2016

The Coherence of Conceptual Primitives

Jamie C. Macbeth
Fairfield University, jmacbeth@smith.edu

Marydjina Barionnette
Fairfield University

Advances in Cognitive Systems 4 (2016) Submitted 4/2016; published 6/2016

The Coherence of Conceptual Primitives

Jamie C. Macbeth JMACBETH@FAIRFIELD.EDU
Marydjina Barionnette MBARIONNETTE@FAIRFIELD.EDU
Department of Electrical and Computer Systems Engineering, Fairfield University, Fairfield, CT 06824 USA

Abstract

Intelligent systems that process natural language need representations of knowledge to support a human-like thought process, and they often use natural language words or phrases to name and represent concepts in a knowledge base. But some theories of cognition claim that language and thought are not the same thing, and that human thought processes occur at a deeper level of representation than words and phrases in language. In this paper we present results of a human subjects study of language-free primitive decomposition as a representation for commonsense knowledge. We found that our subjects could comprehend and use a primitive decomposition representation; they demonstrated a facile understanding of the physical primitives from Conceptual Dependency, matching them reliably to sentences in ways that agreed with our expectations. Our results also show that the set of conceptual primitives we used resemble real human conceptualizations of natural language in ways that were sharp and coherent. Because our human subjects were recruited using a crowdsourcing platform, we claim that crowdsourcing may provide a vast and inexpensive source of conceptual structures based on primitive decomposition.

1. Introduction

Knowledge representation is a key issue for systems that perform natural language understanding, because an understanding program must perform inferences, fill in the gaps, and apply knowledge about commonsense social situations to understand texts in depth (Levesque et al., 2012). Knowledge bases conventionally use words or phrases in natural language to represent objects, concepts, and relations for a variety of intelligent system applications. While words might appear to be isomorphic to concepts, the conceptual structures that a human understander of language manipulates may not be identical to words or phrases in their language. Some theories of cognition and language understanding support the view that the thought process behind language understanding takes place in a "private" realm of "language of thought" that is very different from the spoken language of a human understander (Miller & Johnson-Laird, 1976; Fodor, 1975; Schank, 1975).

One way to build a non-linguistic language of thought is to decompose concepts into complex structures built from simple primitive elements that reflect the representations of acts, events, and states of the world that humans reason with, side-stepping the ambiguity of "surface" linguistic forms. But primitive decomposition systems are not without their own challenges. They may require even more careful hands-on design effort, since they use obscure language-free primitives instead of words. Also, when they are used for natural language understanding and generation, they require complex lexical resources to translate from natural language to the language-free representation and back. More recently, however, the potential for using primitive decomposition methods to build knowledge bases is encouraging, given the astounding increases in compute power and storage capacity over the years, as well as the availability of human computation resources that may help to build these systems inexpensively and accurately with contributions from many people.

This paper addresses a number of scientific questions and engineering challenges regarding primitive decomposition through an empirical study with human subjects. We determined that human subjects can comprehend conceptual primitives, and perform simple tasks of matching descriptions of the conceptual primitives to sentences in natural language. The key empirical tool that we used to measure success is the inter-rater reliability of our subjects' answers; we found nearly perfect levels of inter-rater reliability overall, indicating that the primitives were coherent and well-defined in the minds of our subjects collectively. By asking our subjects to explain their answers to the conceptual primitive matching task, we collected evidence to address what the "right" set of primitives are, whether the primitives reflect the way people think, and whether human subjects research can suggest improvements to primitive decomposition systems. Finally, the success of our study signals the feasibility of using crowdsourcing techniques to collect commonsense knowledge in the form of primitive decompositions from people with little knowledge or expertise in building intelligent systems.

2. Background

Cyc and ConceptNet (Lenat & Guha, 1990; Speer & Havasi, 2013) are examples of large commonsense knowledge bases containing thousands to millions of facts, assertions, and relations structured in a way to support implementation of intelligent systems. Both have been applied to natural language processing tasks such as word-sense disambiguation (Havasi et al., 2010; Barbella & Forbus, 2013). In these systems, concepts are often represented as words or phrases in natural language and connected through logical, taxonomic, or ontological relations.

