the target words seems to rely on a level of semantic similarity which is higher than the one given by a random association, but lower than the one given by words which are predicted by the model to share many semantic attributes.

We ran an ANOVA on the three samples (*Metaphors*, *Random Pairs*, and *Word-Neighbour Pairs*) to confirm that the average CosSims of the three samples were significantly different ($F = 29,464.65$, $df = 2$, $p < 0.005$). Post hoc comparisons using the Tukey HSD test¹⁹ indicated that each of the three means is significantly different from the other two ($p < 0.005$).

Once it was established that metaphor similarity is captured by DM, we proceeded with specific analyses aimed at investigating the role of different variables within the overall sample of 100 metaphors. The following sections report the descriptive and the inferential statistics of the analyses that we conducted (summarized respectively in Tables 2 and 3).

Table 2: Descriptive statistics: average cosines (i. e. average semantic similarity) for different categories (μ : arithmetic mean of the sample; σ : standard deviation of the sample).

	<i>N</i>	Semantic similarity	
		μ	σ
Metaphors	100	0.16	0.09
Linguistic metaphors	50	0.19	0.09
Conventional linguistic metaphors	16	0.22	0.10
Non-conventional linguistic metaphors	34	0.18	0.07
Neighbours space linguistic metaphors	6,050	0.14	0.07
Visual metaphors	50	0.12	0.08
Reformulated visual metaphors	50	0.15	0.09
Neighbours space visual metaphors	6,050	0.11	0.07

¹⁹ The Tukey Honestly Significant Difference test is based on a studentized range distribution q , and it is commonly used in conjunction with ANOVA to evaluate the pairwise significance of the means analysed with ANOVA (Tukey 1949).

Table 3: Inferential statistics. Results of the Tukey test (post-hoc), which was run in conjunction with the one-way ANOVA ($F=38.97$, $df=12,399$, $p<0.001$). In this table we report only the comparisons between samples that we consider relevant for our argumentation (the Tukey test provides a comparison between all the possible pairings of samples).

Sample pair	Tukey HSD Q statistics	Tukey HSD p-value	Tukey HSD inference ($p<0.05$)
Metaphors vs. Linguistic metaphors	2.47	0.46	insignificant
Metaphors vs. Visual metaphors	2.47	0.46	insignificant
Visual metaphors vs. Linguistic metaphors	4.97	0.01	<u>significant</u>
Conventional linguistic metaphors vs. Non-conventional linguistic metaphors	2.12	0.78	insignificant
Visual metaphors vs. Conventional linguistic metaphors	4.99	0.01	<u>significant</u>
Visual metaphors vs. Non-conventional linguistic metaphors	3.55	0.19	insignificant
Reformulated visual metaphors vs. Visual metaphors	2.43	0.65	insignificant
Reformulated visual metaphors vs. Linguistic metaphors	2.54	0.61	insignificant
Reformulated visual metaphors vs. Conventional linguistic metaphors	3.29	0.28	insignificant
Reformulated visual metaphors vs. Non-conventional linguistic metaphors	1.36	0.90	insignificant
Linguistic metaphors vs. Neighbour space linguistic metaphors	4.93	0.01	<u>significant</u>
Visual metaphors vs. Neighbour space visual metaphors	0.61	0.90	insignificant

3.2 Metaphor similarity and modality of expression

We took into account the modality in which the metaphors are expressed, and compared the average semantic similarity computed separately for linguistic and visual metaphors.²⁰ As reported in Table 2, we found that the former category has a mean of 0.19, and the latter of 0.12. As Table 3 shows, this difference between the two modalities is statistically significant. Moreover, Table 3 shows that the difference between the average similarity of a metaphor (in the *Metaphors* sample) is not significantly different from the average similarity of the *Visual metaphors* sample, nor from the average similarity of the *Linguistic metaphors* sample.

²⁰ For the complete lists of metaphors and related similarity values, see Appendix A.

3.3 Conventionality (of linguistic metaphors)

Our third research question was formulated as follows: Does the conventionality of some linguistic metaphors affect the degree of similarity between two terms aligned in linguistic metaphors vs. in visual metaphors?

We hereby define conventional linguistic metaphors as AB pairs of words compared in a metaphor, and identified through the MIPVU procedure, where A, or contextual meaning, is found in dictionaries as one of the possible senses of B (or basic meaning). Consider, for example, the linguistic metaphor *constraint-is-obstacle*, identified through MIPVU on the basis of the metaphorical word *barrier*, whose contextual meaning in the analysed text is *constraint* (target), and basic meaning is (physical) *obstacle* (source). The dictionary suggests that *obstacle* can carry the meaning of *constraint*. We therefore consider this metaphor to be linguistically conventional.

It has been argued that conventional linguistic metaphors such as those used in idiomatic expressions do not necessarily find a corresponding underlying conventional conceptual structure, but can be ascribed to lexical semantics that can be historically explained (Murphy 1996; Glucksberg 2001; Jackendoff 2002; McGlone 2007). Therefore, there is not necessarily a one-to-one correspondence between conventional linguistic metaphors and conventional conceptual metaphors. In addition, it has been argued that conventional linguistic metaphors are processed by means of lexical disambiguation: for the polysemous word *obstacle*, the meaning *constraint* is selected, given a specific context (Bowdle and Gentner 2005).

The tendency of conventional linguistic metaphors to generate polysemy is expected to affect the distributional similarity of the two terms. More specifically, given any metaphor A-is-B, if B acquires some senses of A, and the attributes that typically belong to A can therefore be used also for B, then the distribution over contexts between those two words will also tend to be more similar. In fact, in this case, the two words could pass, at least in some contexts, a substitution test, which is to say that substituting one with the other in a larger linguistic expression does not affect the original meaning of the expression. Therefore, a model which computes semantic similarity as a function of the linguistic contexts of occurrence which are shared by the two words should indeed account for the conventionalization of metaphors as attested in corpora.

We are interested in observing whether the conventionality of an A-is-B metaphor affects the distributional similarity value between A and B. We proceeded to divide linguistic metaphors into conventional and non-conventional ones, by taking the attested existence of a sense of B as one of the senses of A as

a sign of conventionality. To this end, we used two different linguistic resources: the Wordnet²¹ lexical database for English and the MacMillan online dictionary,²² which encompasses many metaphorical senses and is commonly used for identifying linguistic metaphors by scholars and researchers who apply the MIPVU procedure (see for example Steen et al. 2010).

Examples of linguistic metaphors that we classified as conventional²³ are:

- *Aspect-is-surface*
- *Harsh-is-hard*
- *Understand-is-see*

As Table 1 shows, we found that there is a minimal difference between the average similarity of the sample of conventional linguistic metaphors (CosSim = 0.22) vs. the average similarity of the sample of non-conventional linguistic metaphors (CosSim = 0.18). This difference is not statistically significant, and thus might be due to chance (Table 3). However, as Table 3 also shows, visual metaphors similarity is significantly lower than the similarity in conventional linguistic metaphors, but non-significantly lower than the similarity in non-conventional linguistic metaphors. This suggests that visual metaphors are comparable to non-conventional linguistic metaphors, but not to conventional linguistic metaphors.

