Sense Abstractness, Semantic Activation and Word Sense Disambiguation: Implications from Word Association Norms

Oi Yee Kwong

Language Information Sciences Research Centre
City University of Hong Kong
Tat Chee Avenue, Kowloon, Hong Kong

Abstract. Automatic word sense disambiguation (WSD) often draws on a variety of contextual cues, and decides on the most suitable sense by means of some lexical resources or statistical classifiers. While the importance of using multiple types of lexical information is recognised in most systems, not much has been reported on their individual and combined effectiveness in relation to the intrinsic nature of individual words. We attempt to address this cognitive aspect of WSD by examining the psychological evidence regarding the internal lexicon and its compatibility with the information available from computational lexicons. In this study, we compare the responses from a word association task with the lexical associations available from WordNet, to explore the effect of sense abstractness on semantic activation, and thus the implications on the *lexical sensitivity* of WSD. Preliminary results suggest that concrete senses and syntagmatic associations are more readily activated than abstract senses and paradigmatic associations. The results are expected to inform the construction of lexico-semantic resources and WSD strategies.

1 Introduction

Words might have multiple meanings, resulting in *word sense ambiguity*. Getting the right meaning of words in different contexts, otherwise known as *word sense disambiguation* (WSD), is thus an important step in natural language processing (NLP). Automatic WSD often draws on a variety of contextual cues, which are then evaluated against some lexical resources or subject to statistical classifiers, to decide on the most appropriate or most probable sense accordingly.

As Resnik and Yarowsky (1997) remarked, “disambiguation seems highly lexically sensitive, in effect requiring specialised disambiguators for each polysemous word”. Similarly, Ide and Veronis (1998) suggested: “… to date there has been little systematic study of the contribution of different information types for different types of target words. It is likely that this is a next necessary step in WSD work.” In recent years, many research teams all over the world have gained rich experience from the SENSEVAL workshops with their WSD shared tasks. As pointed out by Mihalcea *et al.* (2004), among the 47 participating systems in the SENSEVAL-3 English lexical sample task, “several of the top performance systems are based on combination of

Yee Kwong O. (2007).
Sense Abstractness, Semantic Activation and Word Sense Disambiguation: Implications from Word Association Norms.
In *Proceedings of the 4th International Workshop on Natural Language Processing and Cognitive Science*, pages 169-178
DOI: 10.5220/0002419501690178

multiple classifiers, which shows once again that voting scheme that combine several learning algorithms outperform the accuracy of individual classifiers". However, although the once notorious "knowledge acquisition bottleneck" is partially soothed by statistical methods, the advancement in WSD is rarely accompanied by any extensive account on the cognitive aspects of the lexical sensitivity of the task. Hence the suggestion by Ide and Veronis in what might be considered a somewhat dated source now is nevertheless still valid in a certain sense.

To better understand the contribution of different information types for different types of target words, it is thus important to look at WSD in relation to the very intrinsic nature of the individual words to be disambiguated (or *target words*), in addition to an optimal combination of classifiers alone. We use the concept *Information Susceptibility* (Kwong, 2005) to refer to the relationship between the intrinsic features of a target word and its senses, and the effectiveness of various lexical information to characterise them. While the intrinsic nature of a word and its senses could comprise many factors such as frequency, abstractness, sense relatedness and parts-of-speech (POS), in the current study we focus on the *abstractness* / *concreteness* of individual senses, and analyse the way it corresponds to the responses elicited from word association tests. Since word association norms are generally assumed to reveal the organisation of our mental lexicon, they serve as a bridge between the internal mechanism and the external modelling.

We will start with a discussion in Section 2 on the cognitive aspects of WSD from three perspectives: introspection, psychological evidence, and computational modelling; and how they interact. In Section 3, we present a preliminary study to explore the effect of sense abstractness on semantic activation. In particular, we compare the responses from a word association task with the lexical associations available from WordNet, a widely used computational lexicon. Results and their implications on the lexical sensitivity of WSD will be discussed in Section 4, with future directions, followed by a conclusion in Section 5.

2 Cognitive Aspects of WSD

In this section, we will discuss the need for multiple sources of knowledge for WSD and the evidence of the lexical sensitivity of the task from various perspectives.

