Quantifying the Conceptual Combination Effect on Word Meanings

Nora Aguirre-Celis (naguirre@cs.utexas.edu)
ITESM, Monterrey, Mexico & The University of Texas at Austin
Department of Computer Science, 2317 Speedway
Austin, TX 78712 USA

Risto Miikkulainen (risto@ cs.utexas.edu)
The University of Texas at Austin
Department of Computer Science, 2317 Speedway
Austin, TX 78712 USA

Abstract

How do people understand concepts such as *dog, aggressive dog*, *dog house* or *house dog*? The meaning of a concept depends crucially on the concepts around it. While this hypothesis has existed for a long time, only recently it has become possible to test it based on neuroimaging and quantify it using computational modeling. In this paper, a neural network is trained with backpropagation to map attribute-based semantic representations to fMRI images of subjects reading everyday sentences. Backpropagation is then extended to the attributes, demonstrating how word meanings change in different contexts. Across a large corpus of sentences, the new attributes are more similar to the attributes of other words in the sentence than they are to the original attributes, demonstrating that the meaning of the context is transferred to a degree to each word in the sentence. Such dynamic conceptual combination effects could be included in natural language processing systems to encode rich contextual embeddings to mirror human performance more accurately.

Keywords: Context Effect; Concept Representations; Conceptual Combination; fMRI Data Analysis; Neural Networks; Embodied Cognition

Introduction

In the embodied cognition approach. (Barsalou, 2008, Binder et al., 2009), the meaning of a concept is not a set of verbal features that people associate with the concept, but rather a set of neural processing modalities that are involved while experiencing instances of the concept. This approach provides a direct correspondence between conceptual content and neural representations, and suggests that concepts can be represented through a number of weighted semantic dimensions that correspond to different brain areas. Recently it has become possible to ground this approach to brain imaging. In particular, Binder et al. (2009) identified a distributed large-scale brain network linked to the storage and retrieval of words. This brain network was used as the foundation for the Concept Attributes Representation (CAR) theory (a.k.a. the experiential attribute representation model). CAR theory proposes that words are represented as a set of weighted attributes stimulated by context.

People weigh concept features differently based on context, i.e., they construct a meaning dynamically according to the combination of concepts that occur in the sentence. Such conceptual combination either uses an attribute of one concept to describe another (in attribute combination) or forms some relation between two concepts to create a new one (in relational combination). In case of attribute combination, the modifier features adapt other concepts in the combination to some degree, and as a result, the words involved are alike (Wisniewski, 1998). For example, listeners must realize that *red apple* could mean just a fruit having a certain color by selecting salient features that dominate in the combination. The noun *apple* is defined by color, size, shape, taste, etc. and one or more of those dimensions will be modified during the attribute combination. In relational combination, the modifier features have nothing to do with the combination. For example *apple basket* or *apple pie* contain a variety of relations that often do not include *apple's* features as in *apple baskets* are not edible, red or a fruit. To help understand that *apple pie* is made of apples but *apple baskets* are not, a thematic relation needs to be built based on world knowledge about plausible combinations. Both attribute and relational combinations play an important role in the construction of new or complex concepts (Gagné & Shoben, 1997; Murphy 1990; Pecher, Zeelenberg, & Barsalou, 2004).

This paper focuses on the attribute combination process. It describes how such a dynamic construction of concepts in the brain can be quantified. This question has been studied in previous work anecdotally, by analyzing a few example cases of how the meaning attributes are weighted differently in various contexts for individual concepts, combinations of concepts, and for sentences (Aguirre-Celis & Miikkulainen, 2017, 2018). The current study expands on this prior work by evaluating the robustness and generality of these conclusions across an entire corpus of sentences and semantic roles. A neural network is trained to map brain-based semantic representations of words (CARs) into fMRI data of subjects reading everyday sentences. Backpropagation is then repeated separately for each sentence, reducing the remaining error by modifying only the CARs at the input of the network. As a result, the strengths of the attributes in the CARs change according to how important each attribute is for that sentence context.

The CAR theory is first reviewed, and the sentence collection, fMRI data, and word representation data described. The computational model is presented, followed by the experiments: an example individual case of how

conceptual combinations affect word meanings, and an aggregate study across a corpus of sentences.