A note of caution is in order: the sample of conventional linguistic metaphors contains only 16 items. In experimental research it has been suggested that samples need to have a size of at least 20 items ($N = 20$), in order to have sufficient power to draw valid conclusions when running ANOVA (cf. Simmons et al. 2011 for the argumentation). In this respect, this specific analysis has been reported with the sole aim of inspiring future research, but we are aware of the fact that the samples taken into account might be too small.

As for the visual metaphors, we did not carry out a formal categorization into conventional visual metaphors and non-conventional visual metaphors, because no clear method has yet been provided to define visual metaphor conventionality, nor is there a clear definition of visual conventionality *per se*. We do believe that this is a very interesting line of research and that it will arguably be explored in the coming years. As a matter of fact, some recent pioneering studies have started to show in a qualitative fashion how some of the conventional metaphors extracted from language use can be expressed in visuals as well (see for example Forceville 2011; Hidalgo and Kraljevic 2011),

²¹ <https://wordnet.princeton.edu/>

²² <http://www.macmillandictionary.com/>

²³ For the list of metaphors recognized as conventional, see Appendix B.

as well as how primary metaphors can be realized visually (Ortiz 2011) and even image schemas (Perez Hernandez 2014).

Although we could not categorize visual metaphors into *visually* conventional and *visually* novel, for lack of theoretical definitions as well as lack of methodological tools (such as visual dictionaries), we attempted an informal categorization of the visual metaphors (expressed into A-is-B pairs) into conventional vs. non-conventional ones, by consulting standard *linguistic* dictionaries,²⁴ as we did for the linguistic metaphors. As expected, however, we found that for *none* of the visual metaphors included in our sample was the meaning of the term B included among the senses of the term A. It follows that, considering the semantic information contained in the (linguistic) dictionary, none of the visual metaphors are conventionalized in language use.

3.4 Reformulating visual metaphors

Different communication modalities have different ways of expressing messages, and have inner modality-specific peculiarities and limitations. This phenomenon clearly also applies to messages that involve metaphors. Our third research question was formulated as follows: Does the inner complexity of visual metaphors affect the degree of similarity between two terms aligned in linguistic metaphors vs. visual metaphors?

Because we use a model (Distributional Memory) based on linguistic occurrences, and we refer to the *concepts* compared in metaphors using the *words* that express them, this distinction between modalities needs to be made clear, as the metaphors we analysed come from both language and images. Indeed, visual metaphors might be very different from linguistic metaphors: although conceptual metaphors are assumed to belong to the realm of thought, the production of a message must always respect the needs of the modality in which it is expressed (in this case language or images). Besides, the two modalities, language and images, seem to categorize perceptual experiences in different ways. Thus, the modality in which a metaphor is expressed might possibly influence the choice of the concepts aligned in the metaphor.

Looking at the images from which the sample metaphors are extracted, it is possible to observe the following phenomena, also acknowledged in the field of visual communication:

²⁴ In particular, we used the Macmillan online dictionary, extensively used by metaphor scholars (see for example Steen et al. 2010).

- In visual metaphors, the message is often dense, since it has to be compressed in a single shot (e. g. Mitchell 1994; Forceville 1996). For example, a political cartoon showing an upside down shopping cart on a dark and shadowy background, with people trapped inside the shopping cart, asking for help while trying to break the metal bars to escape, triggers a lot of emotional information and addresses the critical public with a sensitive message about the implications of consumerism for our society. This information can be easily packed in one image, but not so easily and effectively explained through words.
- Visual metaphors are often quite creative and original as a feature of visual communication genres such as satire, advertisement, and art (see for example Forceville and Urios-Aparisi 2009). For this reason, metaphorical mappings are often unexpected, context-dependent features. For example, an advertisement for mouthwash showing the product graphically combined with a bomb can be interpreted negatively by the naive viewer who does not rely on the information conveyed by the specific genre of advertising, seeing as advertisements typically try to convey the *positive* features of their products. With this contextual knowledge in mind, the viewer needs to map the positive effects of the explosive power of the bomb to the mouthwash (i. e. its impact on cleanliness).
- Abstract concepts cannot be represented in images. Because of this limitation, other communication devices are typically used in visual metaphors, such as symbolism (e. g. *dove* for *peace*), or, in the case of generic terms referring to concrete objects, hyponymy (e. g. *bulldozer* for *machine*, cf. Bolognesi et al., under review). Because of this limitation of the visual modality, the viewer needs to disentangle multiple layers of meaning when interpreting the image (e. g. Šorm and Steen under review).
- Similarity achieved through graphic means (e. g. colour, shape) between two objects is often exploited to cue a conceptual similarity between the two (e. g. van Weelden et al. 2012). Thus, the similarity between two terms aligned in a visual metaphor is often played on through shared perceptual features such as colour, shape and perspective, while the conceptual features originally belonging to the source need to be mentally transferred to the target, which does not necessarily possess them in the first place. For example, a commercial campaign for the airline KLM can show a flying swan that clearly resembles an airplane (in colour, shape, and location -in the sky-). Yet the features that need to be mentally transferred from source to target are not merely physical (having wings, being white and able to fly are not the advertised features). Instead, features such as elegance, beauty,

and majesty have to be mapped from the swan to the KLM airline, via the visual resemblance between the bird and an aircraft. As this example shows, visual metaphors often exploit metonymical relations (e. g. aircraft for airline), to cue abstract concepts.

This general complexity of visual metaphor structure can be seen as a potential motivation for the lower than average semantic similarity obtained in our analyses.

We then addressed the modality-specific limitations of visual metaphors described above by reformulating visual metaphors in such a way that their intrinsic complexities were partially simplified (for example, we reformulated *airplane-is-swan* into *airline-is-swan*, given that the image aims at Advertising positive properties of the airline company) With the A-is-B reformulated visual metaphors, we expect to observe an increase in the average semantic similarity score.

We reformulated the list of visual metaphors (which can be found in Appendix C) by applying the following criteria:

1. If the source or the target of a visual metaphor was a conventional symbol, we substituted the symbol with a word that refers to the underlying (abstract) concept.

E. g. In a street-art work showing a white dove carrying an olive branch in its mouth, with a shooting target drawn on its chest, the original metaphor *dove-is-target* was reformulated as *peace-is-target*.

2. If the semantic features relevant for a visual metaphor were shared by the source and the target in virtue of intra- or extra-image contextual information, we reformulated the metaphor in such a way that it is less dependent on the context.

E. g. In an Amnesty International campaign in favour of signing petitions for social aims, a pen is depicted in the shape of a bullet. The original formulation of the visual metaphor was *pen-is-bullet*, while the reformulated version became *signature-is-weapon*.

3. If A or B in a visual metaphor had a more abstract hypernym that still encompassed all the semantic features relevant for the metaphor, we substituted the original word with one referring to the more general concept.

E. g. An advertisement for a painkiller cream showing a man with his hands in the shape of a bulldozer was reformulated from *hand-is-bulldozer* into *hand-is-machine*.

4. If the source or the target of a visual metaphor metonymically cued to an abstract concept, we substituted the original term with abstract one.

