Quantitative Semantic Variation
in the Contexts of Concrete and Abstract Words

Daniela Naumann   Diego Frassinelli   Sabine Schulte im Walde

Institut für Maschinelle Sprachverarbeitung, Universität Stuttgart

`{naumanda,frassinelli,schulte}@ims.uni-stuttgart.de`

Abstract

Across disciplines, researchers are eager to gain insight into empirical features of abstract vs. concrete concepts. In this work, we provide a detailed characterisation of the distributional nature of abstract and concrete words across 16,620 English nouns, verbs and adjectives. Specifically, we investigate the following questions: (1) What is the distribution of concreteness in the contexts of concrete and abstract target words? (2) What are the differences between concrete and abstract words in terms of contextual semantic diversity? (3) How does the entropy of concrete and abstract word contexts differ? Overall, our studies show consistent differences in the distributional representation of concrete and abstract words, thus challenging existing theories of cognition and providing a more fine-grained description of their nature.

1 Introduction

The complete understanding of the cognitive mechanisms behind the processing of concrete and abstract meanings represents a key and still open question in cognitive science (Barsalou and Wiemer-Hastings, 2005). More specifically, the psycholinguistic literature reports extensive analyses of how concrete concepts are processed, however there is still little consensus about the nature of abstract concepts (Barsalou and Wiemer-Hastings, 2005; McRae and Jones, 2013; Hill et al., 2014; Vigliocco et al., 2014).

The *Context Availability Theory* represents one of the earliest theoretical approaches aiming to account for the differences between concrete and abstract concepts (Schwanenflugel and Shoben, 1983). This theory suggests that meaning arises from the ability to create an appropriate context for a concept, which has proven to be more challenging (i.e., enforcing higher reaction times and

larger number of errors) for abstract than for concrete concepts. In a computational study, Hill et al. (2014) quantitatively analysed the distinction between concrete and abstract words in a large corpus. Overall, they showed that abstract words occur within a broad range of context words while concrete words occur within a smaller set of context words. Similarly, Hoffman et al. (2013) and Hoffman and Woollams (2015) analysed the concrete vs. abstract dichotomy in terms of their semantic diversity, demonstrating that concrete words occur within highly similar contexts while abstract words occur in a broad range of less associated contexts (i.e., exhibiting high semantic diversity). These computational findings are fully in line with the Context Availability Theory: the processing time of concrete words is generally shorter than the processing time of abstract words, as abstract words are attached to a broad range of loosely associated words.

More recently, embodied theories of cognition have suggested that word meanings are grounded in the sensory-motor system (Barsalou and Wiemer-Hastings, 2005; Glenberg and Kaschak, 2002; Hill et al., 2014; Pecher et al., 2011). According to this account, concrete concepts have a direct referent in the real world, while abstract concepts have to activate a series of concrete concepts that provide the necessary situational context required to successfully process their meanings (Barsalou, 1999).

These interdisciplinary outcomes are not fully supported by recent computational studies showing different contextual patterns for concrete and abstract words in text compared to the literature (Bhaskar et al., 2017; Frassinelli et al., 2017). It is becoming clear, however, that the inclusion of information regarding the concreteness of words plays a key role in the automatic identification of non-literal language usage (Turney et al., 2011;

Köper and Schulte im Walde, 2016, 2017).

The aim of the current study is thus to provide a contextual description of the distributional representation of these two classes of words, to gain insight into empirical features of abstract vs. concrete concepts. This would represent an essential contribution to the resolution of the debate about meaning representation within the human mind, and thereby also help to enhance computationally derived models that are concerned with meaning derivation from text.

2 Hypotheses

Based on the existing psycholinguistic and computational evidence reported in the previous section, we formulate three hypotheses regarding the distributional nature of concrete and abstract words that we will test in the following studies.

(1) The contexts of both concrete and abstract words are mainly composed of concrete words.

This first hypothesis directly tests the general claim of grounding theories: both concrete and abstract words require the activation of a layer of situational (concrete) information in order to be successfully processed (Barsalou and Wiemer-Hastings, 2005). According to the *Distributional Hypothesis* (Harris, 1954; Firth, 1968), similar linguistic contexts tend to imply similar meanings of words. Thus, we suggest to perform a distributional semantic analysis in order to quantitatively investigate the contexts that concrete and abstract words frequently co-occur within.

(2) Abstract words occur in a broad range of distinct contexts whereas concrete words appear in a limited set of contexts.

Based on the computational study by Hill et al. (2014), we expect to find concrete words appearing in a more restricted set of contexts in comparison to abstract words, which should occur in a broad range of contexts. This second hypothesis is explored by providing two fine-grained analyses of the extension and variety in contexts of concrete and abstract words.

