

Word Similarity Perception: an Explorative Analysis

Alice Ruggeri (ruggeri@di.unito.it)

Centre for Cognitive Science, University of Turin

Loredana Cupi (loredana.cupi@unito.it)

Department of Foreign Languages and Literatures and Modern Cultures, University of Turin

Luigi Di Caro (dicaro@di.unito.it)

Department of Computer Science, University of Turin

Abstract

Natural language is a medium for expressing things belonging to conceptual and cognitive levels, made of words and grammar rules used to carry semantics. However, its natural ambiguity is the main critical issue that computational systems are generally asked to solve. In this paper, we propose to go beyond the current conceptualization of *word similarity*, i.e., the building block of disambiguation at computational level. First, we analyze the origin of the *perceived* similarity, studying how conceptual, functional, and syntactic aspects influence its strength. We report the results of a two-stages experiment showing clear similarity perception patterns. Then, based on the insights gained in the cognitive tests, we developed a computational system that automatically predicts word similarity reaching high levels of accuracy.

Introduction

Words are symbolic entities referring to something which fills a portion of an autopoietic space made of conceptual, cognitive and contextual information. These three aspects are fundamental to understand the meaning ascribed to linguistic expressions.

One of the most important building block in almost all Computational Linguistics tasks is the computation of similarity scores between texts at different levels: words, sentences and discourses. Since manual numeric annotations of word-word similarity revealed a low agreement between the annotators, cognitive studies can help improve computational systems by discovering *what lies behind the perception of similarity* between words and their referenced concepts. The concept of similarity has been extensively studied in the Cognitive Science community, since it is fundamental in the human cognition. We tend to rely on similarity to generate inferences and categorize objects into kinds when we do not know exactly what properties are relevant, or when we cannot easily separate an object into separate properties. When specific knowledge is available, then a generic assessment of similarity is less relevant (G. L. Murphy & Medin, 1985).

Since words co-occur in textual representations (mutually influencing one each other) it is possible to make experiments on contextual information to analyze what and how influence the perception of similarity. Let us consider two words as two mental representations. The intersection between them can be seen as the context that may also help define the correct similarity of the words in a text. For example, *sugar* and *salt* can be easily associated to the context *kitchen*, whereas *salt* and *sea* intersect in another part of the mental representations.

From a computational perspective, being words ambiguous by nature, the disambiguation process (i.e., Word Sense Disambiguation) is one of the most studied tasks in Computational Linguistics. To make an example, the term *count* can mean many things like *nobleman* or *sum*. Using contextual information, it is often possible to make a choice. Again, this choice is done by means of comparisons among contexts, that are still made of words. In other terms, we may state that the *computational* part of almost all computational linguistics research is about the calculus of matching scores between linguistic items, i.e., *words similarity*. But *what's behind words similarity*?

There exist many annotated data related to similarity and relatedness between words, like wordsim-353 (Finkelstein et al., 2001) and SimLex-999 (Hill, Reichart, & Korhonen, 2014). A large part of the proposed computational systems aims at finding *relatedness* between words, instead of similarity. Relatedness is more general than similarity, since it refers to a generic correlation (like *cradle* and *baby*, that are words representing dissimilar concepts which, however, share similar contexts).

One problem of similarity, as often faced in literature or annotated in datasets, is that it cannot be a static value. Indeed, as the authors of these resources state in their works, the agreement between the annotators is usually not high (around 50-70%). The reason is trivial, however: people can give different degrees of importance with respect to the specific characteristics of the concepts to compare. If we ask one to say how much *dog* is similar to *cat*, the right answer can only be *it depends*. While we can all agree about the fact that the concept *dog* is quite similar to *cat*, we cannot say 0.7 rather than 0.9 (in the range [0,1]) with certainty. Different aspects can be taken into account: are we measuring the form of the animal, or its behaviour? In both cases, it depends on which part of the animal and which actions we are considering to make a choice. For instance, dogs use to return thrown objects. From this point of view, dogs and cats are dissimilar.

In the light of this, our contribution provides the basis for understanding what lies behind a similarity between words and their referenced concepts. First, we analyze syntactic, conceptual and functional aspects of the similarity perception; then, we develop a computational system which is able to predict similarity by leveraging contextual information.

The Cognitive Experiment

In this work, we present two tests to analyze how linguistic constructions are perceived by humans in terms of strength of semantic similarity and if there exists a functionality-based connection that has an influence on its perception. The experiment was presented to 96 users, having different ages and professions, without any particular cognitive or linguistic disorder.