E. g. In an advertisement for a coffee brand, a white, empty coffee cup is photographed from above, so that it resembles an electric switch, with the cup handle resembling the switch handle. The original metaphor *cup-is-switch* was reformulated into *coffee-is-switch* (acknowledging the ‘container for content’ metonymy, exploited to suggest the ability of the coffee to provide energy and “turn on” the consumer).

5. If a visual metaphor played on a conventional *linguistic* metaphoric expression, we included this knowledge in the reformulation of the visual metaphor.

E. g. A WWF campaign to protect wildlife shows an open human hand coloured in black and white stripes, as if it was a zebra. The visual metaphor plays on the linguistic expression “give a *hand*” to WWF, which means “provide *help*” to WWF. The original formulation *hand-is-zebra* was thus reformulated into *help-is-hand*.

If necessary, we applied more than one of the above-mentioned modifications. However, this procedure was managed in a such a way that as few modifications as possible are applied to the original pairs of words, stopping as soon as the mapping becomes more abstract and less image-dependent. Moreover, in some cases, no modifications were applied, as it was not possible to reformulate the metaphor without losing the original message.²⁵

Relying on these guidelines, we first reformulated the visual metaphors acting as two independent annotators, and then compared our reformulations and calculated the interrater agreement, as in standard content analyses (Cohen’s $k = 0.74^{26}$). The disagreements on the reformulations were mediated in a discussion until the final reformulations were agreed upon by the two annotators.

We would like to stress that the purpose of our study was not to develop a reliable procedure to reformulate visual metaphors at a more abstract level of expression. Instead, we sought to remove some of the modality-specific phenomena typical of visual communication, in order to investigate possible changes in the similarity scores between metaphor terms when different formulations based on the same image are provided. The analysis reported in Section 3.6 tackles this issue from a different perspective: being aware of the fact that the words chosen to formalize an A-is-B metaphor are crucial for determining high or low similarity values, we compared the similarity values for each AB pair to those obtained with the top 10 neighbours for A and with the top 10

²⁵ For the list of reformulated visual metaphors and related similarity values, see Appendix C.

²⁶ According to the literature, this value has to be interpreted in terms of substantial reliability of the reformulations (cf. Artstein and Poesio 2008).

neighbours for B, retrieved by DM. In this way we aimed to cover different possible verbalizations of the same metaphors (exploiting synonyms, hypernyms, etc.), and could observe whether the similarity values were preserved.

As can be seen in Table 2, the average similarity between metaphor terms in reformulated visual metaphors has increased slightly (CosSim = 0.15, cf. original formulations: CosSim = 0.12). However, as Table 3 shows, the difference with the original formulations is *not* statistically significant. This suggests that the similarity between metaphor terms in visual metaphors is preserved, when comparing the original formulations and the new formulations. In other words, our reformulations, which aimed at simplifying the graphic complex structures of some of the visual metaphors, did not substantially change the nature of the visual metaphors.

Table 3 also shows, however, that the difference between reformulated visual metaphors and linguistic metaphors became insignificant (while the difference between the original formulations of the visual metaphors and linguistic metaphors was significant). This suggests that the reformulations that we performed on the visual metaphors somehow made the visual metaphors more similar to the linguistic metaphors.

3.5 Analysis of shared linguistic contexts

Our fifth research question was formulated as follows: What type of semantic information is shared between metaphorical terms, and captured by DM?

In order to see if and how the relevant common semantic traits of the source and the target of a metaphor arise from the contextual information, we looked into the shared linguistic contexts retrieved by DM in a qualitative fashion. Following the intuition of Miller and Charles (1991), we looked at the semantic constraints of the contexts of use in order to detect, when possible, those semantic features which justify the occurrence of a term in a specific linguistic collocation.

We carried out this type of analysis on a sample of metaphors, looking only at those contexts that we established as more salient for the similarity measurements. Among all the overlapping contexts, we extracted the 100 with highest association measures (LMI) for both A and B.

We observed the shared contexts for different types of metaphors, making our observations on the basis of their modality, conventionality and where their semantic similarity value was with respect to the mean of the category they belong to (i.e. above or below the average similarity value of that category). Here is a summary of our qualitative findings, with examples.

- *Linguistic metaphor with average similarity: emotion-n force-n* (CosSim = 0.18, cf. average CosSim of linguistic metaphors = 0.19).

The conceptual similarity between these two terms is plausibly based on the intensity, power, and influence that both entities have on human experience. Shared linguistic contexts relating to these shared features, retrieved by DM, are: *urgency-n*, *impress-v*, *tremendous-j* *amount-n*, *influence-v*. In addition, semantic collocates that originally seem to belong to the domain of the metaphor target (here the semantic domain of *emotion*), such as *passion-n*, appear as shared contexts, which means that they are acquired as linguistic contexts of the source, *force*.

- *Visual* metaphor with *average* similarity: *hand-n* *bulldozer-n* (CosSim = 0.12, cf. average CosSim of visual metaphors = 0.12).

The conceptual similarity between these two terms is plausibly based on the domain of *working* (where hands are used as tools): some contexts do reflect this aspect, such as *use-v*, *tool-n*, *collect-v* and *digging-n*.

- *Conventional* linguistic metaphor with relatively *high* similarity: *understand-v* *see-v* (CosSim = 0.45 cf. average CosSim of Conventional metaphors = 0.22).

The shared context with highest measure of associativity was *by reasoning-n*. This is quite clearly a linguistic context that originally belongs to the semantic domain of the target, *understand*. However, it appears to be acquired by the source, *see*, which shows quite a high measure of associativity with both source and target terms. Similarly, semantic collocates such as *explanation-n* or *fact-n* are acquired by the source from their original attribution to the target.

- *Conventional* linguistic metaphor with relatively *low* similarity: *opportunity-n* *door-n* (CosSim = 0.09, cf. average CosSim of Conventional linguistic metaphors = 0.22).

In this case, the similarity is quite low. The overlapping contexts for this metaphor seem to be mainly related to the general semantic domain of the source, *door*, rather than to the target, *opportunity* (e. g. *access-n*, *widen-v*, *knock-v*, *entrance-n* or *opening-n*). For this reason, we suggest that the conventional use of this metaphor is not reflected much in the similarity value, because the metaphor is more often realized linguistically, using *opportunity* in linguistic collocates, which are related to the semantic domain of *door*, rather than vice versa.

- *Non-conventional* linguistic metaphor with relatively *high* similarity: *plant-is-person*²⁷ (CosSim = 0.29 cf. average CosSim of non-conventional linguistic metaphors = 0.18).

²⁷ Interestingly, the conceptual metaphor PLANT-IS-PERSON is considered to be conventional. However, according to our definition of linguistic conventionality, the linguistic metaphor *plant-is-person* is not conventional (cf. Section 3.3).

Some of the most salient collocates refer directly to the common conceptual structure of *plant* and *person* (both organisms), thus referring to the semantic area of life, nature, and biological systems: *natural-j*, *live-j*, *molecular-j*, *nature-n*, *genome-n*.

- *Non-conventional* linguistic metaphor with relatively low similarity: *end-is cloth-n* (CosSim = 0.05, cf. average CosSim of non-conventional linguistic metaphors = 0.17).