(3) Abstract words are more difficult to predict than concrete words, due to their higher contextual variability.

Building upon the previous hypothesis and on the studies by Hoffman et al. (2013) and Hoffman and Woollams (2015), we aim to show that concrete words are easier to predict than abstract words. Specifically, we expect higher entropy values for abstract than for concrete contexts, indicating that on average, we need more information to uniquely encode an abstract word than a concrete word (Shannon, 2001). The reason resides within the high context variability of abstract words: there is a large number of highly probable words satisfying these contexts. In contrast, we expect concrete words to occur in a limited set of different contexts because there is only a restricted amount of words that have a high probability to fit a specific context. Thus, we estimate the entropy value of concrete contexts to be lower than the entropy value of abstract contexts.

In the three studies reported in this paper, we systematically test these three hypotheses regarding concrete vs. abstract words, by performing quantitative analyses of the distributional representations across the word classes of nouns, verbs and adjectives.

3 Materials and Method

For our studies, we selected nouns, verbs and adjectives from the Brysbaert et al. (2014) collection of concreteness ratings for 40,000 English words. In total we used 16,620 target words including 9,240 nouns, 3,976 verbs and 3,404 adjectives.[1] Each word in this collection has been scored by humans according to its concreteness on a scale from 1 (abstract) to 5 (concrete).

Our distributional semantic representations of the target words were built by extracting co-occurrences from the POS-tagged version (Schmid, 1994) of the sentence-shuffled English *COW* corpus *ENCOW16AX* (Schäfer and Bildhauer, 2012). We originally constructed three different spaces with window sizes of 2, 10, and 20 context words surrounding the target, and performed parallel analyses for all the three spaces. Since we did not find any relevant differences between the three spaces, we will report only the analyses based on the distributional space from a window size of 20 context words. Moreover, we

---

[1]The reason why we only used a subset of the available targets was that these were also covered in an extensive selection of behavioural measures, such as valency scores (Warriner et al., 2013) and reaction times (Balota et al., 2007) which we aim to include in further analyses.

restricted the dimensions in our matrix to 16,620 × 16,620 (target words × context words). By using the target words also as context words, we had knowledge about the concreteness score of each context word. In a follow-up study, we performed the same analyses extracting co-occurrences from the British National Corpus (Burnard, 2000). Even though both the size and the nature of these two corpora are extremely different, the results did not show any significant difference.

In order to get a clearer picture about empirical distributional differences for concrete vs. abstract targets, we focused some of our analyses only on the most concrete and abstract targets, expecting words with mid-range concreteness scores to be more difficult in their generation by humans and consequently noisier in their distributional representation. For this reason, we analysed the 1,000 most concrete (concreteness range: 4.82 - 5.00) and the 1,000 most abstract (1.07 - 2.17) nouns, the 500 most concrete (4.71 - 5.00) and most abstract (1.12 - 2.21) verbs, and the 200 most concrete (4.34 - 5.00) and most abstract (1.19 - 1.64) adjectives. On the other hand, context was not subset and consisted of the complete set of 16,620 nouns, verbs and adjectives.

4 Study 1: Analysis of Concrete vs. Abstract Co-Occurrences

In this study we test the validity of hypothesis (1): the contexts of both concrete and abstract words are mainly concrete. For this purpose, we analyse the distributions of the 16,620 context dimensions for their concreteness, by the parts-of-speech of target and context words.

Noun Targets Figure 1 reports the distribution of noun, verb and adjective contexts for the 1,000 most abstract target nouns (Figure 1a) in comparison to the 1,000 most concrete target nouns (Figure 1b). As clearly shown in Figure 1a, the majority of contexts of an abstract noun are also abstract: noun, verb and adjective context words all show the maximum peak at low concreteness scores. On the contrary, the distributions of the contexts of concrete nouns shown in Figure 1b vary according to POS. The nouns in the context of concrete noun targets are also very concrete as shown by the high red bar at concreteness 4.5–5.
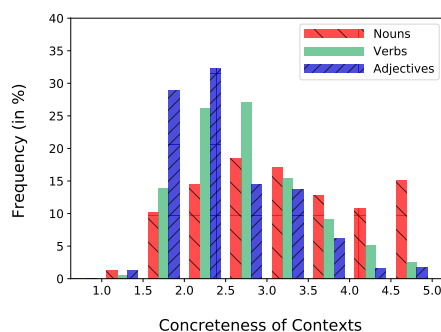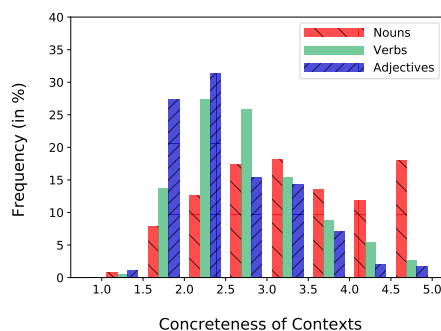