Figure 1: Bar plot of the 66 semantic features for the concept *table*. The values represent average human ratings for each feature. Given that *table* is an object, it gets low weightings on human-related attributes such as Face, Speech, Head, and emotions including Happy, Sad, and Angry, and high weightings on attributes like Vision, Shape, Touch, and Manipulation.

Concept Attribute Representation Theory

CAR theory represents the basic components of meaning defined in terms of observed neural processes and brain systems thereby relating semantic content to systematic modulation of neuroimaging activity (Anderson, et al., 2016; Binder, et al., 2009). They are composed of a list of modalities that correspond to specialized sensory, motor and affective functions, and are therefore not limited to the classical sensory-motor dimensions of most embodied theories.

CARs capture aspects of experience central to the acquisition of event and object concepts, both abstract and concrete. For example, concept ratings on visual and sensory components include brightness, color, size, shape, temperature, weight, pain, etc. These aspects of mental experience model each word as a collection of a 66-dimensional feature vector that captures the strength of association between each neural attribute and the word meaning. For instance, Figure 1 shows the CAR for the concept *table*.

The attributes in CAR theory were selected after an extensive body of physiological evidence based on two assumptions: (1) All aspects of mental experience can contribute to concept acquisition and consequently concept composition; (2) experiential phenomena are grounded on neural processors representing a particular aspect of experience. For a more detailed account of the attribute selection and definition see Binder, et al., (2009, 2011, 2016a, and 2016b). The next section describes how the CAR theory is instantiated by acquiring attribute ratings from human subjects.

Data Preparation

Three data collections were used in this study: A sentence collection prepared by Glasgow et al., (2016), the fMRI images for these sentence by the Medical College of Wisconsin (Anderson, et al., 2016; Binder, et al., 2016), and

semantic Vectors (CAR ratings) for words obtained via Mechanical Turk (Anderson, et al., 2016; Binder, et al., 2009). In addition, fMRI representations were synthesized for individual words from the sentence fMRI. Each of these data collections is described in more detail below.

Sentence Collection

The sentence set was prepared for the fMRI study as part of the Knowledge Representation in Neural Systems Program (KRNS). A total of 240 sentences were composed from two to five content words from a set of 242 words (141 nouns, 39 adjectives and 62 verbs). The words were selected toward imaginable and concrete objects, actions, settings, roles, state and emotions, and events. Examples include *couple*, *author*, *boy*, *theatre*, *hospital*, *desk*, *red*, *flood*, *damaged*, *drank*, *gave*, *happy*, *old*, *summer*, *chicken*, *dog*.

The sentence collection is not fully balanced and systematic, but instead aims to be a natural sample. In order to investigate the effect of context, pairs of contrasting sentences were identified in this collection in an early study. This pairs include differences and similarities such as live mouse vs. dead mouse, family celebrated vs. happy family, and playing soccer vs. watching soccer. The resulting collection of 77 such sentences, with different shades of meaning for verbs, nouns and adjectives, as well as different contexts for nouns and adjectives was used to identify anecdotal examples (Table 1). However, the entire colection of sentences was used in the aggregate study described below.

Table 1: Contrasting Sentences. Sentence examples with differences and similarities in meaning. For instance, the role of the verb *flew* is used in two different contexts, bird and duck flying (animate) vs. plane flying (inanimate). Such sentence pairs illustrate the idea of conceptual combination well. However, the entire set of sentences was used in the aggregate study described in this paper.

SEMANTIC CONTRAST		SENTENCES
GOOD AGGRESSIVE	94	The soldier delivered the medicine.
	112	The soldier kicked the door.
ANIMAL	203	The yellow bird flew over the field.
	207	The duck flew.
OBJECT	210	The red plane flew through the cloud.
BAD PEOPLE	119	The dangerous criminal stole the television.
	152	The mob was dangerous.
NATURE	99	The flood was dangerous.