2.1 An Introspective Account

Being a common and psychologically valid phenomenon of natural languages, word sense ambiguity (or *polysemy*) penetrates our daily language use. Despite the apparent non-discreteness of "sense" as Kilgarriff (1992) argued, human beings used to rely on the predetermined senses in existing lexical resources, especially dictionaries, as a tool for construing senses. We do not seem to have much difficulty, under normal circumstances, to access the intended interpretation of a polysemous word in a given context. For instance, if one says "we will have three *courses* for dinner", it will be unlikely for any hearer, not even vegetarians, to mistake it as eating up golf courses or the grass on them.

Introspection alone might suggest that when processing the above sentence, we could have been following a path like "dinner is a kind of meal, and interpreting *course* as a part of a meal amongst its other possible meanings might be most appropriate". On the other hand, to get the meaning of *courses* in "he enrolled in three *courses*", instead of using paradigmatic relations such as the IS-A and PART-OF relations employed in the previous example, we might rely more on the syntagmatic relations between "enroll" and "course" for disambiguation.

More generally, human beings appear to use a range of cognitive strategies to make sense of concepts of different abstractness. For instance, one can understand a "mirror" as the instrument which reflects the image of oneself; but the simple hypernymy relation might not be as useful for understanding an "injury". It does not refer to some concrete object, although it could often be visualised (such as a bleeding wound or a broken leg). The more abstract a word, the less obvious is its external reference, thus one could imagine how difficult it is to describe what "loss" is.

A corollary from this phenomenon is that the different strategies in making sense of words, including the type and strength of various kinds of semantic association, should be realised in NLP systems, especially in knowledge demanding subtasks like WSD. Hence Quillian's (1968) network model of semantic memory, in which the association amongst concepts can be of very different nature, has inspired and influenced not only the approaches in WSD but also the many computational lexicons created for use in the task, as further discussed below.

2.2 Psychological Evidence

Quillian (1968) proposed a computational model of human memory for storing the "meanings" of words, which remains influential in our conception of the internal lexicon as well as in the construction of computational semantic lexicons. Apparently, our memory stores not isolated but connected information. So a model of semantic memory should have, in Quillian's words, "the ability to use information input in one frame of reference to answer questions in another". Thus his model has a network structure with a mass of nodes interconnected by associative links of different kinds. The model allows two word concepts to be compared and contrasted via the links between them. The association between concepts was found by a method generally known as "spreading activation with marker passing", which also underlies many later WSD programs.

The psychological validity of such a network model is evident from subsequent studies on lexical priming and lexical access. Priming studies (e.g. Collins and Loftus, 1975), suggest that the processing of a concept (in terms of the response time for lexical decision tasks) would be faster if primed by a semantically related concept. Lexical access studies (e.g. Swinney, 1979), on the other hand, propose several hypotheses regarding the processing of multiple meanings in the case of lexical ambiguity. There are no unanimous results, but it seems that multiple meanings are activated at least briefly, and the influence of prior context might interact with the nature of the individual senses, such as dominance or familiarity in terms of frequency.

As far as the nature of the target words is concerned, the relatedness among multiple senses could be another factor. For instance, in the lexical access literature, there is a general finding that visual lexical decisions are faster for words that are semantically ambiguous. Rodd *et al.* (2002) challenged the conventional assumption that this phenomenon of "ambiguity advantage" is a result of ambiguity between multiple, unrelated meanings, instead of multiple related word senses, with several sets of test words more rigorously controlled for the sense relatedness therein.

Thus from the psychological perspective, we see that the access of multiple meanings and the processing of lexical ambiguity are likely to be influenced by the nature of individual target words. It is worthwhile to investigate how such lexical sensitivity could be modelled in automatic WSD, and whether such modelling could substantially benefit the latter.

2.3 Computational Modelling

Automatic resolution of word sense ambiguities has primarily depended on contextual features, which are evaluated against some lexical knowledge sources, or subject to statistical classifiers based on various machine learning algorithms.

In early studies, lexico-semantic knowledge for WSD was often hand-coded for particular systems, e.g. semantic networks (Hirst, 1987), and core and dynamic lexicons (McRoy, 1992). These are serious and rich semantic resources, but at the expense of time, labour and scalability. With the availability of machine-readable dictionaries, thesauri, and large corpora, researchers have explored various ways to (semi-) automatically acquire semantic information from them.