In this case, overlapping contexts are quite generic, and no pattern seems to emerge that suggests a specific common semantic feature between the two concepts.

- *Visual* metaphor with relatively high similarity: *pen-n thermometer-n* (CosSim = 0.34, cf. average CosSim of visual metaphors = 0.12).

The two objects have important semantic features in common: they are both devices, they have similar shapes and they contain liquids. Many of the shared collocates refer to these common semantic features: *instrument-n*, *professional-j*, *use-n*, *speed-n*, *shape-n*, *type-n*, *technology-n*, *durable-j*, *novel-j*, *low-cost-j*, *high-quality-j*, *lightweight-j*. Other contexts seem instead to be due to specific similarities in the way both of the devices work, for example those which contain the expressions *have button-n*, *capillary-j*, *pressure-n*.

- *Visual* metaphor with relatively low similarity: *globe-n ice-cream-n* (CosSim = 0.03, cf. average CosSim of visual metaphors = 0.12).

The aspect that is relevant for the metaphor in the image where it was extracted, is the fact that both the elements are melting (at least part of the globe, namely the glaciers). Except for one context (*liquid-j*), none of them seem to capture this feature. The lower value of similarity also suggests that very few semantic features are shared by the two concepts.

In conclusion, the qualitative analysis of shared contexts between metaphor terms suggests that two factors, among others, have a strong effect:

1. The shared linguistic patterns seem to directly reflect semantic features that are shared by the concepts to which the metaphor terms refer.
2. Within an AB metaphor pair where B is frequently used in place of A, and the A-is-B metaphor therefore is linguistically conventional, then the linguistic contexts shared by A and B are typically attested in a DSM.

Such observation aligns, on one side, with the assumption that distributional similarity measures attributional similarity (Turney and Pantel 2010), and, on the other, with our finding that metaphor conventionality can be a factor that increases distributional similarity.

3.6 Metaphors and neighbour space

Our sixth and final research question was formulated as follows: Is the metaphor formulation still valid when substituting terms with a neighbour in the semantic space? In that case, does the semantic similarity between the terms remain consistent with the trends registered with the original formulation?

Distributional semantics allows us to compute the semantic similarity between *words*, looking at their *linguistic* contexts of use. For this reason, the choice of words used to formulate the metaphor has a crucial role.

In our analyses, we rely on the metaphor formulations derived by the application of MIPVU and VISMIP procedures. Then, we tested the consistency of the similarity values obtained by substituting A and B with semantically closer words retrieved by DM. To this end, we used the list of the top 10 nearest neighbours of each word in DM mentioned above.

Neighbours include synonyms (*person* and *human*), meronyms (*car* and *wheel*), antonyms (*old* and *young*), co-hyponyms (*car* and *van*, which are both hyponyms of *vehicle*), and words that are functionally related or frequently associated (*pencil* and *paper*). Therefore, not all the neighbours can successfully maintain the original meaning of the metaphor.

This analysis was designed to determine whether a new formulation of the metaphor, derived by using neighbours in the semantic space, can preserve the semantic features which are relevant for the metaphor in the original formulation. In that case, we consider the new pair a potential target-source association and an alternative formulation of the original metaphor. This type of determination of *metaphoricity* is carried out on a qualitative basis, by looking at the semantic traits of the concepts that the relevant words refer to.

Moreover, we are interested in observing whether the trends of semantic similarity values that we registered in our previous experiments for visual and verbal metaphors are also consistent with the new formulations derived using the neighbour space.

Let us define the sets of 10 words that are closest in the semantic space to w_1 and w_2 as K_{w_1} and K_{w_2} respectively. For each metaphor m in the form $w_1 - w_2$, the list of alternative formulations is the set S_m defined as follows:

$$S_m := (\{w_1\} \cup K_{w_1}) \times (\{w_2\} \cup K_{w_2})$$

which is the Cartesian product of the two sets of neighbours extended with the words they are related to respectively.

For every metaphor m , the similarity value for each pair in S_m is then measured and compared with the original pair.²⁸

This means that we compute the cosine similarity for all the possible combinations of w_1 or one of its neighbours and w_2 or one of its neighbours.

We also computed the mean of the semantic similarity values obtained for all pairs obtained in this fashion for all visual and verbal metaphors respectively, and compared it with the one obtained with the original pair only.

Table 1 shows the average means (and SD) of the neighbour spaces computed for visual and for linguistic metaphors. Table 3 shows that while the average similarity between the original formulations of the visual metaphors (CosSim = 0.12) and the average similarity computed within their neighbour space (CosSim = 0.11) is not statistically significant, this is not the case for linguistic metaphors (CosSim = 0.19 vs. CosSim = 0.14). In other words, when we substitute a metaphor term of a visual metaphor with a neighbouring term, the metaphor similarity tends to be preserved, but when we do this operation with linguistic metaphors, the metaphor similarity tends to decrease substantially.

The fact that the similarity between A and B in visual metaphors is fairly well preserved within the neighbours' semantic space can be explained by the type of neighbours that are typically retrieved by DM for words appearing in visual metaphors vs. linguistic metaphors. Visual metaphors typically involve concrete objects, and looking at the list of neighbours, we observed that concrete objects (e. g. *pen*, *swan*, *horse*, *bulldozer*) mainly tend to trigger hypernyms and co-hyponyms as neighbours. We also observed, as described below, that *co-hyponyms and direct hypernyms often tend to preserve the metaphoricality of an AB pair*. This can explain why the similarity between metaphor terms *propagates* better among neighbours for visual metaphors than for linguistic metaphors.

Here we summarize some phenomena emerging from the qualitative analysis mentioned above:

- The neighbour k of word w is a *hyponym* and the semantic features that are the basis of the metaphorical potential of the original pair which includes w are also inherited by k in virtue of the taxonomic relation. The distributional similarity remains constant or increases.

In these cases, it is possible that, exploiting the substitution with neighbours, a conventional metaphor formulation is obtained.

²⁸ For some of the words in the original formulation that have different senses due to homonymy (e. g. *match*, *organ*, *tablet*), the model is not able to distinguish among the senses, and the words that are closest in the semantic space are related to one of the senses, which might not be the one used in the metaphor. In cases where lexical ambiguity clearly biases the result, we discard the pair for our analysis.

For example:

Linguistic metaphor: *consequence-is-food* (CosSim = 0.10) → *consequence-is-fruit* (CosSim = 0.10)

Fruit is a hyponym of *food*, and the metaphor potential is maintained. The resulting new linguistic metaphor (*consequence-is-fruit*) is interestingly also one of its conventional realizations.

- The neighbour *k* of word *w* is a *co-hyponym* and the semantic features that are the basis of the metaphorical potential of the original pair which includes *w* are also inherited by *k* in virtue of a common hyponym. The distributional similarity remains constant or increases.

For example:

seaweed-is-plastic (CosSim = 0.22) → *seaweed-is-metal* (CosSim = 0.22)

- The neighbour *k* of word *w* is a hypernym and the semantic features that are the basis of the metaphorical potential of the original pair which includes *w* are inherited by *w* from *k* in virtue of the taxonomic relation, hence *k* has them as well. The distributional similarity remains constant or increases.

