

Contextual Representation of Abstract Nouns: A Neural Network Approach

Katja Wiemer-Hastings (kwiemer@cc.memphis.edu)

Arthur C. Graesser (a-graesser@memphis.edu)

Department of Psychology, The University of Memphis; CAMPUS BOX 526400
Memphis, TN 38152-6400 USA

Abstract

This paper explores the use of an artificial neural network to investigate the mental representation of abstract noun meanings. Unlike concrete nouns, abstract nouns refer to entities that cannot be pointed to. Cues to their meaning must therefore be in their context of use. It has frequently been shown that the meaning of a word varies with its contexts of use. It is more difficult, however, to identify which elements of context are relevant to a word's meaning. The present study demonstrates that a connectionist network can be used to examine this problem. A feedforward network learned to distinguish among seven abstract nouns based on characteristics of their verbal contexts in a corpus of randomly selected sentences. The results suggest that, for our sample, the contextual representation of abstract nouns is in principle sufficient to identify and distinguish abstract nouns, and thus meets the functional requirements of concept representation.

Introduction

In the recent two decades, concept theories have increasingly adopted the view of flexible, context-dependent concept representations (e.g., Anderson, 1990; Barsalou, 1982; Kintsch, 1998). These theories can better account for the empirical findings which suggest that concepts are sensitive to different contexts of use.

The necessity of a context-sensitive representation becomes especially evident when considering abstract nouns such as *goal*, and *idea*. Abstract nouns refer to entities that are not perceivable and cannot be pointed to, as in the case of concrete nouns, like *car*. Therefore, the standard approach of decomposing concepts into features cannot be applied to abstract concepts. Instead, in order to understand abstract nouns, the contexts in which they are used are important (Quine, 1960). In order to understand what *idea* means, for example, it may be useful to know that someone successfully solves a problem after having an idea.

In this study, we explore verbal context. We claim that abstract noun meanings are strongly related to their verbal contexts. The strong form of our hypothesis is that the meaning of abstract nouns may be *determined*, not just influenced, by the context of use. It is motivated by the effect of context on word meanings on the one hand, and by the difficulty to perceive the references of abstract nouns on the other hand. Abstract nouns are acquired relatively late

in language acquisition, and it is therefore possible (but does not necessarily follow) that children use a variety of verbal context cues to infer the meanings of abstract nouns. The claim that context is the source for understanding abstract nouns is intuitively compelling. Still, the claim needs to be investigated to examine the relation of context and word meaning. It seems necessary to examine the role of a broad set of context features.

The fact that context plays a major role for a word's meaning has been widely acknowledged. Evidence that context contains information related to a word's meaning is apparent in studies using the cloze method (for example, Hamberger, Friedman & Rosen, 1996). Subjects are shown a sentence in which a word has been deleted. In its place, there is a blank. The subjects are asked to fill in a word that fits the context best. The information that contexts provide about a missing word's meaning varies. Some contexts provide enough information for the subject to identify the correct missing word, whereas in other cases a wide range of words fits the context.

Lexicographers appreciate the importance of context. In their attempts to define a word's definition, they draw information from example sentences in which the defined word is used. This would be an unproductive task if these contexts did not contain information that constrains the range of possible meanings of the word to be defined.

There is some evidence that humans often derive a word's meaning from its context. Sternberg and Powell (1983) have demonstrated that the meanings of unfamiliar words can be derived from context. McKeown (1985) and many others have argued that word meaning can be learned from context. It is assumed that, after exposure to different contexts, the language learner ideally starts to decontextualize the word meaning.

Our strong hypothesis is further supported by research that points out the relevance of context specifically in the processing of abstract words. Schwanenflugel and Shoben (1983) demonstrated that abstract nouns are processed more easily when presented in a context that activates information relevant to the noun in memory. They demonstrate that the memory advantage of visualizable (that is, concrete) material disappears if both abstract and concrete words are embedded in context. They explain this with the *context availability model*, according to which abstract noun

representations are only weakly connected to associated context information in memory.

Context Information

Miller and Charles (1991) have argued that similar words occur in similar contexts. They propose four types of context information that is stored with concepts in the mental lexicon: collocation, semantic context, syntactic context, and pragmatic context. The subjects' ability to fill in the correct word in studies employing the cloze procedure may be explained with collocation information, that is, associations of words that frequently co-occur in sentences.