Test on single words

The first test of the experiment regards the perception of the similarity between single words¹. In particular, the goal was to analyze how the users focus on the functional links between the words, and more importantly if such functional-based similarity is a preferential perception channel compared to the conceptual-based one.

Words are ambiguous, and many resources have been released with the goal of defining all the possible senses of a word (i.e., WordNet). Word Sense Disambiguation (Bhattacharyya & Khapra, 2012) is the task of resolving the ambiguity of a word in a given context. Notice that, in our experiment, we do not need any disambiguation of the words, since this process is embodied in the human cognition, thus the users of the test will autonomously represent their subjective sense to associate to the words under comparison.

Then, since we wanted to compare conceptual with functional preferences, we designed the test as a comparison between two word pairs, one involving conceptually-related words and one with words linked by direct functionalities. To generalize, let us consider the words a , b , and c with the conceptual word pair $a-b$ and the functional word pair $a-c$. The user is asked to mark the most similar word (among b and c) to associate to a , and so the most correlated word pair. The users were not aware of the goal of the test and of the difference between the word pairs.

Since words and actions present a high variability in terms of conceptual range (or their mental representation), we put particular attention to the choice of the word pairs, according to the following principles:

Conceptual granularity If we think at the words *object* and *thing*, we probably do not have enough information to make significant comparisons due to their large and undefined conceptual boundaries. The same happens in cases when two words represent very specific concepts such as *lactose* and *amino acid*. The word pairs of the proposed test have been selected by considering this constraint (and so they include words which are not too specific nor too general).

Concreteness Words may have direct links with concrete objects such as “*table*” and “*dog*”. In other cases, words such as “*justice*” and “*thought*” represent abstract concepts. Since it is not clear how this may affect the per-

¹Notice that “*similarity between words*” is intended as the similarity between the concepts they bring to mind.

ception of similarity, we decided to keep *concrete* words only.

Semantic coherence Another criterion used for the selection of the words was the level of semantic similarity between the word pairs to compare. To better analyze whether the functional aspect plays a significant role in the similarity perception, we extracted conceptual and functional pairs of words which had similar semantic closeness according to a standard semantic similarity calculation. In the light of this, we used a Latent Semantic Space calculated over almost 1 million of documents coming from the collection of literary text contained in the project Gutenberg page². The selected conceptual and functional word pairs had the property of having a very close semantic similarity (the score differences were less than 0,01 in a [0,1] range).

The test was composed by three word pairs, to leave it simple and to be not affected by users tiredness. Then, instead of randomly selecting three different word pairs, we wanted to consider three cases in which the functional links between the words have distinct levels of importance. Our assumption was that the more the importance of the functional link between two words in a pair, the more its perceived similarity (and thus the user preferences with respect to the conceptual word pair). For this reason we added a final criterion:

Increasing relevance of the functional aspect To estimate the importance of the functional aspect that relates two words we analyzed the number of actions (or verbs) in which they are usually involved with. In our test, the functional word pairs *salt-water*, *nail-polish*, and *ring-finger* have a functional link of 0.003³, 0.0101⁴, and 0.0625⁵ respectively (see Table 1). These values are calculated in the following way: given the total number of existing verbs $NV(rw)$ for the root word rw and the number of effective usages $EU(sw)$ with the second word of the pair sw , we computed the functional link $Fl(rw, sw)$ of the functional pair as $EU(sw) / NV(rw)$.

Table 1: The chosen word pairs in the first test.

Root word rw	Conceptual Pair rw - sw	Functional Pair [F. link] rw - sw [Fl(rw,sw)]
salt	salt - sugar	salt - water [0.003]
nail	nail - finger	nail - polish [0.0101]
ring	ring - necklace	ring - finger [0.0625]

We then considered a backup test set with random word pairs matching with the same above-mentioned criteria, collecting a total of 24 answers on 24 different cases of word

²<http://promo.net/pg/>

³For the verbs “*to put*”, “*to add*” and “*to get*”

⁴For the verbs “*to apply*” and “*to use*”

⁵For the verbs “*to put*” and “*to wear*”

pairs such as the main one of Table 1. This was done to prove the reliability of the test, seeing whether the results and the analyses show a similar trend, being independent from the selection of the words. The results of the whole test is described in the final part of this section.

Test on Phrases

The second test of the experiment concerns the perception of the similarity between phrases, or multi-word linguistic constructions⁶. The goal was to analyze how the syntactic context of a target word influences the perception of similarity among entire phrases. More in detail, we wanted to discover possible differences of such perception along different syntactic roles. We considered a simple syntactic structure of the type *subject-verb-object*.