This phenomenon is typical of visual metaphors. Hypernyms are often quite generic terms (e. g. *bird*), and thus cannot be represented in a visual metaphorical mapping without choosing a particular hyponym (e. g. *swan*), though the semantic features that the last one carries might not necessarily be relevant for the metaphor.

For example:

book-is-tree (CosSim = 0.14) → *book-is-plant* (CosSim = 0.19)

discussion-is-war (CosSim = 0.29) → *discussion-is-conflict* (CosSim = 0.34)

- The neighbour *k* of word *w* is a synonym or near-synonym and the semantic features that are the basis of the metaphorical potential of the original pair are, by the definition of synonyms, shared by *w* and *n*. The distributional similarity remains constant or increases.

For example:

emotion-is-force (CosSim = 0.18) → *feeling-is-force* (CosSim = 0.19)

possibility-is-space (CosSim = 0.27) → *option-is-space* (CosSim = 0.36)

country-is-bomb (CosSim = 0.05) → *nation-is-bomb* (CosSim = 0.10)

- The neighbour *k* of word *w* is related to word *w* by a part-whole relation (meronymy/holonymy) and the new formulation resolves the metonymy in the visual metaphor. The distributional similarity increases.

For example:

pen-is-bullet (CosSim = 0.21) → *pen-is-gun* (CosSim = 0.36)

This qualitative analysis suggests that exploring the neighbourhood of the terms of a metaphor introduces interesting material which could act as a source of potential metaphorical associations. In fact, as shown above, some of the new formulations obtained maintain similar properties to the original formulation, and these include not only alternative formulations of an original metaphor, but also conventional realizations of the former.

Moreover, we also observed a different phenomenon, not directly related to the research question: by substituting one -or both- terms with a neighbour in the semantic space, it is possible to obtain a new pair which might have a metaphorical potential that does not stem from the original metaphor (it is not based on the same semantic features), but a new, creative one, which can be meaningful, at least in certain contexts (e. g. from *emotion – tide*, *anxiety-storm* is derived; from *airplane-swan*, *kite-bird* is derived). This suggests that a system of metaphor generation which aims at producing target-source pairs and exploits both the various semantic relations captured by the neighbours and an existing metaphorical relation could indeed detect pairs with metaphorical potential. Such results point towards an interesting direction for future research in metaphor generation.

4 General discussion

The analyses reported here aim to model the similarity between terms aligned in metaphors by means of distributional semantics. Visual and linguistic modalities of metaphor expression were compared and contrasted and significant differences were found in the degree of similarity captured by Distributional Memory, the DSM applied here.

For our first research question we investigated whether DM can capture and model the similarity between metaphor terms, as opposed to the average similarity between two randomly paired words (belonging to the same part of speech) as well as to highly related words, within the DM model. The analysis reported in Section 3.1 shows that there is indeed a significant difference between the average (cosine) similarity between two metaphor terms, as opposed to two randomly paired words (which have very low similarity scores), as well as between two metaphor terms, as opposed to two semantic neighbours (which have very high similarity scores). Metaphor similarity is ranked somewhere in between these two extremes. Based on the results obtained, we argue that distributional similarity can successfully account for the latent semantic similarity between metaphor terms and is successfully captured by DM.

We then investigated whether the similarity between metaphor terms presents substantial differences in relation to the modality in which the metaphor is expressed (Section 3.2). Our analyses show that the average similarity between metaphor terms is significantly higher for *linguistic* metaphors than for *visual* metaphors. We therefore argued that different modalities of expression construct and represent metaphors in different ways.

We then conducted additional analyses to investigate the reasons for this difference between the two modalities. We looked at two possible variables: Conventionality (the fact that linguistic metaphors are often conventional, while visual are not), and Complexity (the fact that visual metaphors often have modality-dependent inner complexities, while linguistic ones do not).

Our third research question tackled the Conventionality variable: does the conventionality of some of the linguistic metaphors affect the (greater) similarity obtained for this modality, as opposed to the visual metaphors, which are arguably less conventionalized, and in a sense more creative? In Section 3.3 we reported the analysis of the cosine similarities, distinguishing conventional from non-conventional metaphors, where conventionality is determined by whether a given meaning appears in WordNet or in the MacMillan online dictionary (and therefore is fully lexicalized) or not. The results show that the difference between the average similarities between conventional vs. non-conventional linguistic metaphors is not statistically significant, with the non-conventional metaphors displaying a minimally lower degree of similarity between metaphor terms. This phenomenon brings non-conventional linguistic metaphors closer to visual metaphors, as shown by the fact that the difference between these two categories is not statistically significant, while the difference between visual metaphors and conventional linguistic metaphors is statistically significant. We then argue that the Conventionality variable does indeed affect the different results obtained for visual and for linguistic metaphors with our previous analyses. We also specified that, because this analysis was conducted on the original sample of linguistic metaphors, which were then divided into conventional and non-conventional ones, the statistical significance of these results has to be taken with caution, since the sample of conventional metaphors is quite small ($N=13$).

In our third research question we tackled another possible variable that could explain (in combination with the Conventionality variable) why visual metaphor similarity was significantly lower than linguistic metaphor similarity. This is the Complexity variable: visual metaphors often have quite complex structures, motivated by peculiarities of the visual modality. In Section 3.4 we reported a distributional analysis based on a reformulation of the visual metaphors, previously identified and formalized into A-is-B correspondences in

accordance with established procedures. The final reformulation of the visual metaphors, based on cognitive operations reported in Section 3.4, was obtained through a two-step process performed by two independent coders (the authors of this paper), who first reformulated the visual metaphors independently, and then compared the reformulations and mediated the disagreements in a discussion (Cohen's $k = 0.74$). The reformulations rendered the A-is-B visual correspondences more abstract and less image-dependent, and therefore more similar to the linguistic metaphors. The new distributional analyses show that the similarity values of the reformulated visual metaphors do not significantly differ from those computed for the linguistic metaphors (while the original formulations of the visual metaphors did significantly differ from the sample of linguistic metaphors). It follows that the inner complexity of visual metaphors (disentangled in this new analysis that led to their reformulation) affects the degree of similarity retrieved by DM. From these analyses it can be concluded that the complexity of visual metaphors and their less conventionalized nature, compared to linguistic metaphors, can explain the differences between the overall degrees of similarity captured by DM for the two modalities, as reported in Section 3.1.

We then looked into the type of semantic information carried by the linguistic contexts that are shared between the two metaphor terms in a qualitative fashion (Section 3.5). As expected, we observed that many collocates can be put in relation with a shared aspect of the concepts that the two words refer to. This shows that cosine similarity as computed on DM is indeed strongly dependent on the degree of correspondence of semantic properties, an observation which is aligned with and support the assumption that distributional similarity measures attributional similarity. Moreover, we observed that often the shared linguistic contexts, and in particular the semantic collocates retrieved by DM, originally pertain to the target domain of the metaphor, and are then used in relation to the source. In other words, since a metaphor is a correspondence between two terms, where the source acquires part of the meaning of the target (e.g. *see* acquires part of the meaning of *understand*), then it follows that linguistic contexts that are peculiar to the target can also be used with the source, thus generating distributional similarity between the two words.