Two recent models of semantic representation have emphasized the role of word co-occurrence in context: HAL (Hyperspace Analogue to Language, Lund & Burgess, 1996) and LSA (Latent Semantic Analysis, Landauer & Dumais, 1997). In these systems, context is not represented by features, as in the network discussed in the present paper, but by the co-occurrence of words in contexts. Both HAL and LSA impressively demonstrate that word co-occurrence information (or collocation information) of a particular word can go a far way in determining the word's location in semantic space. We argue in this paper that the meaning of abstract nouns is not determined by co-occurring words, but instead by semantic and syntactic context information.

Abstract Concepts in HAL

Burgess and Lund (1997) have demonstrated that a set of abstract concepts can be classified in semantic space on the basis of co-occurrences. According to their model, abstract concepts can be classified in the same way as other types of words (concrete nouns and others). This finding supports the argument that context, in this particular case, word co-occurrence, is a powerful semantic information source. However, the result and the model do not explain why there are differences in the cognitive processing of abstract and concrete concepts, as they have been demonstrated in imagery and comprehension time studies. In this respect, the HAL model does not present a convincing model of abstract concept representation.

The concepts in the study by Burgess and Lund were not exclusively abstract nouns. The emotional connotation words were clearly separated from the legal terms in a multidimensional scaling analysis, but it is not clear whether this is due to their emotional connotation or due to the ambiguity of their word class. Burgess and Lund demonstrate in the same study that HAL can separate syntactic classes of words. We think that this could explain their results for abstract concepts. For example, *happy* is an adjective, and *love* and *hate* are often used as verbs. Francis and Kucera (1982) list a verb frequency of 145 and a noun frequency of 179 for *love*. For *hate*, the verb frequency is 66 and the noun frequency is 10. These frequencies match the semantic distances as presented by Burgess and Lund (p. 62): *Love* is closer to the cluster of law terms (mainly consisting of nouns), and *hate* is at the distant corner. Of

course, a better estimate of frequencies would be based on the text corpus that HAL used for the study.

The possibility that HAL has separated word classes instead of categories of abstract concepts is further supported by the fact that *joy* and *sorrow*, which are clearly emotional terms, are within the cluster of legal terms, whereas the term *judge*, which in Francis and Kucera has a verb frequency of 42 and a noun frequency of 81, is oriented towards the *love* and *hate* cluster. *Murder* is in the 'noun cluster', and according to Francis and Kucera, the verb frequency for *murder* is rather low (16), whereas the noun frequency is 83.

It is difficult to judge, based on this rather small sample, whether the results are really confounded by word class frequencies. The terms with emotional connotation are clearly located at one side of the MDS outcome.

Semantic and Syntactic Context Information

We believe that the contextual representation related to the concept representation of abstract nouns goes beyond word co-occurrence. Of course, co-occurrence reflects meaningful relations among words. However, it is difficult to see how meaning can emerge from merely co-occurrence. Determining word meanings by words in the context can only work if the determining words have meaning independent of other words. Otherwise, the determination would be circular. In this study, we analyze the co-occurring words with respect to their features.

Semantic and syntactic context information can be formulated as abstract features. An example for semantic features is provided by Plaut and Shallice (1991), who have represented abstract and concrete nouns in a neural network by different numbers of features. However, their features did not represent context information.

One advantage of representing context with semantic and syntactic features is that they can describe all verbal contexts. The features can be similar in contexts of a particular word in which the collocation information differs.

We think also that it is possible on this more abstract representational level to depict differences between the representation of abstract and concrete nouns, because the latter, by virtue of their references, do not depend on contextual information.

This analysis is an extension of an earlier study by Wiemer-Hastings (in press). A neural network was trained to identify six abstract nouns. The input was a vector of 53 semantic and syntactic features of the contexts that the abstract nouns occur in. The present study replicates this for seven target nouns. However, the estimates for the relative relevance of the context features looked dramatically different for the present network (for a description of the features see Wiemer-Hastings, in press). Whereas the previous network's performance seemed to mostly rely on verb and world domain information, the present network provided higher estimates for syntax than previously. Verb and knowledge domain got rather low estimates. The

present network was also tested with a larger set of synonyms.

The Network Model

It is assumed that context constitutes important aspects of word meaning. One implication of this assumption is that the contexts of a given word must have features in common. In particular, the contexts of similar words (synonyms) should have features in common (Miller and Charles, 1991), and the contexts of different words must be sufficiently distinct from each other.

The network model that we present was not designed as a model of human word learning. However, we argue that our results, as well as the findings of Burgess and Lund (1997), show that context information is *in principle* sufficient to distinguish a set of abstract concepts. That is, they support the view that word meanings are learned from context.