Neural fMRI Representation of Sentences

To obtain the neural correlates of the 240 sentences, subjects viewed each sentence on a computer screen while in the fMRI scanner. The sentences were presented word-by-word using a rapid serial visual presentation paradigm, with each content word exposed for 400ms followed by a 200ms inter-stimulus interval. Participants were instructed to read the sentences and think about their overall meaning.

Eleven subjects took part in this experiment producing 12 repetitions each. The fMRI data were preprocessed using standard methods, including slice timing and motion correction (AFNI software, Cox 1996). The most stable,

active and discriminative voxels were then selected and Principal Component Analysis and zero mean normalization were performed on them. These transformed brain activation patterns were converted into a single-sentence fMRI representation per participant by taking the voxel-wise mean of all repetitions (Anderson, et al., 2016; Binder, et al., 2016, 2016b). To form the target for the neural network, the most significant 396 voxels per sentence were then chosen (to match six case-role slots of the content words consisting of 66 attributes each) and scaled to [0.2..0.8].

Synthetic fMRI Word Representations

The neural data set did not include fMRI images for words in isolation. Therefore a technique developed by Anderson et al. (2016) was adopted to approximate them. The voxel values for a word were obtained by averaging all fMRI images for the sentences where the word occurs. These vectors, called SynthWords, encode a combination of examples of that word along with other words that appear in the same sentence. Thus, the SynthWord representation for *mouse* contains aspects of running, forest, man, seeing, and dead, from sentence 56:*The mouse ran into the forest* and sentence 60:*The man saw the dead mouse*.

This process of combining contextual information is similar to many semantic models in computational linguistics (Baroni et. al., 2010; Burgess, 1998; Landauer et al., 1997; Mitchell & Lapata, 2010). In other studies, this approach has been used successfully to predict brain activation (Anderson, et al., 2016; Binder, et al., 2016a, 2016b; Just, et al., 2017).

Due to the limited number of combinations, some of SynthWords became identical and were excluded from the dataset. The final collection includes 237 sentences and 236 words (138 nouns, 38 adjectives and 60 verbs).

Semantic CAR Representations for Words

CAR ratings were collected for the original set of 242 words (Glasgow et al., 2016) through Amazon Mechanical Turk. In a scale of 0-6, the participants were asked to assign the degree to which a given concept is associated to a specific type of neural component of experience (e.g., "To what degree do you think of a *table* as having a fixed location, as on a map?"). Approximately 30 ratings were collected for each word. After averaging all ratings and removing outliers, the final attributes were transformed to unit length yielding a 66-dimensional feature vector (Figure 1).

Note that this approach build its representations by directly mapping the conceptual content of a word (expressed in the questions) to the corresponding neural processes and systems for which the CAR dimensions stand. This approach thus contrasts with systems where the features are extracted from text corpora and word co-occurrence (Baroni et. al., 2010; Burgess, 1998; Harris, 1970; Landauer & Dumais, 1997).

Computational Approach

The approach for quantifying the effect of context in the fMRI data is based on the FGREP neural network (Forming Global Representations with Extended BP, Miikkulainen & Dyer, 1991). The idea is to train a neural network to predict what the sentence fMRI should be, based on the CAR representations, and then use FGREP to modify the CARs so that that prediction becomes correct.

Therefore, a simple three-layer neural network is first trained to map the CAR representations to word fMRI (in the left side of Figure 2, the mapping from CARWords, or word attribute ratings, to SynthWords, i.e., fMRI synthetic words).

After training, this network is used to predict what the sentence fMRI would be without the context effects. The SynthWords in the sentence are averaged to form this prediction called SynthSent. The SynthSent is then compared to fMRISent (the original fMRI data) to form an error signal.

That signal is backpropagated through the network (right side of figure 2), but the neural network weights are no longer changed. Instead, the error is used to change the CARWords (which is the FGREP method). This modification can be carried out through multiple iterations until the error goes to zero, or no additional change is possible (because the CAR attributes are already at their max or min limits). Eventually, the revised CARWord represents the word meaning for the current sentence such that when combined with other CARWords in the sentence, the prediction of sentence fMRI is correct.

For the experiments, the FGREP model was trained 20 times with different random seeds for each of the eleven fMRI subjects. A total of 20 different sets of 786 context-based word representations (one word representation for each sentence where the word appear) were thus produced for each subject. Afterwards, the mean of the 20 representations was used to represent each word.