Finally, we investigated whether the similarity between two metaphor terms is preserved within the semantic neighbourhood of the two terms in the analyses in Section 3.6. We conducted these analyses in order to remedy a methodological weakness of the state-of-the-art procedures that we adopted to formulate the metaphors into A-is-B comparisons. Even though the MIPVU and the VISMIP procedures have been shown to be reliable (i.e. different independent annotators can achieve fairly good interrater agreements), it is often hard to agree on one exact and unique formalization for each metaphor. We therefore looked into

the top 10 closest semantic neighbours of each metaphor term, and recomputed all the cosine similarities between A and B to see whether the similarity between two metaphor terms is preserved when the two terms are verbalized through slightly different (but semantically related) words.

In Section 3.6 we reported the analyses of the similarity between the original metaphor terms and all the combinations of As and Bs within the list of their top 10 neighbours. We observed that the list of neighbours for each metaphor term included synonyms, quasi-synonyms, but also hypernyms, hyponyms, related terms, and antonyms, and saw that not all of the neighbours were capable of preserving the original meaning of the metaphor.. However, though a qualitative analysis we show how this substitution procedure is indeed able to retrieve some alternative formulations and conventional linguistic realization of the original metaphors, as well as pairs with a new metaphorical potential. As for the latter group, we found that exploiting one metaphorical relation and neighbours semantic relations it is possible to induce new conceptual associations, which could potentially be used as new target-source pairs. We also observed that the similarity of the original metaphor is preserved better in the neighbours' semantic space for *visual* metaphors than in that of *linguistic* metaphors. This might be due to the fact that visual metaphors tend, on average, to encompass more concrete concepts, which seem to trigger more synonyms, hypernyms, and hyponyms (which, in turn, seem to preserve the similarity in metaphors), while linguistic metaphors tend to encompass averagely more abstract concepts, which trigger different (and more distant) sets of neighbours. Further investigations can shed more light on these differences.

5 Conclusions

This project has aimed to analyse metaphors in terms of cross-domain mappings between concepts, looking at the semantic similarity of their representative words, as computed using a DSM based on corpora of language.

Our analyses formed a bridge between metaphor studies and distributional semantics. Relying on detailed procedures for metaphor identification established in the former field, we observed how such formulations relate with a computational semantic model (DM), which takes linguistic contexts of occurrence to be indicative of word meaning. We then exploited the distributional information harvested in large-scale language evidence in order to derive a quantitative measurement of the similarity between words, and thus the concepts they refer to.

We observed quantitative (and qualitative) differences between the similarity scores that characterize visual and linguistic metaphors. We then predicted that such differences could be explained by two variables: the Conventionality variable (i. e. many linguistic metaphors are conventional, while visual ones are not), and the Complexity variable (i. e. visual metaphors have inner complex *structures* that are inherent of their modality, but when such modality-dependent complexities are disentangled, and the metaphors are simplified, they come closer to linguistic metaphors).

Our results suggest that metaphors expressed in different modalities (language and images) have different modality-specific peculiarities that affect the similarity retrieved by the adopted model. Our findings lead to new research questions for the contrastive study of these two modalities of (metaphor) expression.

Although we did underline phenomena of visual communication that decrease the average similarity of the related concepts computed with a language-based model, a clear and exhaustive explanation of the differences in linguistic and visual metaphors is still needed. Furthermore, while we were able to show some effects of the conventionalization of a linguistic metaphor on the distributional similarity value, a quantitative parameter for the detection of these lexicalized metaphors still needs to be established. This could then provide us with an automatic method for identifying figurative senses, from which the production of large-scale lexical resources would benefit greatly.

In addition, the analysis of similarities using the closest words in the semantic space has shown some interesting phenomena related to metaphoricity. We were able to detect potential metaphorical expressions through a semi-quantitative analysis by using the semantic space to produce the mappings and quantify the similarity, and by then qualitatively evaluating which ones were supposed to have metaphorical potential. However, the analysis suggests that the exploration of the neighbourhood of the terms of a metaphor could be a useful tool for metaphor generation and, in general, computational creativity.

Funding: Seventh Framework Programme, (Grant/Award Number: ‘FP7-PEOPLE-2013-IEF, COGVIM n° 629076’).

References

- Artstein, Ron & Massimo Poesio. 2008. Inter-coder agreement for computational linguistics. *Computational Linguistics* 34(4). 555–596.

- Baroni, Marco, & Alessandro Lenci. 2010. Distributional memory: A general framework for corpus-based semantics. *Computational Linguistics* 36(4). 673–721.
- Black, Max. 1979. More about metaphor. In Andrew Ortony (ed.), *Metaphor and thought*, 19–43. Cambridge: University Press.
- Bolognesi, Marianna. 2016. Using semantic feature norms to investigate how the visual and the verbal modes afford metaphor construction and expression. *Language and Cognition* 27. 1–28.
- Bolognesi, Marianna, Romy van den Heerik & Esther van den Berg. under review. VisMet: An online corpus of visual metaphors. In G. Steen (ed.), *Visual metaphor: Structure and Process*. Amsterdam: John Benjamins.
- Bowdle, Brian & Dedre Gentner. 2005. The career of metaphor. *Psychological Review* 112. 193–216.
- Deerwester, Scott, Susan Dumais & Richard Harshman. 1990. Indexing by latent semantic analysis. *Journal of the American society for information science*. 41(6). 391–407.
- Del Tredici, Marco & Nuria Bel. 2016. Assessing the potential of metaphoricity of verbs using corpus data. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, 4573–4577.
- Firth, John Rupert. 1957. A synopsis of linguistic theory 1930–1955. *Studies in Linguistic Analysis (special volume of the Philological Society)* 1952–1959. 1–32.
- Forceville, Charles. 1996. *Pictorial metaphors in advertising*. London: Routledge.
- Forceville, Charles. 2005. Visual representations of the idealized cognitive model of anger in the Asterix album La Zizanie. *Journal of Pragmatics* 37. 69–88.
- Forceville, Charles. 2011. The JOURNEY metaphor and the source-path-goal schema in Agnès Varda's autobiographical gleaning documentaries. In Monika Fludernik (ed.), *Beyond cognitive metaphor theory: Perspectives on literary metaphor*, 281–297. London: Routledge.
- Forceville, Charles & Eduardo Urios-Aparisi (eds.). 2009. *Multimodal metaphor*. Berlin: Mouton de Gruyter.
- Giora, Rachel. 2008. Is metaphor unique? In Raymond Gibbs, Jr (ed.), *The Cambridge handbook of metaphor and thought*, 143–160. Cambridge, UK: University Press.
- Glucksberg, Samuel. 2001. *Understanding figurative language: From metaphors to idioms*. New York: Oxford University Press.
- Goodall, Catherine, Michael Slater & Teresa Myers. 2013. Fear and anger responses to local news coverage of alcohol-related crimes, accidents, and injuries: Explaining news effects on policy support using a representative sample of messages and people. *Journal of Communication* 63. 373–392.
- Harris, Zellig. 1954. Distributional structure. *Word* 10(2). 146–162.
- Hidalgo, Laura & Blanca Kraljevic. 2011. Multimodal metonymy and metaphor as complex discourse resources for creativity in ICT advertising discourse. In Francisco González García, María Sandra Peña & Lorena Pérez-Hernández (eds.), *Metaphor and metonymy revisited beyond the contemporary theory of metaphor*, 153–178. Amsterdam & Philadelphia: John Benjamins.
- Jackendoff, Ray. 2002. *Foundations of language*. Oxford: University Press.
- Kintsch, Walter. 2000. Metaphor comprehension: A computational theory. *Psychonomic Bulletin & Review* 7. 257–266.
- Lakoff, George & Mark Johnson. 1980. *Metaphors we live by*. Chicago: University Press.