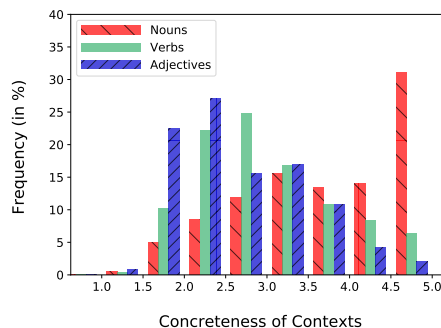(a) Contexts of abstract noun targets.
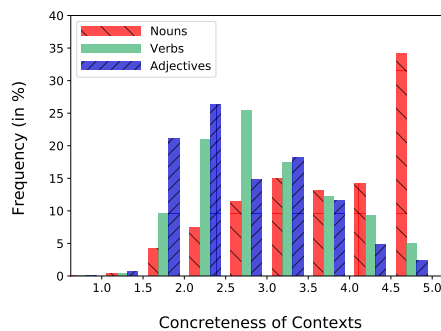
(b) Contexts of concrete noun targets.

Figure 1: Concreteness scores of context words (nouns, verbs, adjectives) of the 1,000 most abstract and concrete noun targets.

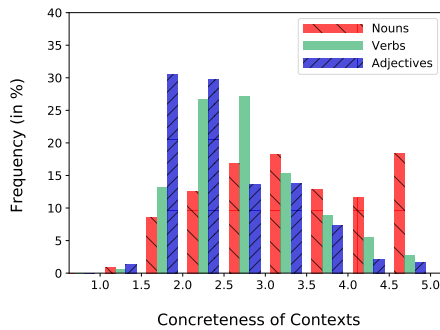(a) Contexts of abstract verb targets.

(b) Contexts of concrete verb targets.

Figure 2: Concreteness scores of context words (nouns, verbs, adjectives) of the 500 most abstract and concrete verb targets.

On the other hand, verbs and adjectives show a similar pattern to Figure 1a: a greater distribution with low concreteness scores.

Verb Targets Figure 2 shows a very comparable pattern to the one described for noun targets. Contexts of abstract verbs are, on average, also abstract, regardless of their POS. On the other hand, the verbs and adjectives in the contexts of concrete verb targets are mainly abstract, while the nouns are mainly concrete.

Adjective Targets Again, Figure 3 shows the same pattern as the one reported for nouns and verbs.

(a) Contexts of abstract adjective targets.
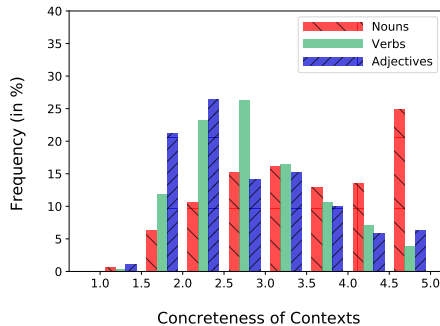
(b) Contexts of concrete adjective targets.

Figure 3: Concreteness scores of context words (nouns, verbs, adjectives) of the 200 most abstract and concrete adjective targets.

Discussion Table 1 reports an overview of the outcomes of this first study. The "X" indicates the predominant contextual class (abstract vs. concrete words) for each target class by POS. All in all, our results partly disagree with our first hypothesis induced from observations in the literature, within the scope of which we expected the context of concrete and abstract words to be mostly composed of concrete words.

Target Words	Context Words					
	abst. NN	abst. V	abst ADJ	conc. NN	conc. V	conc. ADJ
abstract NN	X	X	X			
abstract V	X	X	X			
abstract ADJ	X	X	X			
concrete NN		X	X	X		
concrete V		X	X	X		
concrete ADJ		X	X	X		

X = most frequent context type

Table 1: Evaluation of hypothesis (1).

More specifically, our first hypothesis is confirmed, on the one hand, by the contextual distribution of concrete target nouns, due to the fact that they frequently appear with other concrete nouns. On the contrary, it is rejected by the contextual ratio of abstract nouns as they primarily co-occur with other abstract nouns. Thus, as we based our hypothesis on the theory of embodied cognition, the observed contextual pattern of abstract nouns challenges this theory.

Another evidence in favour of our hypothesis comes from the nouns in the context of concrete verbs and adjectives that are mainly concrete. In contrast, concrete and abstract nouns, verbs and adjectives elicit the same contextual pattern regarding context verbs and adjectives. They co-occur with abstract verbs and abstract adjectives to a large extent, which does not support the expectations based on the existing literature.