Results

Previous work showed (1) that words in different contexts have different representations, and (2) these differences are determined by context (Aguirre-Celis & Miikkulainen 2017, 2018). These effects were demonstrated by analyzing individual sentence cases across multiple fMRI subjects. This paper verifies these same conclusions in the aggregate through a statistical analysis across an entire corpus of sentences. It measures how the CAR representation of a word changes in different sentences, and correlates these changes to the CAR representations of the other words in the sentence. In other words, it quantifies the conceptual combination effect statistically across sentences and subjects.

Figure 2: Using FGREP to quantify the effect of context on word meaning. Brain-based semantic representations (CARs) are changed based on a difference of predicted and actual fMRI image for a sentence. (1) Propagate CARWord to SynthWord. (2) Construct SynthSent by averaging the words into a prediction of the sentence. (3) Compare SynthSent against Observed fMRISent. (4) Backpropagate the error with FGREP for each sentence, freezing network weights and changing only CARWord. (5) Repeat until error reaches zero or CARs reach their upper or lower limits. As a result, the changes in CARs illustrate the effect of context on word meaning.

A detailed individual example of the conceptual combination effect is first presented, followed by the aggregate analysis.

The Conceptual Combination Effect

As discussed above, in CAR theory, concepts' interactions arise within multiple brain networks, activating similar brain zones for both concepts. These interactions determine the meaning of the concept combination (Binder, 2016a, 2016b).

As an example, consider the noun-verb interactions in Sentence 200: *The yellow bird flew over the field*, and Sentence 207: *The red plane flew through the cloud*. Since *bird* is a living thing, animate dimensions related to agency such as sensory, gustative, motor, affective, and cognitive experiences are expected to be activated, including potentially attributes like Speech, Taste, and Smell. In contrast, *plane flew* is expected to activate inanimate dimensions related to perceiving an object, as well as possibly Emotion, Cognition, and Attention.

Figure 3 shows the CARs for the word *flew* in the two sentences after they were modified by FGREP as described in Figure 2 and averaged across all 11 subjects. In Sentence 200 there were indeed high activations on animate attributes like Small, Pain, Smell and Taste, Audition, Music, Speech,

as well as Communication and Cognition. In contrast, Sentence 207 emphasizes perceptual features like Color, Size, and Shape, Weight, Audition, Loud, Duration, Social, Benefit, and Attention.

These results illustrates the effect of conceptual combination on word meaning. As the context varies, the overlap on neural representations create a mutual enhancement, producing a clear difference between animate and inanimate contexts. The FGREP method then encodes this effect into the CAR representations where it can be measured. In other experiments, a similar effect was observed for several other noun-verb pairs, as well as several adjective-noun pairs. In the next section the effect is quantified statistically across the entire corpus of sentences.