- Landauer, Thomas & Susan Dumais. 1997. A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological review* 104(2). 211–240.
- Lenci, Alessandro. 2008. Distributional semantics in linguistic and cognitive research. *Italian journal of linguistics* 20(1). 1–31.
- McGlone, Matthew. 2007. What is the explanatory value of a conceptual metaphor? *Language and Communication* 27. 109–126.
- Miller, George & Walter Charles. 1991. Contextual correlates of semantic similarity. *Language and cognitive processes* 6(1). 1–28.
- Mitchell, William. 1994. *Picture theory: Essays on verbal and visual representation*. Chicago: University Press.
- Murphy, Gregory. 1996. On metaphoric representation. *Cognition* 60(2). 173–204.
- Ng, Carl & Veronika Koller. 2013. Deliberate conventional metaphor in images: The case of corporate branding discourse. *Metaphor and Symbol* 28(3). 131–147.
- Ortiz, Maria. 2011. Primary metaphors and monomodal visual metaphors. *Journal of Pragmatics* 43. 1568–1580.
- Pérez Hernández, Lorena. 2014. Cognitive grounding for cross-cultural commercial communication. *Cognitive Linguistics* 25(2). 203–247.
- Perez-Sobrinho, Paula. 2016. Multimodal metaphor and metonymy in advertising: A corpus-based account. *Metaphor and Symbol* 31(2). 73–90.
- Phillips, Barbara & Edward McQuarrie. 2004. Beyond visual metaphor: A new typology of visual rhetoric in advertising. *Marketing Theory* 4. 113–136.
- Shutova, Ekaterina. 2015. Design and evaluation of metaphor processing systems. *Computational Linguistics* 41(1). 579–623.
- Simmons, Joseph, Leif Nelson & Uri Simonsohn. 2011. False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science* 22(11). 1359–1366.
- Šorm, Ester & Gerard Steen. under review. VISMIP: Towards a method for visual metaphor Identification. In Gerard Steen (ed.), *Visual metaphor: How images construct metaphorical meaning*. Amsterdam: John Benjamins Publishing Company.
- Steen, Gerard. 2013. Deliberate metaphor affords conscious metaphorical cognition. *Journal of Cognitive Semiotics* 5(1). 179–197.
- Steen, Gerard, Lettie Dorst, Berenike Herrmann, Anna Kaal, Tina Krennmayr & Tryntje Pasma. 2010. *A method for linguistic metaphor identification: From MIP to MIPVU*. Amsterdam: John Benjamins.
- Tukey, John. 1949. Comparing individual means in the analysis of variance. *Biometrics* 5. 99–114.
- Turner, Mark & Gilles Fauconnier. 2002. *The way we think. Conceptual blending and the mind's hidden complexities*. New York: Basic Books.
- Turney, Peter. 2006. Similarity of semantic relations. *Computational Linguistics* 32(3). 379–416.
- Turney, Peter & Patrick Pantel. 2010. From frequency to meaning: Vector space models of semantics. *Journal of artificial intelligence research* 37(1). 141–188.
- Utsumi, Akira. 2011. Computational exploration of metaphor comprehension processes using a semantic space model. *Cognitive Science* 35(2). 251–296.
- van Weelden, Lisanne, Alfons Maes, Joost Schilperoord & Marc Swerts. 2012. How object shape affects visual metaphor processing. *Experimental Psychology* 59(6). 364–371.

Veale, Tony, Ekaterina Shutova & Beata Klebanov. 2016. *Metaphor: A computational perspective. Synthesis lectures on human language technologies*. San Raphael, CA: Morgan and Claypool Publishers.

Vecchi, Eva Maria, Marco Baroni & Roberto Zamparelli. 2011. (Linear) maps of the impossible: Capturing semantic anomalies in distributional space. In *Proceedings of the Workshop on Distributional Semantics and Compositionality*, 1–9.

Appendices

A. Similarity values for visual and linguistic metaphors

Linguistic Metaphors

Target	Source	Similarity	Target	Source	Similarity
accumulation-n	river-n	0.07	homeland-n	house-n	0.31
army-n	motion-n	0.13	idea-n	object-n	0.3
aspect-n	surface-n	0.17	idea-n	point-n	0.31
aspect-n	money-n	0.08	institution-n	equipment-n	0.19
attention-n	eye-n	0.2	judgment-n	finger-n	0.07
attitude-n	gas-n	0.09	knowledge-n	brightness-n	0.16
body-n	canvas-n	0.2	manner-n	path-n	0.05
center-n	heart-n	0.21	opinion-n	picture-n	0.24
condition-n	object-n	0.29	opinion-n	eye-n	0.1
consequence-n	food-n	0.1	opportunity-n	door-n	0.1
consideration-n	appearance-n	0.15	organization-n	building-n	0.25
constraint-n	obstacle-n	0.35	partner-n	food-n	0.15
decision-n	movement-n	0.22	plant-n	person-n	0.29
decision-n	path-n	0.16	possibility-n	space-n	0.27
discipline-n	place-n	0.2	provider-n	origin-n	0.09
discussion-n	war-n	0.29	purpose-n	destination-n	0.12
emotion-n	tide-n	0.13	rank-n	location-n	0.18
emotion-n	force-n	0.18	reason-n	point-n	0.16
end-n	cloth-n	0.05	reason-n	location-n	0.18
explanation-n	drawing-n	0.24	rubbish-n	feces-n	0.18
factory-n	organism-n	0.16	situation-n	air-n	0.14
feces-n	food-n	0.32	success-n	condition-n	0.2
food-n	gold-n	0.18	time-n	frame-n	0.17
governance-n	force-n	0.13	toy-n	girl-n	0.29
harsh-j	hard-j	0.26	understand-v	see-v	0.45

