<https://helda.helsinki.fi>

Word Sense Disambiguation in Document Space

Linden, Krister

2002-10

Linden , K & Lagus , K 2002 , Word Sense Disambiguation in Document Space . in
Proceedings of 2002 IEEE International Conference on Systems, Man and Cybernetics .
IEEE International Conference on Systems, Man and Cybernetics , Tunisia , 06/10/2002 .

<http://hdl.handle.net/10138/33868>

submittedVersion

Downloaded from Helda, University of Helsinki institutional repository.

This is an electronic reprint of the original article.

This reprint may differ from the original in pagination and typographic detail.

Please cite the original version.

Word Sense Disambiguation in Document Space

Krister Lindén, Krista Lagus

Abstract— We introduce a method for word sense disambiguation that uses an existing topical document map created with an unsupervised method (WEBSOM [1]) on a very large document collection. Results on the SENSEVAL-2 corpus indicate that the method is statistically significantly better than the baselines and on a par with supervised methods.

The method uses the document map as a representation of the semantic space of word contexts. The assumption is that similar meanings of a word have similar contexts, which are located in the same area on the self-organized document map. The results confirm this assumption.

The benefit of the proposed method is that a single general purpose representation of the semantic space can be used for all words and their word senses.

Keywords— Word sense disambiguation, Self-organized document maps, Large corpora, Semantic space

I. INTRODUCTION

WORD sense disambiguation is the task of automatically determining the appropriate senses of a word in context. It is an important and difficult problem with many practical consequences for language-technology applications in information retrieval, document classification, machine translation, spelling correction, parsing, speech synthesis as well as speech recognition. In order to create domain-independent applications in these areas word sense disambiguation is essential. So far many applications solve the problem by letting users manually switch between tailor-made domain-specific vocabularies. Successful word sense disambiguation could at its best provide a seamless transition between domains in natural language processing.

Word sense disambiguation has been approached by traditional AI methods, such as hand-made rule sets or semantic networks, by knowledge-based methods using dictionaries or thesauri, and by corpus-based methods [2]. For a textbook introduction to word sense disambiguation, see [3]. The methods may vary in how different levels of context are selected and encoded. From a linguistic point of view the information included in the representation of context corresponds to approximations of morphological, syntactic and discourse context. Yarowsky [4] noted that there seems to be only one sense per collocation and that words tend to keep the same sense during a discourse. Later Martinez and Agirre [5] confirmed that the one-sense-per-collocation hypothesis holds, but it is weaker on fine-grained sense distinctions. In addition they showed that different topics tend to have different sets of collocations, which reduces the applicability of collocations over topic or genre variations. In this paper we evaluate genre

and topic independent disambiguation effects. We therefore rely heavily on the one-sense-per-discourse hypothesis for single words. We take one discourse to be one document.

Corpus-based methods use supervised or unsupervised machine learning. For a recent comparison of machine learning algorithms for word sense disambiguation see [6] and [7]. For some results of combining supervised and unsupervised methods see [8]. Unsupervised methods operate on untagged words in context, whereas supervised methods use sense-tagged words in context. Schütze [9] has proposed an unsupervised method using clusters of ambiguous words in context. He achieved very good results on a small set of words using only the two main senses of each word. He created the word clusters for the same genre and topics in the training data as in the test data. We propose a hybrid technique which uses a self-organizing map to create a representation of semantic space from a massive independent document collection and then calibrates the representation with a small amount of hand-tagged data from the genre of the test data. The training data for the semantic space and the test data are from unrelated genres.

A mathematical structure of semantic space is discussed in [10]. Formally it is a quadruple $\langle A, B, S, M \rangle$, where B is the set of basis elements, A is the mapping between particular basis elements and each word in the language, S is the similarity measure between vectors of basis elements, and M is a transformation between two semantic spaces, eg. a dimensionality reduction. In [11], Steyvers and Tenenbaum show that large-scale natural language semantic structures such as thesauri are characterized by sparse connectivity and strong local clustering. In [12], Martinetz and Schulten show that self-organizing maps often are a local neighborhood preserving projection from a high dimensional space to a low dimensional display. This leads to an assumption that a document map created by the WEBSOM method [1] could serve as a locally accurate representation of a high dimensional semantic space, since it preserves the strong local clustering characteristic of large semantic structures. A self-organized document map represents the semantic space as ordered clusters of documents. The ordering thereby locally reflects the topical and discourse contexts of words.