Aggregation Analysis

So far, the conceptual combination effect has been demonstrated in a number of example cases, like the one above, and others in earlier work (Aguirre-Celis & Miikkulainen 2017, 2018). The goal of the aggregation study in this paper is to demonstrate that the effect is robust and general across the entire corpus of sentences and case roles. The hypothesis is that similar sentences have a similar effect, and this effect is consistent across all words in the sentence.
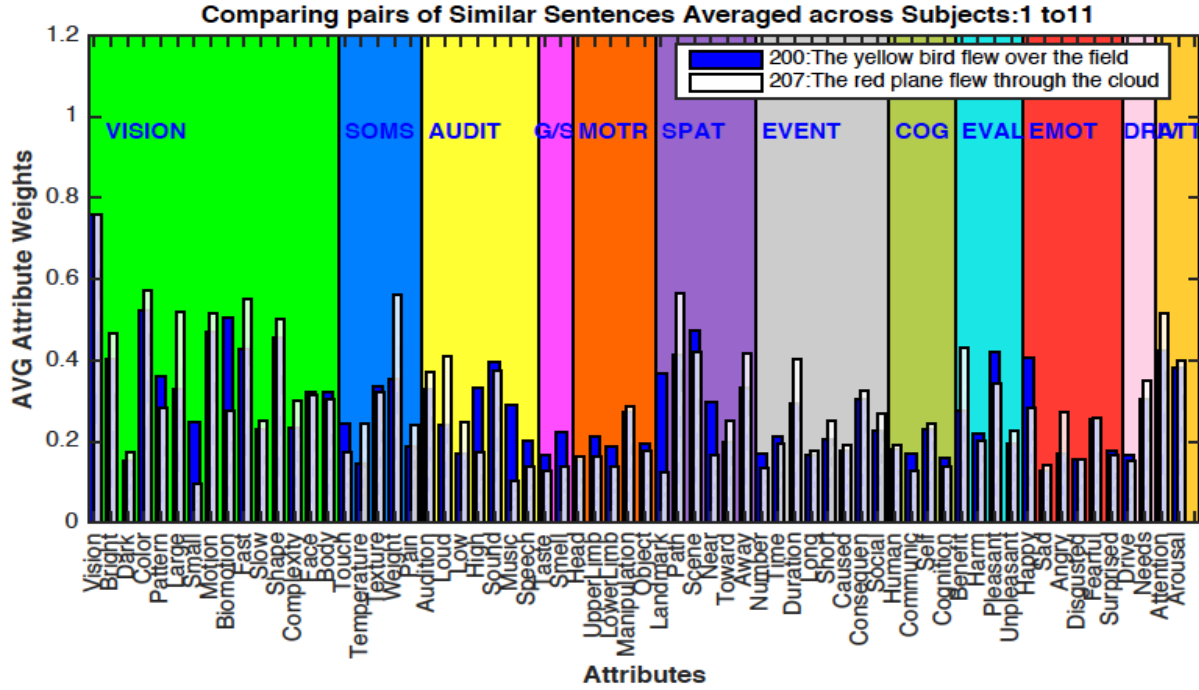
Figure 3: Contrasting the conceptual combination effect in two different sentences. In Sentence 200 (blue bars), the CAR representation modified by FGREP for the word *flew* has salient activations on animate features, presumably denoting *bird* properties like Pain, Small, Smell and Taste, and Communication. In Sentence 207 (white bars), it has high activations on inanimate object features, describing a Loud, Large, and Heavy object such as a *plane*.

This hypothesis was verified in the following process:

1. For each subject, modified CARs for each word in each sentence were formed through FGREP as described in Figure 2.
2. A representation for each sentence, SynthSent, was assembled by averaging the modified CARs.
3. Clusters of sentences were formed by running the Matlab function **linkage** on the set of SynthSents. Linkage measure the distance between clusters using the Ward method and the distance between elements with Euclidean distance. It treats each sentence as a single cluster at the beginning and then successively merges pairs of clusters. The process was stopped at 30 clusters, i.e., at the point where the granularity appeared most meaningful (e.g., sentences describing open locations vs. closed locations).
4. For each cluster, CAR representations with similar roles (agent, verb patient) were identified.
5. For each word in each such role, the differences between the modified CAR representations and the original CARs were calculated and averaged, and statistical significance of the difference measured using t-test across the entire set for each CAR dimensions.
6. The CARs of the other words in the sentence were averaged.

7. Pearson's Correlations were then calculated between the modified CARs and the averages CARs of other words across all the dimensions.
8. Similarly, correlations were calculated for the original CARs.
9. These two correlations were then compared. If the modified CARs correlate with the CARs of other words in the sentence better than the original CARs, there is evidence of context effect based on conceptual combination

In other words, this process aims to demonstrate that changes in a word CAR originate from the other words in the sentence. As in the example presented in the previous subsection, the noun-verb combination of *bird flew* and *plane flew* showed how some of the noun properties (animate/inanimate) were transferred to the verb, adapting the combination to the extent that the words share similar features. For example, if the other words in the sentence have high values in the CAR dimension for *Small*, then that dimension in the modified CAR should be higher than in the original CAR for that word. The correlation analysis measures this effect across the entire CAR representation. It measures whether the word meaning changes towards the context meaning.