Visual Metaphors

Target	Source	Similarity	Target	Source	Similarity
<i>airplane-n</i>	<i>bread-n</i>	0.1	<i>globe-n</i>	<i>ice-cream-n</i>	0.04
<i>airplane-n</i>	<i>swan-n</i>	0.17	<i>hand-n</i>	<i>fork-n</i>	0.19
<i>america-n</i>	<i>crocodile-n</i>	0.06	<i>hand-n</i>	<i>zebra-n</i>	0.06
<i>bank-n</i>	<i>beggar-n</i>	0.14	<i>hand-n</i>	<i>bulldozer-n</i>	0.12
<i>barcode-n</i>	<i>jail-n</i>	0.02	<i>jeep-n</i>	<i>rhino-n</i>	0.08
<i>beggar-n</i>	<i>bomb-n</i>	0.07	<i>kid-n</i>	<i>piglet-n</i>	0.22
<i>bomb-n</i>	<i>flower-n</i>	0.16	<i>missile-n</i>	<i>dove-n</i>	0.1
<i>book-n</i>	<i>tree-n</i>	0.14	<i>mouth-n</i>	<i>onion-n</i>	0.03
<i>bottle-n</i>	<i>bullet-n</i>	0.18	<i>mouthwash-n</i>	<i>bomb-n</i>	0.09
<i>brain-n</i>	<i>newspaper-n</i>	0.11	<i>newspaper-n</i>	<i>manhole-n</i>	0.05
<i>brain-n</i>	<i>turtle-n</i>	0.12	<i>octopus-n</i>	<i>tire-n</i>	0.08
<i>car-n</i>	<i>dolphin-n</i>	0.16	<i>pen-n</i>	<i>bullet-n</i>	0.21
<i>car-n</i>	<i>horse-n</i>	0.43	<i>pen-n</i>	<i>thermometer-n</i>	0.35
<i>car-n</i>	<i>pepper-n</i>	0.08	<i>perfume-n</i>	<i>doorway-n</i>	0.04
<i>cart-n</i>	<i>jail-n</i>	0.1	<i>president-n</i>	<i>sun-n</i>	0.05
<i>cart-n</i>	<i>tank-n</i>	0.34	<i>president-n</i>	<i>lion-n</i>	0.12
<i>cigarette-n</i>	<i>maze-n</i>	0.03	<i>radio-n</i>	<i>beggar-n</i>	0.08
<i>coke-n</i>	<i>dandelion-n</i>	0.06	<i>seagull-n</i>	<i>book-n</i>	0.06
<i>country-n</i>	<i>bomb-n</i>	0.06	<i>seaweed-n</i>	<i>plastic-n</i>	0.22
<i>country-n</i>	<i>drain-n</i>	0.1	<i>skin-n</i>	<i>match-n</i>	0.06
<i>cream-n</i>	<i>dandelion-n</i>	0.1	<i>sun-n</i>	<i>yolk-n</i>	0.03
<i>cup-n</i>	<i>switch-n</i>	0.2	<i>sweater-n</i>	<i>gorilla-n</i>	0.05
<i>dollar-n</i>	<i>carpet-n</i>	0.1	<i>terrorist-n</i>	<i>match-n</i>	0.07
<i>dove-n</i>	<i>target-n</i>	0.09	<i>tablet-n</i>	<i>trainer-n</i>	0.11
<i>elephant-n</i>	<i>trumpet-n</i>	0.07	<i>wheel-n</i>	<i>clock-n</i>	0.32

B. Conventional linguistic metaphors

Target	Source	Target	Source
<i>aspect-n</i>	<i>surface-n</i>	<i>idea-n</i>	<i>point-n</i>
<i>attention-n</i>	<i>eye-n</i>	<i>manner-n</i>	<i>path-n</i>
<i>center-n</i>	<i>heart-n</i>	<i>opinion-n</i>	<i>picture-n</i>
<i>constraint-n</i>	<i>obstacle-n</i>	<i>opportunity-n</i>	<i>door-n</i>
<i>discussion-n</i>	<i>war-n</i>	<i>possibility-n</i>	<i>space-n</i>
<i>explanation-n</i>	<i>drawing-n</i>	<i>purpose-n</i>	<i>destination-n</i>
<i>governance-n</i>	<i>force-n</i>	<i>reason-n</i>	<i>point-n</i>
<i>harsh-j</i>	<i>hard-j</i>	<i>understand-v</i>	<i>see-v</i>

C. Reformulated visual metaphors

Target	Source	Similarity	Target	Source	Similarity
<i>gift-n</i>	<i>separation-n</i>	0.17	<i>earth-n</i>	<i>ice-cream-n</i>	0.04
<i>airline-n</i>	<i>swan-n</i>	0.17	<i>help-n</i>	<i>hand-n</i>	0.16
<i>america-n</i>	<i>crocodile-n</i>	0.06	<i>help-n</i>	<i>hand-n</i>	0.16
<i>bank-n</i>	<i>beggar-n</i>	0.14	<i>hand-n</i>	<i>machine-n</i>	0.27
<i>market-n</i>	<i>jail-n</i>	0.1	<i>jeep-n</i>	<i>rhino-n</i>	0.08
<i>poverty-n</i>	<i>bomb-n</i>	0.06	<i>kid-n</i>	<i>piglet-n</i>	0.22
<i>war-n</i>	<i>peace-n</i>	0.17	<i>war-n</i>	<i>peace-n</i>	0.17
<i>knowledge-n</i>	<i>plant-n</i>	0.19	<i>breath-n</i>	<i>onion-n</i>	0.03
<i>alcohol-n</i>	<i>weapon-n</i>	0.21	<i>mouthwash-n</i>	<i>explosive-n</i>	0.19
<i>intelligence-n</i>	<i>reading-n</i>	0.32	<i>news-n</i>	<i>rubbish-n</i>	0.16
<i>brain-n</i>	<i>turtle-n</i>	0.12	<i>tentacle-n</i>	<i>tire-n</i>	0.12
<i>car-n</i>	<i>dolphin-n</i>	0.16	<i>signature-n</i>	<i>weapon-n</i>	0.32
<i>car-n</i>	<i>horse-n</i>	0.43	<i>pen-n</i>	<i>thermometer-n</i>	0.35
<i>car-n</i>	<i>pepper-n</i>	0.08	<i>perfume-n</i>	<i>escape-n</i>	0.12
<i>market-n</i>	<i>jail-n</i>	0.1	<i>leader-n</i>	<i>sun-n</i>	0.09
<i>market-n</i>	<i>tank-n</i>	0.19	<i>president-n</i>	<i>predator-n</i>	0.16
<i>smoking-n</i>	<i>maze-n</i>	0.03	<i>radio-n</i>	<i>beggar-n</i>	0.08
<i>low-fat-j</i>	<i>light-j</i>	0.13	<i>reading-n</i>	<i>flight-n</i>	0.33
<i>country-n</i>	<i>bomb-n</i>	0.06	<i>seaweed-n</i>	<i>rubbish-n</i>	0.19
<i>country-n</i>	<i>drain-n</i>	0.1	<i>skin-n</i>	<i>fire-n</i>	0.1
<i>cream-n</i>	<i>dandelion-n</i>	0.1	<i>sun-n</i>	<i>life-n</i>	0.09
<i>coffee-n</i>	<i>switch-n</i>	0.14	<i>sweater-n</i>	<i>gorilla-n</i>	0.05
<i>economy-n</i>	<i>carpet-n</i>	0.07	<i>terrorist-n</i>	<i>fire-n</i>	0.15
<i>peace-n</i>	<i>target-n</i>	0.16	<i>technology-n</i>	<i>trainer-n</i>	0.22
<i>elephant-n</i>	<i>trumpet-n</i>	0.07	<i>wheel-n</i>	<i>clock-n</i>	0.32