Given a root sentence such as “*Mario sings the song*”, we created three variations by changing the subject, the verb, and the direct object. For example, by changing “*Mario*” with “*The bird*” we obtained “*The bird sings the song*”. The complete set of replacements are shown in Table 3.

We presented to the 96 users a total of 4 sentences (see Table 2), that with its 4 variations produce a total of 16 pairs of sentences to be analyzed by the users in terms of perceived similarity, as in the first test. For each sentence, the users had to indicate the degree of similarity of the original sentence with one of its variations using a value in the range [0,10]⁷. 0 means no semantic similarity between the two phrases and 10 means total equality. The grammatical changes made on the original sentences were chosen maintaining the semantic validity (i.e., all the sentences represent valid mental representations).

Table 2: The chosen phrases in the second test.

Phrase ID	Phrase
(a)	Mario sings the song
(b)	Alan drives the car
(c)	Alice writes the book
(d)	Marco does the homeworks

Table 3: The word replacements for subjects (SC), verbs (VC) and direct objects (OC).

Replacement	(a)	(b)	(c)	(d)
SC	bird	robot	computer	software
VC	writes	cleans	cleans	gives
OC	verse	band	sheet	pasta

⁶Even in this case, “*similarity*” is intended as the similarity between the concepts related to the phrases.

⁷We used a [0,10] range instead of a [0,1] range as in the previous test because it represents a more human-understandable and intuitive notation.

Interpretation of the results

In this section, we give a preliminary interpretation of the results on the collected answers.

In the first test on single words, we can state that, generally, conceptual and functional word pairs are differently perceived according to the importance of the functional link in the functional word pair. This shows that words and their referenced concepts are mainly compared in terms of conceptual similarity, but when there exists important functionalities between them, this influences the users preference towards the functional word pair ??.

For example, the similarity of the *sugar-salt* pair results to be stronger compared to the *water-salt* one, since the action *to add/put the salt in the water* is “*a needle in a haystack*” with respect to all the actions related to *water* and *salt* independently. This means that there is no exclusive action between *water* and *salt* (i.e., there are many actions that involve *water*). An opposite example is represented by the word pair *ring-finger*, since the action *to put/wear the ring on the finger* is much more exclusive than in the previous case. Such preference could be explained by stating that all word pairs, especially with words that underlie actions, have a strong visual representation that makes them quickly perceivable.

Table 4: Results showing the percentage of preferences in the choice of the most (perceived) correlated word pairs of the first test.

Case	Word pairs	N. of preferences	%
1a.	salt - sugar	75	78%
1b.	salt - water	21	22%
2a.	nail - finger	44	46%
2b.	nail - polish	52	54%
3a.	ring - necklace	16	17%
3b.	ring - finger	80	83%

Table 5: Results on the backup of the first test (with 24 different cases including 8 low-FL cases, 8 medium-FL cases and 8 high-FL cases). The results are in line with the ones of the main test shown in Table 4.

Funct. Pair	Pref. w.r.t conceptual pair
Funct. Pair (low FL)	1 out of 8 (12.5%)
Funct. Pair (medium FL)	5 out of 8 (62.5%)
Funct. Pair (high FL)	7 out of 8 (87.5%)

This result is also in line with what stated by (Cohen et al.,2002), i.e., words that have a functionality-based relationship can have a more complex visual component that makes such correlation weaker.

In Figure 1, we show the users preferences for the second test. In the case of verb replacement (VC) we can notice a high meaning change in terms of similarity perception (similarity values close to 0), so the verb represents the real root

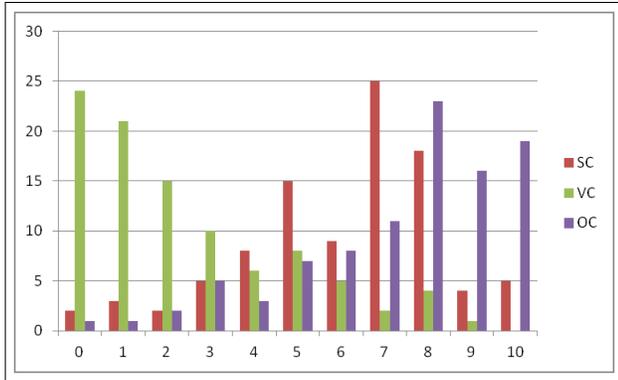


Figure 1: Results of the second test, showing the change scores in terms of word similarity perception after subject, verb and object replacements. SC stands subject change, VC for verb change, and OC for object change.

of the mental representations. The case of the subject change (SC) shows a less important decrease of similarity perception, while the object change (OC) resulted to be the less relevant syntactic role influencing the meaning of the whole phrase.