In the same sense that a general language thesaurus is domain independent, we take the particular map created in [1] and evaluate our disambiguation method using the fine-grained word sense distinctions of the English SENSEVAL-2 data set [6].

To our knowledge, this is the first time a word sense disambiguation method using an independent large-scale representation of semantic space is reported.

The rest of this article is organized as follows: First we

Krister Lindén is with the University of Helsinki, Department of General Linguistics, P.O.Box 9 (Siltavuorenpenger 20 A), FIN-00014 University of Helsinki, Finland. E-mail: Krister.Linden@helsinki.fi. Krista Lagus is with the Helsinki University of Technology, Neural Networks Research Centre, P.O.Box 9800 (Tammasaarekatu 3), FIN-02015 HUT, Finland. E-mail: Krista.Lagus@hut.fi.

Fig. 1. Word sense disambiguation with a WEBSOM map

briefly discuss the WEBSOM method, then introduce our disambiguation method in Section II. The SENSEVAL-2 corpus and its preprocessing are described in Section III and the disambiguation experiments and results in Section IV. Sections V and VI present the discussion and conclusion, respectively.

A. WEBSOM document maps

The WEBSOM method [13], [1] uses the Self-Organizing Map algorithm [14] to organize a large document collection onto a two-dimensional display called the map. Documents are encoded using the bag-of-words model with word weighting. For computational reasons the dimensionality of the representation is reduced using random projection. The cosine measure (dot product for normalized vectors) is used for measuring similarities between documents.

Documents similar in content are located near each other on the ordered map display. The visualized display along with an HTML interface can be used for exploring the document collection. The SOM consists of a set of map units ordered on a two-dimensional lattice. By virtue of a model vector stored with each map unit, searches can be performed on the map by locating the most similar model vector for a new document or short context.

It has been shown that the WEBSOM method is applicable to various kinds of text collections, including very colloquial ones. The method is scalable, as is shown by the organization of the map of nearly 7 million patent abstracts [1]. For our disambiguation experiments we use this existing large document map.

The size of the two-dimensional patent abstract map used in this experiment is 1,002,240 map units. Mathematical symbols and numbers have been converted to dummy symbols. Originally there were 733,179 different tokens in the document collection, but a cutoff threshold was applied

leaving out words occurring less than 50 times, as well as 1,355 frequent words on a stop list. The words were reduced to their base forms and the remaining number of types was 43,222 words. Each document was positioned in a 43,222-dimensional space. For efficiency of map creation, the dimensionality of the vector space was reduced to 500 by the random projection procedure. [1]

II. WORD SENSE DISAMBIGUATION METHOD

Let us assume that there is a vast general purpose document collection containing a representative sample of many different areas of human endeavor, discussions or other texts representing varying discourses and different topics. One can then organize such a collection on a large document map obtaining a representation of semantic space.

Such a representation of semantic space can be used in word sense disambiguation as follows: A word in context is treated as a small document. As a *calibration* step, each sample of a sense-tagged word in context is encoded as a document vector and positioned on the best-matching map location corresponding to that word and its context. The word and its context are both used in the matching.

When an ambiguous word needs disambiguating the word in context is processed identically as the tagged samples, and the best-matching location in semantic space is found. The tag for the word is selected by a majority vote among the k nearest sense-tagged samples of the same word on the map lattice, cf. Figure 1. In case of a tie between several senses, the globally most frequent of them is chosen. A restriction¹ on the maximum allowed distance d_{max} is set; if a sample lies outside of this limit, it is not con-

¹The restriction on the maximum distance to the nearest neighbors is motivated by the fact that the WEBSOM map is likely to be only a piecewise locally accurate projection of the document clusters in the multi-dimensional semantic space.

sidered in the vote. If no samples lie near enough (within d_{max}), instead of the local decision strategy a global strategy is applied: a majority vote is taken among all the sense tagged samples of that word.