Context Representation

In the network, context was represented by a vector of 53 semantic and syntactic context features. The features were manually extracted from natural sentences in which the target nouns occur. The features were assigned binary values (1, 0) depending on the context. Context information covered the verb and adjective that are directly related to the abstract nouns, ontological status information about the target as far as evident from the context, the case roles of the nouns, and the information expressed by the whole sentence (including domains of world knowledge). We included information that could be derived from the context only. Thus, our knowledge of the meaning of the target noun did not enter into the context analysis. Features that could not be determined from the context of a given sentence were set to zero. The features are described in more detail in Wiemer-Hastings (in press).

The context feature vectors were the network input. The network was trained to identify the target noun that fits the sentence represented by a particular input vector. The performance of the network was used to assess

- (1) whether the information in the contexts is sufficiently distinctive to classify abstract nouns, and
- (2) whether the set of features cover the relevant information.


Training and Test Corpus

Seven abstract noun targets were randomly selected from a corpus: *Attention*, *concept*, *consultation*, *goal*, *idea*, *impression*, and *wisdom*. For each of the seven target nouns, 125 isolated sentences were selected from the NexisLexis online database. We discarded semantically depleted sentences, that is, very general sentences that do not contain any information about the word's meaning. 875 training cases with 53 features were constructed from these sentences.

The Network Architecture

Figure 1 presents the architecture of the network. It has a 53-21-7 fully connected feedforward architecture, with 53 input units, 21 hidden units, and 7 output units. The network's program forced the network to select only one output for each input vector. Thus, multiple classification was not possible.

Figure 1: The network architecture.



Training

The network was trained with backpropagation, using NevProp 3 (Goodman, 1996), in two phases. In the first phase, only 775 sentences of the training set were used for training. The other 100 cases were used to estimate how well the network generalizes during training. After this phase, the program computed the average square error for which the generalization was best. This error was stored and used as target error for training phase 2 to prevent overtraining. In phase 2, the network was trained with the full training set. Training stopped automatically when the target error was reached.

Results

The network achieved the target error of .085 after 120 training epochs. Generalization was tested with the feature vectors of 71 new sentence that the network had not been trained with. This test is comparable to a cloze experiment, where a sentence frame is presented and humans have to identify a missing word. Furthermore, the network was tested with context features from sentences that contained synonyms or words related to the target nouns. This test is motivated by the assumption that similar words occur in similar contexts (see, e.g., Miller & Charles, 1991). It should

be possible for the network, given the context features of a synonym's context as input, to identify the target noun that is most similar to the synonym. This test is constrained by the existence of words that are really similar in meaning. The test can show whether the context features used in this study really cover the aspects of context which are relevant to the *meanings* of the target nouns.

Generalization to Sentences with Target Words

The network could classify most of the test sentences (70%) correctly. Performance at chance level would be a classification rate of 14%. The Average Square Error (ASE) for the corpus of test sentences was .107. The generalization rate is very good. The high correct classification rate indicates that the features representing the sentence contexts are, in most cases, sufficient to distinguish among the target nouns. This suggests that word meanings can indeed be distinguished and identified on the basis of their contexts. Table 1 presents the correct classification rates segregated by the target nouns.

Table 1: Number of test sentences per target noun that are classified correctly by the network.

| Target noun | Classification rate absolute and in % |
|--------------|--|
| Attention | 9/11 (81.8%) |
| Concept | 7/10 (70%) |
| Consultation | 8/10 (80%) |
| Goal | 7/10 (70%) |
| Idea | 4/10 (40%) |
| Impression | 8/10 (80%) |
| Wisdom | 7/10 (70%) |

For none of the target nouns, the network could classify all sentences correctly. This is not too surprising, considering the fact that the amount of relevant context information varies across sentences.

The network had difficulties classifying the sentence frames for *idea*. Only 40% of the test sentences were classified correctly. This phenomenon deserves some

attention. In the next section we present results that demonstrate that the network can classify a good percentage of sentence frames that are the contexts of synonyms of the targets. But even some of the targets are closely related in meaning. In particular, Roget's lists *idea* as the synonym for both *concept* and *impression*. The network misclassifies the sentence frames for *idea* as *concept* once, and twice as *impression*. As the goal is for the network to identify abstract nouns' meanings, it should make these errors for words that are synonymous to alternative targets.

Table 2 describes the specific misclassifications of the network. The leftmost column presents the correct target nouns fitting the test sentence contexts; the upper row presents the target nouns that the network selects. For example, contexts for *attention* were misclassified as *wisdom* twice.

Test with Synonyms of the Targets

The network was tested on contexts of *synonyms* of the target nouns. Synonyms were selected from Roget's Thesaurus and WordNet (Miller, 1990). For most target nouns (e.g. *attention*, *concept*, *wisdom*), there are no satisfying synonyms but only related terms. Furthermore, some target nouns have synonyms in Roget's thesaurus in common. We have therefore used some of the synonyms as test words for more than one target noun.