The Computational Analysis

In the previous section we studied the role of the context (on different levels) within the process of word similarity perception. Since the results indicated that both functional aspects and syntactic roles have an impact on how people perceive similarity, we experimented a computational approach for the automatic estimation of the similarity based on functional and syntax-aware contextual information.

In particular, we used the large and freely-available semantic resource ConceptNet⁸. A partial overview of the semantic knowledge contained in ConceptNet is illustrated in Table 6. ConceptNet is a resource based on common-sense rather than linguistic knowledge since it contains much more function-based information (e.g., all the actions an object can or cannot do) contained in even complex syntactic structures. The idea is also to exploit users perception of reality (the actual origin of ConceptNet) instead of the result of top-down expert building of ontologies (e.g., WordNet). ConceptNet contains important semantic problems related to coverage, utility of semantic information and coherence, but we used it as a black box due to its largeness and common-sense nature. A deep analysis of this resource is out of the scope of this paper.

The experiment started from the transformation of a word-word-score similarity dataset into a context-based dataset in which the words are replaced by sets of semantic information taken from ConceptNet. The aim was to figure out which semantic facts *make* the similarity between two words *perceivable*.

We used the dataset SimLex-999 (Hill et al., 2014) that contains one thousand word pairs that were manually annotated with similarity scores. The inter-annotation agreement is 0.67

⁸<http://conceptnet5.media.mit.edu/>

Table 6: Some of the existing relations in ConceptNet, with example sentences in English.

Relation	Example sentence
IsA	NP is a kind of NP.
LocatedNear	You are likely to find NP near NP.
UsedFor	NP is used for VP.
DefinedAs	NP is defined as NP.
HasA	NP has NP.
HasProperty	NP is AP.
CapableOf	NP can VP.
ReceivesAction	NP can be VP.
HasPrerequisite	NP—VP requires NP—VP.
MotivatedByGoal	You would VP because you want VP.
MadeOf	NP is made of NP.
...	...

(Spearman correlation). We leveraged ConceptNet to retrieve the semantic information associated to the words of each pair, then keeping the intersection. For example, considering the pair *rice-bean*, ConceptNet returns the following set of semantic information for the term *rice*:

[hasproperty-edible, isa-starch, memberof-oryza, atlocation-refrigerator, usedfor-survival, atlocation-atgrocerystore, isa-food, isa-domesticateplant, relatedto-grain, madeof-sake, isa-grain, receivesaction-cook, atlocation-pantry, atlocation-ricecrisp, atlocation-supermarket, ...]

Then, the semantic information for the word *bean* are:

[usedfor-fillbeanbagchair, atlocation-infield, atlocation-can, usedfor-nutrition, usedfor-cook, atlocation-atgrocerystore, usedfor-grow, atlocation-foodstore, isa-legume, usedfor-count, isa-domesticateplant, atlocation-cookpot, atlocation-beansoup, atlocation-soup, isa-vegetable, ...]

Finally, the intersection produces the following set:

[atlocation-atgrocerystore, isa-domesticateplant, atlocationpantry]

At this point, for each non-empty intersection, we created one instance of the type:

<semantic information>, <similarity score>

and computed a standard *term-document* matrix, where the *term* is a semantic term within the set of semantic information retrieved from ConceptNet and the *document* dimension represents the word pairs of the original dataset. After this preprocessing phase, the *score* attribute is discretized into two bins:

- *non-similar* class - range in the dataset [0, 5]

- *similar* class - range in the dataset [5.1, 10]

The splitting of the data into two clusters allowed us to experiment a classic supervised classification system, where a Machine Learning tool (a Support Vector Machine, in our case) has been used to learn a binary model for automatically classifying *similar* and *non-similar* word pairs. The result of the experiment is shown in Table 7. Noticeably, the classifier has been able to reach a quite good accuracy (65.38% of correctly classified word pairs), considering that the inter-annotation agreement of the original data is only 0.67 (Spearman correlation).

Table 7: Classification results in terms of Precision, Recall, and F-measure. The total accuracy is 65.38%.