The distance between two words in context is defined as the map lattice distance between their locations on the map. The distance between neighboring map units is one².

A variant of the method uses fuzzy positions on the WEBSOM map. A fuzzy position consists of the l best matching map locations corresponding to the word in context. For small l (e.g. $l < 10$) these locations are generally near each other, but especially for larger l several clusters may emerge. In this variant, the distance between words w_1 and w_2 is defined as the minimum distance between the map locations corresponding to these two words.

Once a suitable map exists, the computational complexity of the method is low. The computational cost is dominated by the search for the best matching location for the calibration and test samples. The search requires $\mathcal{O}(m * w * M)$ operations, where m is the number of calibration and test samples, w is the average number of words in each sample, and M is the number of map units of the map. When the best matching units have been located on the map, the disambiguation requires an additional $\mathcal{O}(n * k)$ operations, where n is the number of test samples, and k is the average number of nearest neighbors to be processed for each sample. For each test sample a constant time is needed for locating the nearest neighboring sense-tagged samples, if an index is used. The fuzzy variant of the disambiguation requires $\mathcal{O}(n * k * l^2)$ operations, where n and k are as before and l is the number of map locations for each sample.

III. DATA AND PREPROCESSING

A. Data set

In order to compare the method with other systems we used the data set from the English lexical sample task of the SENSEVAL-2 competition [6], which took place in July 2001. The data set consisted of 8611 training samples and 4328 test samples of 73 words in context.

The lexicon used for the sense inventory was WordNet 1.7 and the word instances were mostly from the British National Corpus with some from the Wall Street Journal. Inter-tagger-agreement was 85.5% [15]. There were between 70 and 455 instances per word³ (divided 2:1 between training data and test data) with an average of 121:61 for nouns, 91:46 for adjectives and 122:62 for verbs.

The words in context comprised 29 nouns (art, authority, bar, bum, chair, channel, child, church, circuit, day, detention, dyke, facility, fatigue, feeling, grip, hearth, holiday, lady, material, mouth, nation, nature, post, restraint,

sense, spade, stress, yew), 15 adjectives (blind, colourless, cool, faithful, fine, fit, free, graceful, green, local, natural, oblique, simple, solemn, vital) and 29 verbs (begin, call, carry, collaborate, develop, draw, dress, drift, drive, face, ferret, find, keep, leave, live, match, play, pull, replace, see, serve, strike, train, treat, turn, use, wander, wash, work).

Examples of two noun senses for the word church are the building sense: "The little church has suffered from the rigours of time and town planning, normally a lethal combination, and survived.", and the congregation sense: "On the other hand to disassociate the Church from the Kingdom breaks the nerve-cord of hope and destroys the community of commitment to Christ as Saviour and Lord."

Before disambiguation the words had an average of 10.9 senses per word in the training data (8.7 senses per word in the testing data). Nouns had 8.2 (6.7) senses, adjectives 4.4 (4.1) and verbs 15.6 (12.2) on the average. The high number of senses is partly due to the fact that multi-word expressions are counted as separate senses of single words. There were 105 (67) compound nouns and 155 (106) phrasal verbs, which accounted for 8.5 % (7.5 %) of the noun and 9.7 % (10.1 %) of the verb data, respectively. Typically a compound noun would have only one sense and a phrasal verb would have one or two senses. Without the compounds and phrasal verbs the noun and verb sense ambiguity was 4.6 (4.4) and 10.2 (8.5) senses per word, respectively.

B. Preprocessing of the sense inventory

The sense inventory provided by the SENSEVAL organizers contained senses both for single words as well as for multi-word expressions or collocations such as compound words and phrasal verbs. Spotting multi-word expressions was left to preprocessing with a phrase filter. E.g. the word church was 4-ways ambiguous, but the collocation 'church bell' had only one sense. Without phrase filtering the word church would in this case have been regarded as 5-ways ambiguous. Due to the phrase filter a compound word or phrasal verb would sometimes be falsely recognized, e.g. "The first ladies' mile in ...", where the correct interpretation is the 'first mile' and not the 'first lady'. The falsely recognized multi-word expressions were so few (< 1 %) in the SENSEVAL training data that this was not considered important for the experiment.