5 Study 2: Semantic Diversity of Context

In this study, we test our second hypothesis: abstract words occur in a broad range of distinct contexts whereas concrete words appear in a limited set of different contexts. In the following sections we report two studies where we analyse (i) the number of non-zero dimensions in the representation of concrete vs. abstract words, and (ii) the degree of semantic variability in their contexts.

5.1 Non-Zero Dimensions

The analysis of the number of non-zero dimensions in the vector representation of concrete and abstract words provides a first indicator of the contextual richness of our targets. Based on Hill et al. (2014), we expect concrete target words to have significantly less diverse context dimensions than abstract target words, as the former should co-occur within a restricted set of context words. Therefore, we expect the portion of non-zero context dimensions to be smaller for concrete than for abstract target words.

The following analyses compare the proportions of non-zero context dimensions between the 1,000 highly concrete (blue boxes) and highly abstract (red boxes) target nouns, 500 verbs, and 200 adjectives, based on raw frequency counts. For each POS, we compared the proportion of non-zero dimensions in the full vectors of 16,620 context words for concrete and abstract target words (left side), and the number of non-zero dimensions with the same part-of-speech of the target (respectively, 9,240 context nouns, 3,976 context verbs, 3,404 context adjectives). The star (⋆) indicates the mean number of non-zero dimensions.
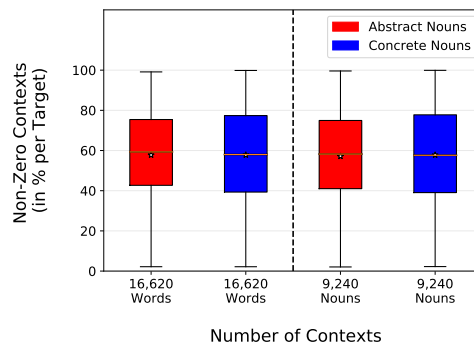
Noun Targets As shown in Figure 4, the comparison of non-zero context dimensions of concrete (M = 57.80, SD = 23.07) and abstract (M = 57.78, SD = 22.57) target nouns does not show any significant difference (t(33238) = -0.02, p = 0.98). This result indicates that concrete and abstract target nouns co-occur with a similar amount of context words. We can observe the exact same pattern when we restrict the contexts to nouns only: no significant difference between the number of non-zero context noun dimensions for concrete (M = 32.12, SD = 12.98) and abstract (M = 31.78, SD = 12.76) target nouns (t(18478) = -0.59, p = 0.56).
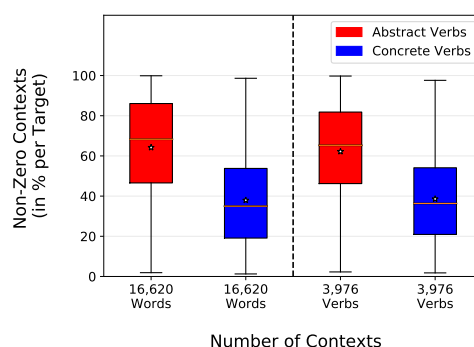
Verb Targets Figure 5 reports the number of non-zero dimensions for concrete and abstract verbs. When considering the full set of contexts (left side), concrete words (M = 37.93, SD = 22.5) have significantly less active contexts than abstract words (M = 64.2, SD = 25.73; t(33238) = 17.18, p < 0.001). The exact same outcome is shown when focusing only on verbs as contexts (t(7950) = 16.3, p < 0.001).

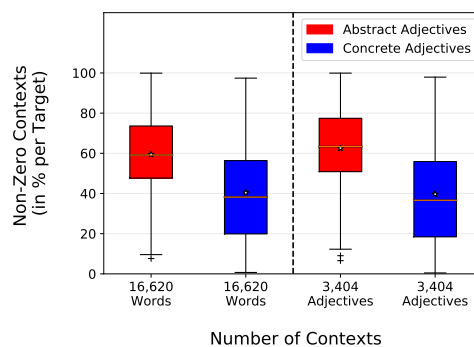**Adjective Targets** The analysis of the adjectives in Figure 6 indicates that the number of non-zero dimensions for concrete and abstract adjectives follows the same pattern as the verbs. When considering the full set of contexts (left side), concrete adjectives (M = 40.4, SD = 24.7) have significantly less active contexts than abstract adjectives (M = 59.46, SD = 19.11, t(33238) = 8.63, p < 0.001). The exact same outcome is shown when focusing only on adjectives as contexts (t(6806) = 10.15, p < 0.001).