We selected *thought*, *plan*, and *visualization* for *concept*, *counseling* and *advice* for *consultation*; *intention*, *motive*, *plan* and *objective* for *goal*; *thought* and *plan* for *idea*; *belief* for *impression*; and *knowledge*, *experience*, *maxim* and *understanding* for *wisdom*. For the target noun *attention*, no acceptable synonym was found, therefore it was not tested in this analysis.

Table 3 gives an overview of the results for synonyms. For each synonym, Table 3 lists all the target nouns that the network selected as the output, along with the number of times the target was selected for the synonym. For some synonyms, for example, *objective*, the results are rather promising.

Table 2: Erroneous classifications made by the network on the test sentence corpus.

| Target | Attention | Concept | Consultation | Goal | Idea | Impression | Wisdom |
|--------------|-----------|---------|--------------|------|------|------------|--------|
| Attention | (9) | | | | | | 2 |
| Concept | | (7) | | | 2 | | 1 |
| Consultation | | 1 | (8) | | | | 1 |
| Goal | | | | (7) | 2 | 1 | |
| Idea | | 1 | 1 | | (4) | 2 | 2 |
| Impression | | | | 1 | | (8) | 1 |
| Wisdom | 1 | 1 | | 1 | | | (7) |

For other synonyms, either a large variety of targets was selected or a target was consistently selected which was not

predicted by our sources. Even though some of the outputs of the network are not consistent with our predictions, the

results are encouraging.

Table 3: Synonyms for the target nouns and the target nouns they were classified as, with information about how often each synonym was classified as a particular target.

| Synonym | Classified as... |
|-------------------|--|
| Counseling | Consultation |
| Thought | Wisdom |
| Knowledge | Goal |
| Experience | Wisdom |
| Intention | Concept |
| Plan (14) | Concept 4, idea 3, consultation 2, attention 2, goal 2, wisdom 1 |
| Motive (14) | Consultation 4, attention 4, goal 2, idea 2, wisdom 1 |
| Visualization (3) | Attention 3 |
| Objective (5) | Goal 4, consultation 1 |
| Maxim (7) | Idea 2, concept 2, impression 1, attention 1, wisdom 1 |
| Understanding (4) | Attention 2, goal 1, impression 1 |
| Belief (7) | Impression 4, wisdom 2, idea 1 |
| Advice (7) | Attention 3, wisdom 2, consultation 1, goal 1 |

As stated above, there are hardly any ‘real synonyms’. The network often succeeds in providing an output noun that is at least similar to the word whose context was the input.

Automatic Relevance Determination

The second purpose of this study was to obtain estimates of the relevance of the context feature domains. We used Automatic Relevance Determination (ARD; Neal, 1996) for this analysis. ARD estimates the relevance of the input features based on the amount of penalty assigned to the connections weights (see Goodman, 1996, NevProp Manual). The resulting values indicate how much impact each feature has on the network output *relative* to the impact of the other features. The values are given in percent and ranged from .73% to 6.36% for the single features.

In the previous study, the ARD results suggested that the verb and world knowledge domain were of major relevance for the task. This result was confirmed by a subsequent sensitivity analysis as well as by a discriminant analysis of the training data. Interestingly, the average ARD values in this study did not differ a lot among modules, with the exception of the module of ontological status. The verb categories and the world knowledge domain rank among the least relevant modules for the present network. It is conceivable that in a future model with yet more target nouns these values may change again.

The ARD estimates are listed in Table 4. They were averaged within each module to provide an impression of the relative relevance of the modules. A somewhat surprising result is the high ARD estimate for the syntax module. In the previous analysis, most syntactic features had

ARD values below 1%. These features are more discriminative for the present set of target nouns.

Another domain with a high ARD value is ‘person / reference’, which indicates whether the target noun has an explicit reference within a sentence. Examples are the sentences “The idea was to (...)”, or “It is the company’s goal to achieve (...)”. The other feature in this category indicates whether the target is something that a person possesses, e.g., “I like *your* idea.”