Precision	Recall	F-Measure	Class
0,697	0,475	0,565	<i>non-similar</i>
0,633	0,815	0,713	<i>similar</i>
0,664	0,654	0,643	weighted total

Notice that *similar* word pairs are generally easier to identify with respect to *non-similar* ones.

Related Work

This paper presents an idea which combines linguistic, cognitive and computational perspectives. In this section, we mention those theoretical and empirical methods that inspired our motivational basis.

Linguistic Background

The difficulty of defining the *meaning of meaning* has to do with some tricky issues like lexical ambiguity and polysemy, vagueness, contextual variability of word meaning, etc. As a matter of fact, words are organized in lexicon as a complex network of semantic relation which are basically subsumed within Saussure’s paradigmatic (the axis of combination) and syntagmatic (the axis of choice) axes (Saussure,1983).

Some authors (Chaffin & Herrmann,1984) have already suggested theoretical and empirical taxonomies of semantic relations consisting of some main families of relation (such as contrast, similars, class inclusion, part-whole, etc.). As Murphy points out (M. L. Murphy,2003), lexicon has become more central in linguistic theories and, even if there is no a widely accepted theory on its internal semantic structure and how lexical information are represented in it, the semantic relations among words are considered in scholarly literature as relevant to the structure of both lexical and conceptual information and it is generally believed that relations among words determine meaning.

Cognitive Background

Although words perception could seem immediate, the input we perceive is recognized and trasformed mediating background and contextual information, within a dynamic and cooperative process. The well-known semiotic triangle (Ogden,

Richards, Malinowski, & Crookshank,1946) introduced by different authors over time represents a first reference for our study. People use symbols (our words) to communicate meanings (the effective content). The meaning is something untangible, which can be though even without any concrete presence. The last point is then the physical reference, i.e., the object in the reality⁹. Note that there is no connection between symbols and references, since only imagined meanings can allow the two to be linked.

Interaction is another important aspect that has been investigated in literature. Indeed, the actions change the type of perception of an object, which models itself to fit with the context of use. Then, the Gestalt theory (Köhler,1929) contains different notions about the perception of meaning according to interaction and context. In particular, the core of the model is the complementarity between the *figure* and the *ground*. In our case, a word is the figure and the ground is the context that lets emerge its specific sense. Finally, James Gibson introduced the concept of affordances as the cognitive cues that an object exposes to the external world, indicating ways of use (Gibson,1977). In cognitive and computational linguistics, this theory can be inherited to model words as objects and contexts as their interaction with the world.

Computational Background

In this section, we review the main works that are related to our contribution from a computational perspective. Natural Language Processing represents an active research community whose focus is letting machines communicate by understanding semantics within linguistic expressions. Ontology Learning (Cimiano,2006) is the task of automatic extracting structured semantic knowledge from texts, and it well fits the scope of this paper. Nevertheless, Word Sense Disambiguation (WSD) (Stevenson & Wilks,2003) is maybe the most related NLP task, whose aim is to capture the correct meaning of a word in a context. Generally speaking, many other tasks have the problem of comparing linguistic items in order to make choices to pass from syntax to semantics. Named Entity Recognition (NER) (Nadeau & Sekine,2007;Marrero, Urbano, Sánchez-Cuadrado, Morato, & Gómez-Berbís,2013) is the task of identifying entities like people, organizations and locations in texts. This is often done by comparing words in contexts to some learned patterns. In general, many other NLP tasks are based on the evaluation of similarity scores (Manning & Schütze,1999).

Nowadays, there exists a large set of available semantic resources that can be used in Natural Language Processing techniques in order to understand the hidden meaning of perceived similarity between two words or concepts. For example, ConceptNet contains semantic information that are usually associated with common terms (even if not correctly disambiguated). By analyzing the relationship between annotated similarity scores and semantic information it is possi-

⁹The existing terminology is quite varying: symbol-thought/reference/referent (Aristotele); object-representation-interpretant (Peirce); signified-sign-referent (De Saussure)

ble to create predictive models which automatically deduce words similarity by dynamically weighting words features based on their mutual interaction.