Figure 4: Non-zero dimensions in the contexts of the 1,000 most abstract (red boxes) and concrete (blue boxes) noun targets.

Figure 5: Non-zero dimensions in the contexts of the 500 most abstract (red boxes) and concrete (blue boxes) verb targets.

Figure 6: Non-zero dimensions in the contexts of the 200 most abstract (red boxes) and concrete (blue boxes) adjectives.

5.2 Semantic Diversity of Context

Based on hypothesis (2), we expect the contexts of concrete words to be more similar among themselves than the contexts of abstract words. We test this hypothesis by computing the semantic diver-

sity of the contexts of concrete and abstract targets. Semantic diversity corresponds to the inverse of the average semantic similarity of each pair of context dimensions of a word (Hoffman et al., 2013). In order to control for pure frequency effects, we transformed the co-occurrence frequency counts into local mutual information (LMI) scores (Evert, 2005).

The study reports the average cosine similarity between context dimensions for concrete and abstract words; the analysis is conducted incrementally, including the top-k most associated context dimensions (from 5 to 16,620 associates) sorted by their LMI scores.

Noun Targets Figure 7 reports the average semantic similarity between the context dimensions of the 1,000 most concrete (blue boxes) and the 1,000 most abstract (red boxes) target nouns. The analysis is performed step-wise from left to right, starting with the average similarity between the 5 most associated contexts and moving up to the average similarity between all 16,620 context dimensions. Overall, while increasing the number of dimensions, both the mean similarity and also the differences in mean between concrete and abstract words drop, while remaining significant. The difference between the mean cosine similarity of the most associated contexts of concrete (M = 0.32, SD = 0.14 at k = 5) and abstract (M = 0.20, SD = 0.13 at k = 5) target nouns is significant ($p < 0.001$ at k = 5).

Verb Targets As shown in Figure 8, there are no significant differences (p = 0.38 at k = 5) in the similarity of the context dimensions of the 500 most concrete (M = 0.23, SD = 0.15 at k = 5) and most abstract (M = 0.23, SD = 0.16 at k = 5) verb targets.

Adjective Targets When analysing the similarity of the contexts of the 200 most concrete and abstract adjectives we see (Figure 9) the same pattern as shown for nouns. The average similarity of the most associated contexts is significantly higher (p<0.001 at k = 5) for concrete (M = 0.26, SD = 0.14 at k = 5) than for abstract (M = 0.17, SD = 0.12 at k = 5) target adjectives.
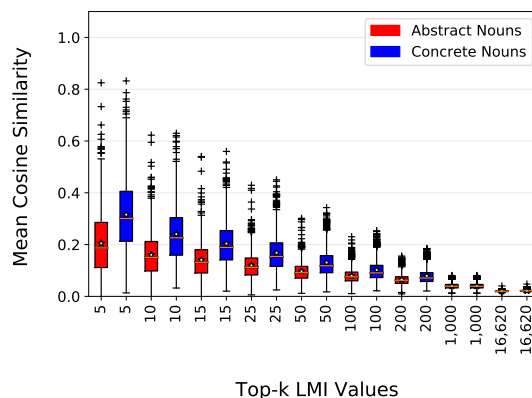

Figure 7: Mean cosine similarities between contexts of 1,000 noun targets.
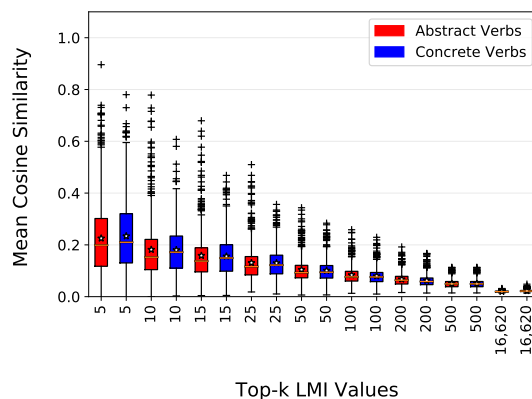

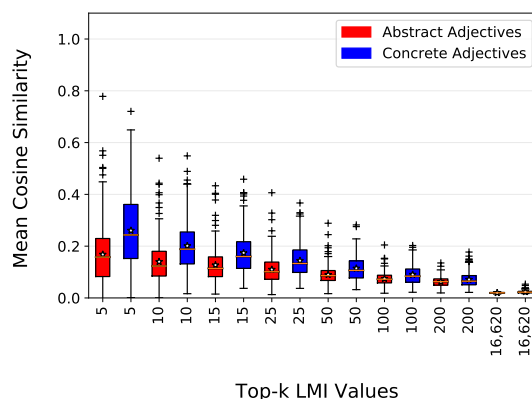Figure 8: Mean cosine similarities between contexts of 500 verb targets.