Table 4: ARD values for the modules, averaged across the features.

| Feature module | Number of features | ARD in % |
|--------------------|--------------------|----------|
| World knowledge | 10 | 1.53 |
| Ontological status | 4 | 4.40 |
| Adjective | 6 | 1.41 |
| Verb | 14 | 1.58 |
| Dynamics | 3 | 1.98 |
| Case role | 6 | 1.50 |
| Time reference | 3 | 1.99 |
| Syntax | 5 | 2.27 |
| Person / reference | 2 | 2.16 |

Discussion

We have presented a neural network that can distinguish seven abstract nouns on the basis of their verbal context, as represented by a set of 53 context features. In the majority of test cases, the network could identify the correct target. With respect to our strong hypothesis, the results seem to support the view that context determines abstract noun meanings. The evident difficulties of the network to distinguish similar target nouns is not inconsistent with such a view. The network’s performance further demonstrates that the semantic and syntactic features of the contexts of abstract nouns contain distinctive information. The findings support our hypothesis that the meaning of abstract nouns can be derived from and may be constructed from context information other than merely co-occurrence data.

We have further tested whether the features are related to the target nouns’ *meaning*, rather than to particular words, by testing the network’s performance on sentences that contained synonyms of the target nouns. For a few target nouns, highly similar nouns could be found. The network could classify many contexts of these similar nouns as their synonymous target nouns.

The results provided by the network provide insight into which elements of context, or what kinds of context information, provide the critical information. The ARD values suggest that for the present network, all modules are involved in the computation of the output. The ARD values have not been stable between the six-target (Wiemer-Hastings, in press) and the present seven-target noun network. These values might further change as we add new target nouns to the network. It is possible that the values will stabilize at some number of target nouns.

It is conceivable that there is an upper limit in the number of abstract concepts that a network, such as the one discussed here, can distinguish based on this set of features. It is unclear, however, where such a limit would be. We are currently working on extending the network to more target nouns. The performance of the network on highly related targets like *idea* and *concept* suggests that the identification task is increasingly difficult as more interrelated target nouns are included. On the other hand, a network with similar target nouns *only* could provide us with valuable information about how similar nouns can be distinguished. Possibly, more features and finer distinctions within the modules would be necessary.

We are currently working on developing a network that will be trained to distinguish concrete nouns based on the same set of features. Our hypothesis is that the contextual representation of abstract nouns and concrete nouns might not contain the same information, that is, concrete and abstract nouns might be related to different aspects of their contexts. It is also possible that the contextual representation of concrete nouns does not have the same function for concrete concept representation as it does for abstract nouns. We assume that context features as the ones discussed here do not determine the meanings of concrete nouns, as they have intrinsic representations. That is, it is possible that a network cannot learn to distinguish between concrete nouns based on context.

Acknowledgments

This work was supported in part by the National Science Foundation under grant number SBR-9720314, Learning and Intelligent Systems Unit.

References

- Anderson, R. C. (1990). Inferences about word meaning. In A. C. Graesser & G. H. Bower (Eds.), *Inferences and text comprehension*. San Diego: Academic Press.
- Barsalou, L. (1982). Context-independent and context-dependent information in concepts. *Memory & Cognition*, 10, 82-93.
- Burgess, C., & Lund, K. (1997). Representing abstract words and emotional connotation in a high-dimensional memory space. *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Francis, & Kucera (1982). *Frequency analysis of english usage*. Boston: Houghton Mifflin.
- Goodman, P. (1996). *NevProp software, version 3*. Reno, NV: University of Nevada.
- Hamberger, M. J., Friedman, D., & Rosen, J. (1996). Completion norms collected from younger and older adults for 198 sentence contexts. *Behavior Research Methods, Instruments & Computers*, 28, 102-108.
- Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. Cambridge University Press.

- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104, 211-240.
- Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, and Computers*, 28, 203-208.
- McKeown, M. G. (1985). The acquisition of word meaning from context by children of high and low ability. *Reading Research Quarterly*, 20, 482-496.
- Miller, G. A. (1990). WordNet: An on-line lexical database. *International Journal of Lexicography*, 3.
- Miller, G. A., & Charles, W. G. (1991). Contextual constraints of semantic similarity. *Language and Cognitive Processes*, 6, 1-28.
- Neal, R. M. (1996). *Bayesian learning for neural networks*. New York: Springer.
- NexisLexis [Online database].
- Plaut, D. & Shallice, T. (1991). Effects of word abstractness in a connectionist model of deep dyslexia. *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Quine, W.V.O. (1960). *Word and object*. Cambridge: MIT Press.
- Roget's international thesaurus (4th Rev. ed.). (1977). New York: Crowell Company.
- Schwanenflugel, P. J., & Shoben, E. J. (1983). Differential context effects in the comprehension of abstract and concrete verbal materials. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 9, 82-102.
- Sternberg, R., & Powell, J. (1983). Comprehending verbal comprehension. *American Psychologist*, 38, 878-893.
- Wiemer-Hastings, K. (in press). Abstract noun classification: using a neural network to match word context and word meaning. *Behavior Research Methods, Instruments, and Computers*.