WordNet (Fellbaum, 1998) is probably the first broad coverage general computational lexical database. It defines word senses via synonymy, linked by relational pointers (e.g. hypernym, antonym, etc.), forming semantic nets. It is, however, known for the lack of syntagmatic relations, and researchers started to address this gap with various means to enrich the lexicon with topic associations and other broader semantic relations to enhance word access (e.g. Ferret and Zock, 2006).

Following the upsurge of corpus-based and empirical methods, statistical approaches become the common practice in automatic WSD. Multiple knowledge sources are modelled computationally as a variety of features from topical and local contexts. The prevalence of machine learning approaches in WSD is evident from the recent SENSEVAL workshops (Mihalcea *et al*., 2004).

Thus knowledge-based methods for WSD address the need for multiple types of lexical knowledge by using semantic networks containing different kinds of semantic relations (e.g. IS-A, PART-OF, thematic relatedness, etc.), and statistical methods address the issue by getting an optimal combination of the various knowledge sources for individual target words (e.g. Mihalcea, 2002). However, it is interesting to note that there is somehow no comprehensive qualitative and objective account of the relation between the disambiguation results and the nature of individual target words underlying the apparent lexical sensitivity of the task.

2.4 When They Meet

To say that different information types contribute variably to different target words is essentially presupposing that different types of lexical information vary in their effectiveness to characterise a sense and distinguish it from other senses of the same word. Thus it is not enough to just conceptualise senses by a certain dimension (e.g. a certain semantic relation) across the board. The very intrinsic nature of a given word/sense and its relation with different semantic dimensions must also be thoroughly examined.

Leacock *et al*. (1998), for example, observed that “the benefits of adding topical to local context alone depend on syntactic category as well as on the characteristics of the individual word”. Such “characteristics” are equivalent to the *intrinsic properties* of the target words in our discussion, which might include abstractness, frequency, sense relatedness, POS, amongst others, and as we propose, are critical for understanding the lexical sensitivity of the WSD task.

Information Susceptibility (Kwong, 2005) thus refers to the relation between the intrinsic properties of a word and the effectiveness of various types of lexico-semantic knowledge to characterise it. Such information is absent from existing lexical resources. Based on the performance of a spectrum of semantic relations to disambiguate a set of target words, it was observed that senses involving more abstract thinking tend to be disambiguated only with broader semantic relations. This observation also coincides with findings from human word association tests. For instance, in the Birkbeck word association norms (Moss and Older, 1996), “loss” triggers associations like “death” and “grief”, which cannot be related via a simple IS-A relation, in contrast to responses like “magic” triggered by “trick” which are simply synonymous. Hence, from the cognitive perspective, the knowledge on the information susceptibility of individual target words is important for fine-tuning WSD systems and informing the optimal combination of disambiguation cues. To provide this knowledge in existing lexical resources, we need to examine the nature of target words (in terms of frequency, abstractness, sense relatedness, POS, etc.) in the context of lexical access and WSD.

Hence, in the current study, we focus on one aspect of the intrinsic nature of words, namely *sense abstractness*, and explore how it varies with the kind of lexical association in our mental lexicon and how it might affect the effectiveness of various kinds of semantic knowledge in disambiguation. The study is based on data from word association norms, and we compare the responses gathered in Hirsh and Tree’s (2001) study with the lexical association available from the widely used semantic lexicon WordNet. Since word association is a commonly used method to probe the organisation and structure of the internal lexicon, and computational lexicons and ontologies are assumed to model human conceptualisation, the comparison is expected to allow us to better understand the human semantic repertoire and thus the computational information demand for individual words in the lexically sensitive disambiguation process.

3 A Preliminary Study

In this section, we present our preliminary study on the effect of sense abstractness on semantic activation, by comparing the responses from a word association task with the lexical associations available from WordNet. The word association responses are assumed to be reflective of the organisation of the internal lexicon, and WordNet information is used for operationalising the type and strength of semantic association links. The objectives are two-fold: (1) to investigate the effect of one aspect of the intrinsic nature of words on semantic activation, and (2) to study the implications of such target-dependent semantic association patterns, if any, on the lexical sensitivity of WSD.