Figure 9: Mean cosine similarities between contexts of 200 adjective targets.

5.3 Discussion

According to hypothesis (2), we expected abstract words to occur in a broader range of distinct contexts and concrete words to appear in a more lim-

ited set of different contexts. Moreover, the contexts of concrete words should be more restricted and more similar to each other compared to the contexts of abstract words. The results discussed only partially support this hypothesis.

The analysis of the number of non-zero context dimensions for concrete and abstract target verbs and adjectives show results in line with hypothesis (2). On the contrary, concrete and abstract target nouns share the same number of non-zero dimensions. The analysis of the similarity between contexts of concrete and abstract target nouns and adjectives supports our hypothesis; while we do not see any significant difference when analysing the verbs.

6 Study 3: Entropy of Concrete and Abstract Words

In this study we test our third hypothesis: abstract words are more difficult to predict than concrete words, due to their higher contextual variability. In study 2 we already started investigating this phenomenon using semantic diversity. In the current study we will use entropy as a measure of variability (Shannon, 2001):

$$H(X) = -\sum_{x \in X} p(x)\log_2 p(x) \qquad (1)$$

Based on the assumption that abstract words occur within a high number of distinct contexts, we expect the entropy of abstract words to be higher than the entropy of concrete words.

Noun Targets   Figure 10 reports the average entropy in the context of the top 1,000 most abstract (on the left side) and most concrete (on the right side) target nouns. Regarding the 1,000 most abstract target nouns, the entropy of the 1,000 most abstract context nouns (M = 7.42, SD = 0.58) is significantly higher (p < 0.001) than the entropy of the 1,000 most concrete context nouns (M = 6.44, SD = 0.77). A similar pattern emerges in the analysis of the entropy of the contexts of the 1,000 most concrete target nouns: the difference between concrete (M = 6.64, SD = 0.61) and abstract contexts (M = 7.21, SD = 0.54) is statistically significant (p < 0.001).

Verb Targets   Similarly to nouns (see Figure 11), also the abstract contexts of both concrete and abstract target verbs show significantly (p < 0.001) higher entropy (concrete target:

M = 6.1, SD = 0.58; abstract target: M = 6.55, SD = 0.49) than the entropy of their concrete contexts (concrete target: M = 4.70, SD = 0.89; abstract target: M = 5.50, SD = 0.86).
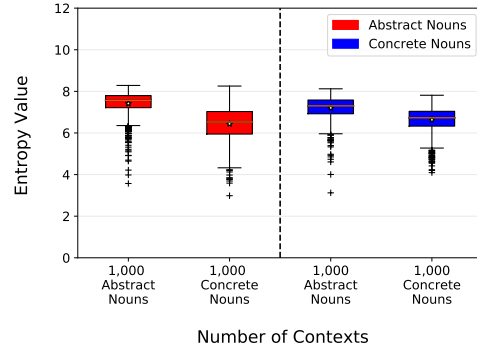

Figure 10: Entropy of 1,000 most abstract (left side) and 1,000 most concrete (right side) noun targets.
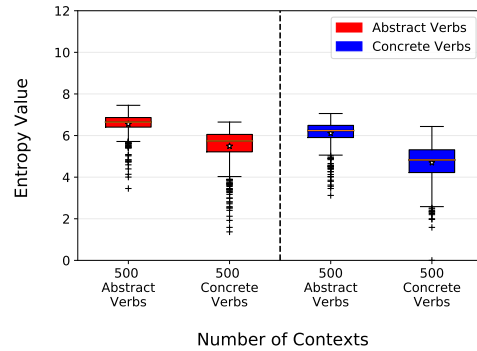

Figure 11: Entropy of 500 most abstract (left side) and 500 most concrete (right side) verbs targets.
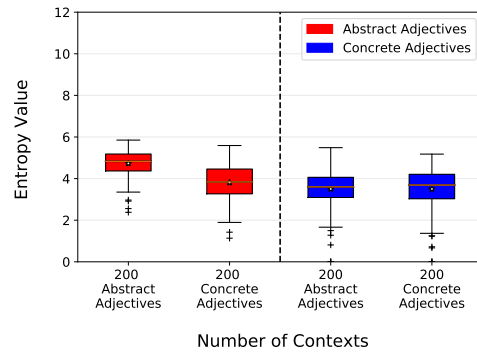

Figure 12: Entropy of 200 most abstract (left side) and 200 most concrete (right side) adjectives targets.

Adjective Targets   The same pattern seen for nouns and verbs (see Figure 12) describes also

the entropy of target concrete and abstract adjectives. Abstract contexts show significantly ($p < 0.001$) higher entropy (concrete target: M = 3.5, SD = 0.88; abstract target: M = 4.73, SD = 0.61) than the entropy of their concrete contexts (concrete target: M = 3.50, SD = 0.98; abstract target: M = 3.81, SD = 0.87).