After the phrase filter the words had an average of 7.2 senses per word in context in the training data (7.4 senses per word in the testing data). Nouns had 5.5 (5.6) senses, adjectives 6.1 (6.1) and verbs 9.6 (9.8) on the average.

C. Baselines

The expectation value baseline in the test data was 14 % correct sense assignments. This is the baseline, which would be gained by equally distributed random selection of a word sense from a full sense inventory of a word. The majority baseline was 47.6 % correct, which would be achieved by always selecting the majority sense.

The phrase filter changes the baselines, because the full range of senses need no longer be considered. In the prepro-

²Strictly speaking this is true only for the rectangular SOM lattice; for the hexagonal lattice the calculation is slightly more complicated. In any case, the same distance calculation has been used as in defining the map neighborhood function in SOM learning, cf. [14]

³The aim of the organizers was to have $75 + 15n + 6m$ instances for each word, where n is the number of senses and m is the number of multi-word expressions for a word.

cessing we left out the proper noun and unknown senses of the words. We also discarded samples with more than one tag in the training data. This already changed the expectation value and majority baseline to 28 % and 51.3 %, respectively. In addition, the phrase filter assigned the compound and phrasal verb senses only to words identified as compounds and phrasal verbs. This further increased the majority baseline to 52.3 %.

We realized that our phrase filter introduces approximately 4.2 % errors in the test data by not correctly identifying all compounds or phrasal verbs. This is quite a high percentage, but no time was spent fixing the ad hoc phrase filter. The 181 faulty samples were disregarded. The majority baseline for the remaining 4147 samples was 54.6 %.

IV. DISAMBIGUATION EXPERIMENTS

We used the training samples of the SENSEVAL-2 data, which were disambiguated words in context, for calibrating the WEBSOM map. For the parameter selection we used 10-fold cross-validation on the calibration data to find the best-performing parameter combinations. Finally, we disambiguated the test data using the best parameters.

A. Parameter selection

The parameter selection was carried out in two stages. A small initial study was carried out in order to restrict the parameter ranges. Cross-validation was utilized for selecting the final test parameters.

A.1 Extent of the disambiguation context

The task was to assign the proper word sense to a ambiguous word in context. As the context of a word we used alternatively a paragraph, a sentence, a clause, or a 7-word window in which the word occurred. The best performance in the initial study was achieved when the sentence context was used both for calibration and disambiguation.

A.2 Selection of k

The ambiguous word in context was assigned the same sense as the majority of the k nearest sense-disambiguated neighbors. The parameter k was varied between 1 ... 20 and d_{max} between 1 ... 10. Maximum performance was achieved when $k = 3$ and $d_{max} = 2$.

A.3 Fuzzy vs. crisp positions on the map

We also evaluated the method using fuzzy positions on the WEBSOM map. A fuzzy position consisted of l best locations chosen in order of how well they matched the word in context. We evaluated the method using 1, 2, 5, 10 or 15 best locations for each word in context. Fuzzy locations did not improve the overall result when we cross-validated the calibration data, even though it improved the result on the small initial study, so we chose $l = 1$.

B. Test results

The final test results, were obtained using a separate test data set, namely the English lexical task test corpus of the SENSEVAL-2 competition. We evaluated the test data

with the parameter values obtained with the experiments described in Section IV-A on the fine-grained sense distinctions in the SENSEVAL test corpus yielding an overall performance of (2297/4147) 55.39 % correct results with a standard deviation of 0.77 %. This was 65.3 % (std 1.8 %) for adjectives, 59.6 % (std 1.2 %) for nouns and 46.9 % (std 1.2 %) for verbs.

If we include the errors due to the phrase filter, we obtain an overall result of (2297/4328) 53.07 % (std 0.76 %).

C. Importance of the result

The results are statistically significant compared to their majority baselines on the 5 % significance level. Compared to other methods reported in the SENSEVAL-2 competition the described method would have achieved an overall position of 16 among the 21 supervised methods.