3.1 Materials

We used the 90 stimulus words in Hirsh and Tree's (2001) word association test as our target words in this study, and focused on the top five responses elicited from the young cohort. (Hirsh and Tree analysed the difference between the responses from young adults and those from older adults.)

As mentioned, WordNet organises word senses in the form of synsets (i.e. sets of synonyms) with relational pointers linking among different synsets to form some sort of a semantic hierarchy. The synsets are also organised under 45 lexicographer files based on syntactic category and logical groupings. WordNet was created for psycholinguistic studies of the mental lexicon to start with but turned out to be an electronic resource widely used by computational linguists. Thus, in this study, we used WordNet 2.1: (1) as a dictionary to provide information on the number of senses for a word, (2) as a computational model of the internal lexicon in the form of a semantic network, despite its known bias toward paradigmatic relationship in general, and (3) as a means to distinguish between concrete and abstract concepts.

3.2 Method

The 90 target words were first checked against WordNet 2.1 for the number of senses they have, and each sense against the lexicographer files to which they belong, to determine whether they correspond to concrete or abstract concepts.

The top five responses from the young cohort were taken and compared to two groups of word associations obtained from WordNet. The first, which we will call WNAsso1 below, consists of all *words* in the synsets (words composing the synsets only, excluding glosses and examples) directly related to the synset(s) to which the target word belongs. These directly related synsets include antonyms, hypernyms, hyponyms, holonyms, meronyms, and coordinate (or sister) terms. The second, which we will call WNAsso2 below, consists of the words in the glosses and examples in these related synsets. Thus the first could be taken as the cluster of words corresponding to mostly paradigmatic relations with the target words, and the second more likely to be broader semantic relations and associations, including some syntagmatic relations. In the current study, we only looked at the noun hierarchy, and ignored senses of the target words under other POS.[1]

[1] It is, however, possible that Hirsh and Tree's respondents might not have always responded to a stimulus word as a noun, given that they were not specifically instructed to do so.

4 Results and Discussion

Among the 90 target words, 16 were monosemous and 73 were polysemous according to the WordNet noun database.[2] The polysemous words have 2 to 15 senses, with an average of 4.7 senses. There was one target word (i.e. "sly") which was found only in the adjective database in WordNet, and was ignored in subsequent analysis.

There are 45 lexicographer files under which the synsets are organised based on syntactic category and logical groupings, 26 of which are relevant to noun senses. We identified 7 concrete classes and 19 abstract classes. The concrete classes include animal, artifact, body, food, object, person, and plant. The rest are the abstract classes, including attribute, cognition, feeling, motive, and so on. This dichotomous distinction may have certain limitation, which will be further discussed below.

Thus for the remaining 89 target words, all (16) monosemous words and 20 (out of 73) polysemous words only have concrete or tangible senses, 3 polysemous words only have abstract senses (they are "bunch", "traffic" and "wedding"), and 50 polysemous words have both tangible and abstract senses. This results in altogether 222 tangible senses and 136 abstract senses. The fact that more tangible senses are observed is expected because Hirsh and Tree (2001) had indicated in their study that their stimuli were "mostly names of concrete or picturable objects or likely to elicit the name of a concrete object". However, they did not mention with respect to which sense the "picturability" was determined in the event of polysemy.

Table 1 shows the results for comparing the association responses with WNAsso1 and WNAsso2. The figures show the number of target words found under the various overlapping scenarios. The overlapping could correspond to one or more senses of a given target word. Thus WNAsso1 was assumed to contain mostly paradigmatically related words and WNAsso2 broader associations including some syntagmatically related words. It can be seen that for all sense types, the "WNAsso2" and "Both" columns make up the majority, and only three cases overlap with purely paradigmatic responses.[3] This is in consensus with Hirsh and Tree's analysis, where they observed more syntagmatic responses. There are, however, a few exceptional cases which have none of their responses overlapping with any of our WordNet data. Some preliminary qualitative analysis of the results is discussed below, regarding the relationship between the numerical figures and the abstract/concrete nature of the words.

Table 1. Results for comparing the association responses with WordNet data.