Discussion The results of this study support the predictions of hypothesis (3): concrete contexts have significantly lower entropy than abstract contexts irrespective of the POS of their target words.

7 Conclusions

The aim of this work was to provide a very detailed description of the contextual representation of concrete and abstract English nouns, verbs and adjectives. Table 2 summarises the most important findings. 1) Concrete target nouns, verbs and adjectives mainly co-occur with concrete nouns and with abstract verbs and adjectives, while abstract target words always co-occur with abstract words. 2a) The contexts of abstract target verbs and adjectives are broader (less non-zero dimensions) than those of concrete targets verbs and adjectives. On the other hand, concrete and abstract target nouns have a similar number of non-zero dimensions. 2b) The most associated contexts of concrete nouns and adjectives are significantly more similar to each other than the contexts of abstract nouns and adjectives. However, no difference emerges between the contexts of verbs. 3) The concrete contexts of concrete and abstract targets (nouns, verbs, adjectives) have significantly lower entropy values than their abstract contexts. Overall, hypotheses (1) and (2) are not fully supported by our analyses; on the contrary, the predictions made in hypothesis (3) are confirmed.

The three studies described in this paper thus show consistent differences in the contexts of concrete and abstract words and yield patterns that challenge the grounding theory of cognition. In their analyses on noun and verb comprehension, Barsalou (1999) and Richardson et al. (2003) suggest that humans process abstract concepts by creating a perceptual representation. These representations are inherently concrete because they are stored as "experiential traces" generated through the exposure to real world situations using our five senses (Van Dam et al., 2010). In the instructions of their norming study, Brysbaert et al. (2014, p. 906) describe concrete words in a similar way:

"some words refer to things or actions in reality, which you can experience directly through one of the five senses". On the contrary, our study is aligned more with recent theories claiming a representational pluralism that includes both perceptual and non-perceptual features (Dove, 2009).

While the reported cognitive theories describe general patterns emerging from the distinction between concrete and abstract words, the novelty of our study is to provide a fine-grained analysis of the distributional nature of these words and an attempt to explain their similarities and differences from a data-driven perspective. In our opinion, the detection of the precise properties of concrete and abstract words makes an extremely valuable contribution to the long-lasting debate about meaning representation in the human mind and to the use of this knowledge to significantly improve the performance of computational models.

Acknowledgments

References

David A. Balota, Melvin J. Yap, Michael J. Cortese, Keith A. Hutchison, Brett Kessler, Bjorn Loftis, James H. Neely, Douglas L. Nelson, Greg B. Simpson, and Rebecca Treiman. 2007. The English Lexicon Project. *Behavior Research Methods*, 39(3):445–459.

Lawrence W. Barsalou. 1999. Perceptual Symbol Systems. *Behavioral and Brain Sciences*, 22:577–660.

Lawrence W. Barsalou and Katja Wiemer-Hastings. 2005. Situating Abstract Concepts. In D. Pecher and R. Zwaan, editors, *Grounding Cognition: The Role of Perception and Action in Memory, Language, and Thinking*, chapter 7, pages 129–163. Cambridge University Press, New York.

Sai Abishek Bhaskar, Maximilian Köper, Sabine Schulte im Walde, and Diego Frassinelli. 2017. Exploring Multi-Modal Text+Image Models to Distinguish between Abstract and Concrete Nouns. In *Proceedings of the IWCS Workshop on Foundations of Situated and Multimodal Communication*.

Marc Brysbaert, Amy Beth Warriner, and Victor Kuperman. 2014. Concreteness Ratings for 40 Thousand generally known English Word Lemmas. *Behavior Research Methods*, 64:904–911.

POS	Primary Co-Occurrence	Semantic Diversity of Context	Entropy
concrete nouns	*concrete* nouns *abstract* verbs *abstract* adjectives	similar number of contexts; high context similarity	low
abstract nouns	*abstract* nouns *abstract* verbs *abstract* adjectives	similar number of contexts; low context similarity	high
concrete verbs	*concrete* nouns *abstract* verbs *abstract* adjectives	low number of contexts; similar context similarity	low
abstract verbs	*abstract* nouns *abstract* verbs *abstract* adjectives	high number of contexts; similar context similarity	high
concrete adjectives	*concrete* nouns *abstract* verbs *abstract* adjectives	low number of contexts; high context similarity	low
abstract adjectives	*abstract* nouns *abstract* verbs *abstract* adjectives	high number of contexts; low context similarity	high

Table 2: Overview of our main findings.