The results are shown in detail in Table 2. The correlations are significantly higher for new CARs than for the original CARs across all subjects and all roles. As a summary, the average correlation was 0.3201 (STDEV

0.020) for original CAR representations and 0.3918 (STDEV 0.034) for new CAR representations. The results indeed confirm that the conceptual combination effect occurs reliably across subjects and sentences, and it is possible to quantify it by analyzing the fMRI images using the FGREP method on CAR representations.

Table 2: Correlation results. Average correlations analyzed by word class for 11 subjects comparing the original and new CARs vs. the average of the other words in the sentence. A moderate to strong positive correlation was found between new CARs and the other words in the sentence suggesting that features on one word are transferred to other words in the sentence during conceptual combination.

AVERAGE CORRELATIONS PER SUBJECT (3 ROLES)						
	ORIGINAL			NEW		
SUBJECTS	AGENT	VERB	PATIENT	AGENT	VERB	PATIENT
5051	0.2956	0.2884	0.3138	0.3908	0.3760	0.4147
5146	0.3272	0.3103	0.3476	0.3854	0.3585	0.4096
9322	0.3097	0.3049	0.3209	0.3746	0.3661	0.3905
9324	0.3264	0.3021	0.3456	0.3613	0.3373	0.3800
9362	0.3595	0.3029	0.3252	0.3918	0.3621	0.3959
9637	0.3195	0.3076	0.3391	0.3585	0.3319	0.3755
9655	0.3306	0.3045	0.3176	0.3627	0.3435	0.3835
9701	0.3839	0.3046	0.3074	0.4360	0.3992	0.4383
9726	0.3311	0.3064	0.3075	0.4185	0.3989	0.4258
9742	0.3410	0.3119	0.3250	0.3941	0.3682	0.4203
9780	0.3377	0.3023	0.3046	0.4706	0.4483	0.4610
AVERAGE	0.33293	0.30417	0.32312	0.39494	0.37182	0.40865
STDEV	0.02364	0.00611	0.01525	0.03464	0.03355	0.02670

Discussion and Future Work

This study aimed to verify the hypothesis that during sentence comprehension, people adjust the word meanings according to the combination of the concepts that occur in the sentence. This effect had been demonstrated in individual cases before, and the goal was to demonstrate it more broadly across many subjects, and entire corpus of sentences, and different semantic case roles in the sentence. The correlation results indeed demonstrated that the effect is robust, and can be quantified by analyzing fMRI images through the FGREP mechanism.

These findings are significant considering that the dataset was limited and was not designed to answer the question of dynamic effects in meaning. In the future, it may be possible to extend the data with identical contexts and contrasting contexts, and such fully balanced stimuli could be used to test the hypothesis more systematically.

Similarly, it would be desirable to extend the data with fMRI images of individual words. The current approach of synthetic words (SynthWords) is an approximation often used in computational linguistic (Baroni et. al., 2010; Burgess, 1998; Landauer et al., 1997; Mitchell & Lapata, 2010) and neural activity prediction research (Anderson, et al., 2016; Binder, et al., 2016a, 2016b; Just, et al., 2017). The FGREP process of mapping semantic CARs to SynthWords and further to sentence fMRI, refines the synthetic representations by removing noise. Still, such representations blend the meanings of many words in many sentences, therefore including word fMRI should lead to stronger and clearer results.

One important advantage of CAR theory is that it is grounded on brain representations, and therefore a good choice when mapping semantic representations to fMRI. In the future, it would be interesting to compare whether similar effects can be observed with semantic representations based on co-occurrence in text corpora, or perhaps even a combination of the two. Another important direction of future work is to take advantage of this effect in an artificial natural language processing system. The vector representations for words can be modified dynamically based on context. Such a process should match human behavior better, and result in a more effective and robust system.

Conclusion

This paper shows how word meanings change dynamically depending on context. Using FGREP as a mechanism it was possible to show that the difference between the expected and observed fMRI images can indeed be explained by a change in CARs. Across an entire corpus of sentences, the new CARs are more similar to the other words in the sentence than to the original CARs, demonstrating how features of the context are transferred to each word in the sentence. In the future it may be possible to utilize such dynamic representations in an artificial natural language processing system, by making the word embeddings more sensitive to the semantic meanings that humans actually perceive.