Word Type	Sense Abstractness	Overlapping			
		WNAsso1 Only	WNAsso2 Only	Both	None
Monosemous	All Tangible	1	6	8	1
Polysemous	All Tangible	2	5	11	2
	All Abstract	0	1	1	1
	Both T & A	0	22	27	1

[2] Hirsh and Tree (2001) claimed to have 41 unambiguous nouns. This was more than what we found with WordNet senses, which might have more fine-grained senses.

[3] Note that this observation does not preclude any syntagmatic responses for the three cases in the word association test, which might not be found in our limited syntagmatic associations obtained from WordNet.

The lexical access literature suggests that multiple senses might be at least briefly activated in the case of polysemy, but has not systematically explored the sense abstractness factor. For 21 out of the 50 polysemous words with both tangible and abstract senses, the association responses overlap with WNAsso1 or WNAsso2 corresponding to one or more of their tangible senses *only*. For instance, the stimulus word "zip" has four noun senses in WordNet, including "zero", "postcode", and "vigour" which are abstract, and "zipper" which is tangible. The top five responses are "trouser(s)", "fly", "button(s)", "jacket" and "clothes". All except "fly" were found among the WordNet associations corresponding to the "zipper" sense. This is in contrast to 6 (out of 50) with responses overlapping with WordNet associations corresponding to their abstract senses only. For example, the stimulus word "safety" has two tangible senses and four abstract senses in WordNet. Only the response "security" overlaps with WNAsso1 and WNAsso2 for one of the abstract senses referring to "a state of being certain that adverse effects will not be caused". This observation suggests that in the case of polysemy with both tangible and abstract senses, the tangible concepts seem to be relatively more accessible from the internal lexicon, assuming word association responses reflect the closest and strongest associations in the internal lexicon.

Notwithstanding the above observation, the preference for tangible senses might also be a result of frequency or familiarity. However, the frequency effect is not obvious from the current study. While WordNet senses are ordered by frequency, there is no significant pattern to show that the responses are necessarily related to the first few senses. There are several cases where the overlapping corresponds to the top one or two senses of a word, but no conclusive remarks could be made at this stage, and further investigation with better control on the sense frequency would be required.

As mentioned earlier and evident from Table 1, syntagmatic associations appear to be more prevalent than paradigmatic ones. This is not surprising given the much broader possibilities with syntagmatic associations. Nevertheless, about 38% of all target words have responses overlapping with WNAsso2 only. So what underlies the absolute dominance of syntagmatic associations in these cases? Could it be related to the specificity and concreteness of the senses? However, looking at the six monosemous words under this category, they are nevertheless located at a position in the WordNet hierarchy as deeply branched as the other monosemous target words, and thus they appear similarly specific. At the same time, the apparent inferiority of paradigmatic responses might be an artifact of the WordNet classification itself. For instance, the hypernym of "ankle" is "gliding joint", and that for "kennel" is "outbuilding", which might be too specialised for daily usages and conception. Hence, even though the top response for "ankle" is "foot", they are not straightforwardly related in the WordNet database. The concreteness hypothesis is not supported either, given that all the monosemous target words are tangible concepts, there is still a substantial portion of them dominated by syntagmatic responses. One limitation, however, is that our dichotomous distinction between concreteness and abstractness might be too coarse, whereas abstractness / concreteness could be a continuum. Another drawback of using the lexicographer files for the distinction is that even seemingly tangible classes like "animal" could also be abstract with words like "Animalia". We definitely need to address this issue in future studies.

Thus it seems that although paradigmatic relations like IS-A or PART-OF are an important part of our semantic knowledge, spreading activation seems to favour broader associations. What makes the syntagmatically related words stronger links would be our focus in future study. In particular, we plan to extend our comparison with corpus-based data sources, for more syntagmatic relations and general associations. We will also refine our definition of sense abstractness. In addition to the concrete/abstract distinction, other factors like frequency, relatedness among senses, POS and possibly others might all contribute to the intrinsic nature of words, and target words need better control on these dimensions in future work. Moreover, given our preliminary findings on sense abstractness and semantic activation, one important future direction is to further examine disambiguation performance on concrete and abstract senses and to investigate their respective information demand for WSD.