L. Burnard. 2000. Reference Guide for the British National Corpus (world edition).

Guy Dove. 2009. Beyond Perceptual Symbols: a Call for Representational Pluralism. *Cognition*, 110(3):412–31.

Stefan Evert. 2005. *The Statistics of Word Co-Occurrences: Word Pairs and Collocations*. Ph.D. thesis, Institut für Maschinelle Sprachverarbeitung, Universität Stuttgart.

John R. Firth. 1968. A Synopsis of Linguistic Theory, 1930–55. In Frank R. Palmer, editor, *Selected Papers of J.R. Firth 1952–59*, Longman's Linguistics Library, pages 168–205. Longmans.

Diego Frassinelli, Daniela Naumann, Jason Utt, and Sabine Schulte im Walde. 2017. Contextual Characteristics of Concrete and Abstract Words. In *Proceedings of the 12th International Conference on Computational Semantics*, Montpellier, France.

Arthur M. Glenberg and Michael P. Kaschak. 2002. Grounding Language in Action. *Psychonomic Bulletin and Review*, 9(3):558–565.

Zellig S. Harris. 1954. Distributional structure. *Word*, 10(2-3):146–162.

Felix Hill, Anna Korhonen, and Christian Bentz. 2014. A Quantitative Empirical Analysis of the Abstract/Concrete Distinction. *Cognitive Science*, 38(1):162–177.

Paul Hoffman, Matthew A. Lambon Ralph, and Timothy T. Rogers. 2013. Semantic Diversity: A Measure of Semantic Ambiguity Based on Variability in the Contextual Usage of Words. *Behavior Research Methods*, 45(3):718–730.

Paul Hoffman and Anna M. Woollams. 2015. Opposing Effects of Semantic Diversity in Lexical and Semantic Relatedness Decisions. *Journal of Experimental Psychology: Human Perception and Performance*, 41(2):385.

Maximilian Köper and Sabine Schulte im Walde. 2016. Distinguishing Literal and Non-Literal Usage of German Particle Verbs. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 353–362, San Diego, California, USA.

Maximilian Köper and Sabine Schulte im Walde. 2017. Improving Verb Metaphor Detection by Propagating Abstractness to Words, Phrases and Individual Senses. In *Proceedings of the 1st Workshop on Sense, Concept and Entity Representations and their Applications*, pages 24–30, Valencia, Spain.

Ken McRae and Michael Jones. 2013. Semantic Memory. *The Oxford Handbook of Cognitive Psychology*, 206.

Diane Pecher, Inge Boot, and Saskia Van Dantzig. 2011. Abstract Concepts. Sensory-Motor Grounding, Metaphors, and Beyond. *Psychology of Learn-*

ing and Motivation – Advances in Research and Theory, 54:217–248.

Daniel C. Richardson, Michael J. Spivey, Lawrence W. Barsalou, and Ken McRae. 2003. Spatial Representations Activated during Real-Time Comprehension of Verbs. *Cognitive Science*, 27(5):767–780.

Roland Schäfer and Felix Bildhauer. 2012. Building Large Corpora from the Web Using a New Efficient Tool Chain. In *Proceedings of the 8th International Conference on Language Resources and Evaluation*, pages 486–493, Istanbul, Turkey.

Helmut Schmid. 1994. Probabilistic Part-of-Speech Tagging Using Decision Trees. In *Proceedings of the 1st International Conference on New Methods in Language Processing*.

Paula J Schwanenflugel and Edward J Shoben. 1983. Differential context effects in the comprehension of abstract and concrete verbal materials. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 9(1):82.

Claude E. Shannon. 2001. A Mathematical Theory of Communication. *ACM SIGMOBILE Mobile Computing and Communications Review*, 5(1):3–55.

Peter Turney, Yair Neuman, Dan Assaf, and Yohai Cohen. 2011. Literal and Metaphorical Sense Identification through Concrete and Abstract Context. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 680–690, Edinburgh, UK.

Wessel O. Van Dam, Shirley-Ann Rueschemeyer, Oliver Lindemann, and Harold Bekkering. 2010. Context effects in embodied lexical-semantic processing. *Frontiers in Psychology*, 1.

Gabriella Vigliocco, Stavroula-Thaleia Kousta, Pasquale Anthony Della Rosa, David P. Vinson, Marco Tettamanti, Joseph T. Devlin, and Stefano F. Cappa. 2014. The Neural Representation of Abstract Words: The Role of Emotion. *Cerebral Cortex*, 24(7):1767–1777.

Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. 2013. Norms of Valence, Arousal, and Dominance for 13,915 English Lemmas. *Behavior Research Methods*, 45(4):1191–1207.