This is truly remarkable, considering that the semantic ordering of the map is formed in a fully unsupervised manner, and optimized for text exploration, not for word sense disambiguation. It can be expected that if the application task is taken into account in the construction of the map even better disambiguation performance may be obtained. In this sense, the current work is only a preliminary study of the potential of the method for the word sense disambiguation problem. In contrast, most of the participants in the SENSEVAL-2 competition were complex hybrid systems specifically optimized for that particular application. Our results show that a representation of semantic space can be created independently of the domain and the application.

V. DISCUSSION

As mentioned in Section II, the map was used only for finding neighbors within a d_{max} distance of a test word in context. If no neighbors were within reach, the word in context was disambiguated with a global majority vote. When $k = 3$ and $d_{max} = 2$, the map was used for disambiguating 953 samples. The portion of correct results among the samples, where the map was applied, was 60.4 % (std 1.58 %). It might seem problematic that the map was used only for 953 of the 4157 context data samples, but this is due to the small amount of calibration data. Only one word had more than 400 calibration samples and the average is 112 samples. For words with an average of approx. 200 samples of training data we get an overall performance of 57.3 % correct results and with an average of 250 samples we get 60.6 %.

In [16] Ng reports on a supervised method using a nearest neighbor algorithm directly on the training samples with 100-1300 training samples per ambiguous word. The smallest amount of training data Ng reports is 100 and the results around 250 training samples are similar to ours despite the fact that we use only a bag-of-words with base-forms, whereas Ng uses multiple knowledge sources, such as part-of-speech and verb-object relations as reported in [17]. The comparison needs further investigation, but it strengthens the assumption that the map has induced a

topical and discourse ordering, which is leveraged in the word sense disambiguation task.

The map used in the evaluation was created from 7 million patent abstracts, which means that the domain and the genre of the texts were fairly different from the evaluation data. Some of the words in the SENSEVAL-2 data were not even in the vocabulary when the WEBSOM map was created. Even when the word was in the patent vocabulary and the word sense distinction was of a topical nature, some senses might not be represented at all on the map. An example of this is the building and the congregation sense of the word church. There are no patents pertaining to the congregation sense of church, but there are many relating to the building sense. The result for church was 55.6 % correctly disambiguated (std 6.31 %), which was 4.75 % below the majority baseline for church. This is, however, not statistically significantly lower than the baseline. Only the word *bum* performs statistically significantly below the baseline, whereas the 25 words *art, bar, blind, bum, chair, circuit, detention, dyke, fine, green, grip, natural, sense, vital, call, carry, draw, keep, play, pull, serve, strike, train, turn, wash* all perform statistically significantly above the baseline. It might still be interesting as a future direction to study the effects of a smaller more general purpose document collection for the self-organized map.

The WEBSOM map aims at representing topic and discourse information and disregards morphosyntactic information. This seems to work for the disambiguation contexts and words of the 25 words mentioned above. However, in some cases morphosyntactic information provides important cues for disambiguation. An even better method would perhaps be obtained by combining the strengths of the topically ordered WEBSOM map with some form of morphosyntactic analysis. A simple way of doing this would be to let a linguistic preprocessor select the senses to be disambiguated on the basis of the local context leaving only the topic and discourse related ambiguities for the word-sense disambiguation in semantic space.

As reported e.g. in [1] the WEBSOM map has two different means for locating a map position for a word in context: document search (or *content-addressable search*) and keyword search. In the future our intent is to examine whether the keyword search mode would provide an even better basis for the disambiguation than the currently used document search.

VI. CONCLUSION

We have shown that an independent large-scale representation of semantic space can be utilized in word sense disambiguation. The current study also shows that the distributed knowledge of a WEBSOM map can be used as a representation of semantic space.

Based on the experiments it is apparent that the existing large WEBSOM map distilled from a massive patent abstract collection provided a successful representation of semantic space, even though the map and the original document collection were not initially intended for word sense disambiguation. Constructing a map of a large balanced

document collection might yield further improvements.

What is even more interesting, the current work succeeds in showing experimentally what has earlier been only a hypothesis, namely that the topical ordering obtained using the WEBSOM method is truly a general semantic representation. The word sense disambiguation problem as well as other NLP problems dealing with word semantics can provide us indirect information regarding the cognitive semantic apparatus that humans utilize. The fact that a semantic representation is found that works across problem types and application domains is thus remarkable from a cognitive modeling point of view: it provides strong indirect evidence in support of the hypothesis that similar representational structures may be employed also by the human brain.