Thus our current preliminary study has at least the following implications on the lexical sensitivity of WSD and the classification of senses in computational lexicon for WSD: (1) Tangible concepts seem to be more easily activated in the internal lexicon, and even in the case of polysemy, tangible senses appear to be more accessible than abstract senses, although frequency and familiarity might also play a role. (2) While paradigmatic associations form an important part of our semantic knowledge, the observed dominance of syntagmatic associations might inform the computational modelling of the internal lexicon, such that different weights might be attached to different kinds of associations for words with different nature. To this end, it is worth to investigate the feasibility of enriching existing lexical resources like WordNet as well as the possibility of an alternative classification of word senses based on the intrinsic nature of words, in addition to conventional conceptual classifications in existing lexical databases.

5 Conclusion

In this study we have analysed word association responses with respect to the lexical associations obtained from a widely used computational lexicon, namely WordNet. In particular, we explored the effect of sense abstractness on semantic activation, and thus the implications on the lexical sensitivity of automatic WSD. Preliminary results suggest that tangible senses are more readily accessed and syntagmatically related senses are apparently more strongly associated. The results do not only reinforce the significance of the intrinsic nature of individual target words in WSD, but also inform the computational modelling of the internal lexicon and semantic knowledge for the task.

Acknowledgements

The work described in this paper was fully supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. CityU 1508/06H).

References

1. Collins, A.M. and Loftus, E.F. (1975) A spreading-activation theory of semantic processing. *Psychological Review, 82(6)*:407-428.
2. Fellbaum, C. (1998) *WordNet: An Electronic Lexical Database*. MIT Press.
3. Ferret, O. and Zock, M. (2006) Enhancing electronic dictionaries with an index based on associations. In *Proceedings of COLING-ACL 2006*, Sydney, Australia, pp.281-288.
4. Hirsh, K.W. and Tree, J.J. (2001) Word association norms for two cohorts of British adults. *Journal of Neurolinguistics, 14*:1-44.
5. Hirst, G. (1987) *Semantic interpretation and the resolution of ambiguity*. Cambridge, UK: Cambridge University Press.
6. Ide, N. and Veronis, J. (1998) Introduction to the special issue on word sense disambiguation: The state of the art. *Computational Linguistics, 24(1)*:1-40.
7. Kilgarriff, A. (1992) *Polysemy*. Ph.D. thesis, University of Sussex.
8. Kwong, O.Y. (2005) Word Sense Classification Based on Information Susceptibility. In A. Lenci, S. Montemagni and V. Pirrelli (Eds.), *Acquisition and Representation of Word Meaning*. Linguistica Computazionale, pp.89-115.
9. Leacock, C., Miller, G.A. and Chodorow, M. (1998) Using Corpus Statistics and WordNet Relations for Sense Identification. *Computational Linguistics, 24(1)*:147-166.
10. McRoy, S.W. (1992) Using multiple knowledge sources for word sense disambiguation. *Computational Linguistics, 18(1)*:1-30.
11. Mihalcea, R.F. (2002) Word sense disambiguation with pattern learning and automatic feature selection. *Natural Language Engineering, 8(4)*:343-358.
12. Mihalcea, R., Chklovski, T. and Kilgarriff, A. (2004) The SENSEVAL-3 English Lexical Sample Task. In *Proceedings of SENSEVAL-3*, Barcelona, Spain.
13. Moss, H. and Older, L. (1996) *Birkbeck Word Association Norms*. Hove, U.K.: Psychology Press.
14. Quillian, M.R. (1968) Semantic memory. In M. Minsky (Ed.), *Semantic Information Processing*. Cambridge, MA: MIT Press.
15. Resnik, P. and Yarowsky, D. (1997) A perspective on word sense disambiguation methods and their evaluation. In *Proceedings of SIGLEX'97 Workshop: Tagging Text with Lexical Semantics: Why, What, and How?*, Washington, D.C., pp.79-86.
16. Rodd, J., Gaskell, G. and Marslen-Wilson, W. (2002) Making Sense of Semantic Ambiguity: Semantic Competition in Lexical Access. *Journal of Memory and Language, 46*:245-266.
17. Swinney, D.A. (1979) Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior, 18*:645-659.