Biologically Inspired Cognitive Architectures (2015) **12**, 121–133

Available at www.sciencedirect.com

ScienceDirect

journal homepage: www.elsevier.com/locate/bica

RESEARCH ARTICLE

Human-inspired semantic similarity between sentences

J. Ignacio Serrano [a,*], M. Dolores del Castillo [a], Jesús Oliva [b], Rafael Raya [a]

[a] *Group of Neural and Cognitive Engineering (gNeC), Centro de Automática y Robótica, Consejo Superior de Investigaciones Científicas (CSIC), Spain*
[b] *BBVA Data & Analytics, Spain*

Received 23 March 2015; received in revised form 10 April 2015; accepted 10 April 2015

KEYWORDS
Cognitive linguistics;
Computational linguistics;
Semantic similarity

Abstract
Following the Principle of Compositionality, the meaning of a complex expression is influenced, to some extent, not only by the meanings of its individual words, but also the structural way the words are assembled. Compositionality has been a central research issue for linguists and psycholinguists. However, it remains unclear how does syntax influence the meaning of a sentence. In this paper, we propose an interdisciplinary approach to better understand that relation. We present an empirical study that seeks for the different weights given by humans to different syntactic roles when computing semantic similarity. In order to test the validity of the hypotheses derived from the psychological study, we use a computational paradigm. We incorporate the results of that study to a psychologically plausible computational measure of semantic similarity. The results shown by this measure in terms of correlation with human judgments on a paraphrase recognition task confirm the different importance that humans give to different syntactic roles in the computation of semantic similarity. This results contrast with generative grammar theories but support neurolinguistic evidence.

Introduction

We, humans, are continuously assessing the similarity of objects in our daily life. As explained by Goldstone (1994), when humans try to judge the similarity of visual scenes, we take into account the structure of the compared objects and the different relationships between the different parts.

* Corresponding author.
E-mail addresses: jignacio.serrano@csic.es (J.I. Serrano), md.delcastillo@csic.es (M.D. del Castillo), jesus.oliva1984@gmail.com (J. Oliva), rafael.raya@csic.es (R. Raya).

http://dx.doi.org/10.1016/j.bica.2015.04.007
2212-683X/

So humans use the structural information in the comparison of general objects, but could this conclusion be generalized to language? To what extent do the different parts of the hierarchical structure of a sentence influence the global meaning? The relations between syntax and semantics have been studied for several years. In particular, compositionality has been largely studied since it was first proposed as the notion that the meaning of an expression is determined by the meaning of its constituents and the way they are combined. This is clearly shown by sentences made up by the same words but with very different semantic interpretations like: ''*The dog bit the man*'' and ''*The man bit the dog*''.

Despite the great amount of work about the Principle of Compositionality and its different interpretations, it is still not clear the real influence of different syntactic roles on the representation of meaning. See the following examples:

(a) That movie made me cry quickly.
    That Movie made me cry slowly.
(b) Than movie made me cry quickly.
    That movie made me laugh quickly.

It seems clear that the two sentences in (a) are more semantically similar than the ones in (b). However, in both cases we have just replaced one word by one of its antonyms. So, how does our brain compute semantic similarity? Are some parts of a sentence more important than others in the computation of semantic similarity? Most studies have centered on the dominant effect of verbs on the meaning of a sentence but there is a lack of work about the relative influence of different syntactic roles. The only study in this direction is the one presented by Wiemer-Hastings (2004). On that paper, Wiemer-Hastings points that human judges tend to ignore similarities between segments with different functional roles, denoting the importance of syntactic structure analysis in the computation of semantic similarity, and claiming that different syntactic roles have a different level of signification in the calculation of semantic similarity by humans.

In this paper, we propose an interdisciplinary approach to better understand how our mind computes semantic similarity and, in particular, the different importance that humans give to different syntactic roles in the computation of semantic similarity. Acoording to Cambria and White (2014), the work presented here aims at explaining how can we jump from the syntactics and semantics curves to the pragmatics one. To this end, we performed a psychological study about how humans compute semantic similarity between sentences and then we use a computational paradigm in order to test the validity of the hypotheses derived from that study. First of all, we present an empirical study that seeks for the different weights given by humans to different syntactic roles when computing semantic similarity. According to the results obtained by Wiemer-Hastings (2004), we start from the hypothesis that different syntactic roles have different importance in the calculation of semantic similarity by humans. Going a step forward, we made a deeper empirical study, based on two experiments with human judges, that complement the experiments carried out by Wiemer-Hastings (2004) and overcome some of their limitations. Our experiments are not restricted to specific domains while the work of Wiemer-Hastings (2004) is focused on the two specific domains: computer literacy and psychological research methods. Moreover we give a more accurate quantitative measure of the different weights given by humans to different syntactic roles while computing semantic similarity.

In order to assess the validity of the conclusions obtained with the experiments carried out with humans, we used a computational paradigm. We incorporated the results of the empirical study to a psychologically plausible semantic similarity measure (Oliva, Serrano, Del Castillo, & Iglesias, 2011) that takes into account the influence of different syntactic roles on the overall sentence meaning. The semantic similarity measure was applied to a paraphrase recognition task (Dolan, Quirk, & Brockett, 2004) using two different combinations of weights obtained from the human judgments for a semantic similarity task and from the human judgments for the same paraphrase recognition task. The results obtained with both versions confirm the different contribution of different syntactic roles on semantic similarity computation. The different variations of the similarity measure tested with the two combinations of weights outperformed their non-weighted counterparts. Moreover, they obtained results similar to the ones reported by Islam and Inkpen (2008) and Mihalcea, Corley, and Strapparava (2006) on the paraphrase recognition task. Furthermore, four of the six approaches tested outperform significantly the method of Mihalcea et al. (2006) and the results of three of them are similar to the ones reported by Islam and Inkpen (2008). Finally, we compared the different weights given by humans to different syntactic roles on different tasks that involve semantic similarity computation. The weights obtained from the human ratings of semantic similarity and the ones obtained from the paraphrase recognition task were very similar, showing that humans tend to use the same weights through different tasks. The interdisciplinary character of this work is not only assessed by the combination of experimental techniques derived from psycholinguistics and computational sciences. Moreover, the contributions of this paper are of interest both from a theoretical and a practical point of view. The importance of sentence semantic similarity measures in natural language research is increasing due to the great number of applications that are arising in many text-related research fields. For example, in text summarization, sentence semantic similarity is used (mainly in sentence-based extractive document summarization) to cluster similar sentences and then find the most representing sentence of each cluster (Aliguliyev, 2009). Also in web page retrieval sentence similarity can improve the effectiveness by calculating similarities of titles of pages (Park, Ra, & Jang, 2005). Conversational agents can also benefit from the use of sentence similarity measures reducing the scripting process by using natural language sentences rather than patterns of sentences (Allen, 1995). These are only a few examples of applications whose effectiveness could improve with sentence semantic similarity calculation. Therefore, our work not only sheds light on the theoretical question of how humans use syntactic roles when computing semantic similarity. It also shows how those results can be straightforwardly used in a psychologically plausible computational system with many practical applications.