ACKNOWLEDGEMENT

We are grateful to the WEBSOM team at the Neural Networks Research Centre of the Helsinki University of Technology for the permission to use the patent abstract map for this experiment.

REFERENCES

- [1] Teuvo Kohonen, Samuel Kaski, Krista Lagus, Jarkko Salojärvi, Vesa Paatero, and Antti Saarela, "Organization of a massive document collection," *IEEE Transactions on Neural Networks, Special Issue on Neural Networks for Data Mining and Knowledge Discovery*, vol. 11, no. 3, pp. 574–585, May 2000.
- [2] Nancy Ide and Jean Veronis, "Introduction to the special issue on word sense disambiguation: The state of the art," *Computational Linguistics*, vol. 24, no. 1, pp. 1–40, 1998.
- [3] Christopher D. Manning and Hinrich Schütze, *Foundations of Statistical Natural Language Processing*, The MIT Press, Cambridge, Massachusetts, 1999.
- [4] David Yarowsky, "Unsupervised word-sense disambiguation rivaling supervised methods," in *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics (ACL '95)*, Cambridge, MA, 1995, pp. 189–196.
- [5] David Martinez and Eneko Agirre, "One sense per collocation and genre/topic variations," in *Proceedings of the Joint SIG-DAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora*, Hong Kong, 2000.
- [6] SENSEVAL-2, "Training and testing corpora," <http://www.cis.upenn.edu/~cotton/senseval/corpora.tgz>, 2001.
- [7] Gerard Escudero, Lluís Màrquez, and German Rigau, "A comparison between supervised learning algorithms for word sense disambiguation," in *Proceedings of CoNLL-2000 and LLL-2000*, Claire Cardie, Walter Daelemans, Claire Nedellec, and Erik Tjong Kim Sang, Eds. 2000, pp. 31–36, Lisbon, Portugal.
- [8] E. Agirre, G. Rigau, J. Atserias, and L. Padró, "Combining supervised and unsupervised lexical knowledge methods for word sense disambiguation," *Computers and Humanities*, vol. 34, no. 1-2, 2000.
- [9] Hinrich Schütze, "Automatic word sense discrimination," *Computational Linguistics*, vol. 24, no. 1, pp. 97–123, 1998.
- [10] Will Lowe, "Towards a theory of semantic space," in *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*, J. D. Moore and K. Stenning, Eds., Mahwah NJ, 2001, pp. 576–581, Lawrence Erlbaum Associates.
- [11] M. Steyvers and J. B. Tenenbaum, "The large-scale structure of semantic networks: statistical analyses and a model of semantic growth," *Cognitive Science*, to appear.
- [12] Thomas Martinetz and Klaus Schulten, "Topology representing networks," *Neural Networks*, vol. 7, no. 3, pp. 507–522, 1994.
- [13] Timo Honkela, Samuel Kaski, Krista Lagus, and Teuvo Kohonen, "Newsgroup exploration with WEBSOM method and browsing interface," Tech. Rep. A32, Helsinki University of Technology, Laboratory of Computer and Information Science, Espoo, Finland, 1996.
- [14] Teuvo Kohonen, *Self-Organizing Maps*, Springer, Berlin, 1995.

- [15] Adam Kilgarriff, “English lexical sample task description,” <http://www.itri.bton.ac.uk/events/senseval/englexsamp.ps>, 2001.
- [16] Hwee Tou Ng, “Getting serious about word sense disambiguation,” in *Proceedings of the ACL SIGLEX Workshop on Tagging Text with Lexical Semantics: Why, What, and How?*, Washington, D.C., USA, 1997, pp. 1–7.
- [17] Hwee Tou Ng and Hian Beng Lee, “Integrating multiple knowledge sources to disambiguate word sense: An exemplar-based approach,” in *Proceedings of the Thirty-Fourth Annual Meeting of the Association for Computational Linguistics*, Arivind Joshi and Martha Palmer, Eds., San Francisco, 1996, pp. 40–47, Morgan Kaufmann Publishers.