Acknowledgments

References

Aguirre-Celis, N., Miikkulainen R. (2017). From Words to Sentences & Back: Characterizing Context-dependent Meaning Rep in the Brain. Proc.39th Annual Meeting of the Cognitive Science Society, London, UK. 1513-1518.

Aguirre-Celis N., Miikkulainen R. (2018) Combining fMRI Data and Neural Networks to Quantify Contextual Effects in the Brain. In: Wang S. et al. (Eds.). Brain Informatics. BI 2018. Lecture Notes in Computer Science. 11309, 129-140. Springer, Cham.

Anderson, A. J., Binder, J. R., Fernandino, L., Humpries C. J., Conant L. L., Aguilar M., Wang X., Doko, S., Raizada, R. D. (2016). Perdicting Neural activity patterns associated with sentences using neurobiologically

motivated model of semantic representation. *Cer. Cortex*. 1-17. DOI:10.1093/cercor/bhw240.

Baroni, M., Murphi, B., Barbu, E., Poesio, M. (2010). Strudel: A Corpus-Based Semantic Model Based on Properties and Types. *Cognitive Science, 34*(2), 222-254.

Barsalou, L. W. (2008). Grounded Cognition. *Annual Review of Psychology, 59*, 617-845.

Binder, J. R., Desai, R. H., Graves, W. W., Conant L. L. (2009). Where is the semantic system? A critical review and meta-analysis of 120 functional neuroimaging studies. *Cerebral Cortex*, 19:2767-2769.

Binder, J. R., Desai, R. H. (2011). The neurobiology of semantic memory. *Trends Cognitive Sci*, 15(11):527-536.

Binder, J. R., Conant L. L., Humpries C. J., Fernandino L., Simons S., Aguilar M., Desai R. (2016a). Toward a brain-based componential semantic representation. *Cog. Neuropsychology, 33*:(3-4), 130-174.

Binder, J. R. (2016b). In defense of abstract conceptual representations. *Psychonomic Bulletin & Review*, 23.

Burgess, C. (1998). From simple associations to the building blocks of language: Modeling meaning in memory with HAL model. *Behavior Research Methods, Inst. & Com., 30*, 188-198.

Cox, R.W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Comput Biomed Res., 29*:162-173.

Gagné, R., Shoben, E. (1997). Influence of Thematic Relations on the Comprehension of Modifier-Noun Combinations. *Journal of Experimental Psychology Learning Memory and Cognition, 23*, 71-87 DOI: 10.1037/0278-7393.23.1.71

Glasgow, K., Roos, M., Haufler, A. J., Chevillet, M., A., Wolmetz, M. (2016). *Evaluating semantic models with word-sentence relatedness*. arXiv:1603.07253.

Harris, Z. (1970). Distributional Structure. *In Papers in Structure and Transformational Linguistics, 775*-794.

Just, M.A., Wang, J., & Cherkassky, V. (2017). Neural representations of the concepts in simple sentences: Concept activation prediction and context effects. *NeuroImage, 157*, 511-520.

Kiefer M, Pulvermüller F. (2012). Conceptual representations in mind and brain: theoretical developments, current evidence and future directions. *Cortex*. 48, 805–825.

Mitchell J, Lapata M. (2010). Composition in distributional models of semantics. *Cogn Sci*. 38(8), 1388–1439. DOI: 10.1111/j.1551-6709.2010.01106.x

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.

Miikkulainen, R., Dyer, M., G. (1991). Natural Language Processing with Modular PDP Networks and Distributed Lexicon. *Cognitive Science*, 15, 343-399.

Murphy, G. (1990). Noun Phrase Interpretation and Conceptual Combination. *Journal of Memory and Language*, 29, 259-288.

Pecher, D., Zeelenberg, R., Barsalou, L. W. (2004). Sensorimotors simulations underlie conceptual representations: Modality-specific effects of prior activation. *Psychonomic Bulletin & Review, 11*, 164-167.

Wisniewski, E., Property Instantiation in Conceptual Combination. (1998). *Memory & Cognition*, 26, 1330-1347. https://doi.org/10.3758/BF03201205