If we consider the objects / agents / actions to be terms in text sentences, we can try to extract their meaning and semantic constraints by using the idea of affordances. For instance, let us think to the sentence “*The squirrel climbs the tree*”. In this case, we need to know what kind of subject “*squirrel*” is to figure out (and visually imagine) how the action will be performed. According to this, no particular issues come out from the reading of this sentence. Let us now consider the sentence “*The elephant climbs the tree*”. Even if the grammatical structure of the sentence is the same as before, the agent of the action is different, and it obviously creates some semantic problems. In fact, from this case, some constraints arise; in order to climb a tree, the subject needs to fit to our mental model of something that can climb a tree. In addition, this also depends on the mental model of “*tree*”. Moreover, different agents can be both correct subjects of an action whilst they may produce different meanings in terms of how the action will be mentally performed. Consider the sentences “*The cat opens the door*” and “*The man opens the door*”. In both cases, some implicit knowledge suggests the manner the action is done: while in the second case we may think at the cat that opens the door leaning to it, in the case of the man we probably imagine the use of a door handle. A study of these language dynamics can be of help for many NLP tasks like Part-Of-Speech tagging as well as more complex operations like dependency parsing and semantic relations extraction. Some of these concepts are latently studied in different disciplines related to statistics. Distributional Semantics (DS) (Baroni & Lenci, 2010) represents a class of statistical and linguistic analysis of text corpora that tries to estimate the validity of connections between subjects, verbs, and objects by means of statistical sources of significance.

Conclusions and Future Work

In this paper, we proposed a combined analysis of linguistic, cognitive and computational aspects to assess the nature of words similarity. First, we studied how word similarity perception is influenced in terms of conceptual, functional and syntactic roles. In future work, our aim is to extend the sample of users on more specific cases. Still, by changing the language of the users, we can have results that take into account the cultural ground, understanding if and how word similarity depends on it. On the other side, we stressed the importance of computational understanding of similarity to improve Computational Linguistics tasks which are based on it, usually without any analysis of contextual information. In particular, we used the large semantic knowledge contained in ConceptNet to create a Support Vector Machine classifier to predict word similarity based on an annotated dataset. In future work, we will extend our experimental analysis to validate existing similarity datasets and to produce predictive models for the automatic identification of human-readable

similarity scores.

References

- Baroni, M., & Lenci, A. (2010). Distributional memory: A general framework for corpus-based semantics. *Computational Linguistics*, 36(4), 673–721.
- Bhattacharyya, P., & Khapra, M. (2012). Word sense disambiguation. *Emerging Applications of Natural Language Processing: Concepts and New Research*, 22.
- Chaffin, R., & Herrmann, D. J. (1984). The similarity and diversity of semantic relations. *Memory & Cognition*, 12(2), 134–141.
- Cimiano, P. (2006). *Ontology learning from text*. Springer.
- Cohen, L., Lehericy, S., Chochon, F., Lemer, C., Rivaud, S., & Dehaene, S. (2002). Language-specific tuning of visual cortex? functional properties of the visual word form area. *Brain*, 125(5), 1054–1069.
- Finkelstein, L., Gabrilovich, E., Matias, Y., Rivlin, E., Solan, Z., Wolfman, G., et al. (2001). Placing search in context: The concept revisited. In *Proceedings of the 10th international conference on world wide web* (pp. 406–414).
- Gibson, J. J. (1977). The Theory of Affordances. In R. Shaw & J. Bransford (Eds.), *Perceiving, acting, and knowing: Toward an ecological psychology*. Lawrence Erlbaum.
- Hill, F., Reichart, R., & Korhonen, A. (2014). Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *arXiv preprint arXiv:1408.3456*.
- Köhler, W. (1929). Gestalt psychology.
- Manning, C. D., & Schütze, H. (1999). *Foundations of statistical natural language processing*. MIT press.
- Marrero, M., Urbano, J., Sánchez-Cuadrado, S., Morato, J., & Gómez-Berbís, J. M. (2013). Named entity recognition: Fallacies, challenges and opportunities. *Computer Standards & Interfaces*, 35(5), 482–489.
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological review*, 92(3), 289.
- Murphy, M. L. (2003). *Semantic relations and the lexicon*. Cambridge University Press.
- Nadeau, D., & Sekine, S. (2007). A survey of named entity recognition and classification. *Linguisticae Investigationes*, 30(1), 3–26.
- Ogden, C. K., Richards, I. A., Malinowski, B., & Crookshank, F. G. (1946). *The meaning of meaning*. Harcourt, Brace & World New York.
- Proffitt, D. (2006). Embodied perception and the economy of action. *Perspectives on psychological science*, 1(2), 110–122.
- Saussure, F. d. (1983). Course in general linguistics, trans. R. Harris, London: Duckworth.
- Stevenson, M., & Wilks, Y. (2003). Word-sense disambiguation. *The Oxford Handbook of Comp. Linguistics*, 249–265.