However, a body of cognitive psychology research suggests that not all human reasoning arises from rigid ontologies and logical relations (Rosch, 1975; Wason & Johnson-Laird, 1972). For natural language-understanding tasks, these knowledge bases may be enhanced by supporting sources of knowledge that represent decompositions of natural language expressions into primitive constructions. Wilks' Preference Semantics (Wilks & Fass, 1992), Schank's Conceptual Dependency (Schank, 1972), and Jackendoff's Lexical Conceptual Structures (Jackendoff, 1983) are all examples of semantic primitive systems. These systems attempt to build structures that would reflect the pictures and representations in people's minds about real-world acts and events that happen, as well as the changes in the state of the world that result from them. According to these theories of cognition regarding language, natural language expressions are translated to conceptual structures in understanding, and translated from conceptual structures back to language in generation. Their virtues and drawbacks were discussed vigorously in their time (Wilks, 1996; Winograd, 1978).

One of the chief benefits of decomposing concepts into primitives is that, particularly in language applications, it allows concepts to be related to other concepts directly, rather than through syntactic or logical transformations of language expressions onto other language expressions. For our study of human subjects we chose Schank's Conceptual Dependency (CD) system because it intends to be a language-free representation, because it has relatively few primitives, and because of its successful application to natural language understanding systems over several decades (Schank & Abelson, 1977; Schank & Riesbeck, 1982; Lytinen, 1992). We discuss CD further in the methodology section of this paper.

3. Research Questions

Our interest in conceptual primitives raised a number of interconnected questions that we sought to address in our human-subjects study. As a precursor to even conducting a human-subjects study on conceptual primitives, one question we asked was this: Can people comprehend the conceptual primitives without any prior introduction to Conceptual Dependency theory or any other theory of primitive decomposition of language? Assuming our human subjects understood the primitives, could they identify their presence in a conceptualization of a natural language sentence or utterance?

Our human subjects study also addresses questions related to the realism and validity of primitive decomposition systems. We wondered if representations based on primitive decompositions reflect the ways people actually conceptualize events and process natural language. Related to this, we were curious about implications our study would have for primitive decomposition system design: Could the primitives be generalized or simplified? Do the primitives complement each other? Are there primitives that overlap or cases where one primitive could be subsumed into another?

Our main empirical method is to present human subjects with descriptions of the primitives and simple sentences and have the human subjects match the sentences with the primitives. To gather quantitative evidence toward answering these questions we sample a population of language users and see how much they agreed when performing the matching task. High levels of agreement would indicate that there is a coherence to the "conceptual signal" indexed by the primitives in the minds of our subjects. This raised the following further questions about the behavior of a sample: For a given sentence, did a majority of subjects match a particular primitive to that sentence? Were there some sentences for which more than one primitive was popular? Quantitative analysis of the sample will help address our more general questions, as agreement serves as a proxy indicator that subjects understood the primitives and could reliably locate a particular primitive in their conceptualization of a sentence.

Finally, because we use a crowdsourcing platform to recruit human subjects for our study, the results will have obvious implications for efforts to collect knowledge based on primitive decomposition of conceptualizations through human computation.

4. Methodology

To investigate these questions, we designed and performed a formative human-subjects study focusing solely on matching descriptions of conceptual primitives to sentences in natural language. Because we assumed subjects would have no knowledge of conceptual primitives, we had them perform the matching as a "categorization" task. First, we gave the subjects descriptions of the conceptual primitives, calling them "categories". Then we presented subjects with a set of very short, simple sentences, which we called *target sentences*, describing simple acts, each of which corresponded to one of the conceptual primitives. We asked subjects to read the short sentences and to choose which category was a general match for the "main action" in each sentence. Finally, we instructed subjects to provide a one-sentence explanation of each of their answers.

The set of conceptual primitives we chose for this study were the six physical primitives provided by Conceptual Dependency (Schank, 1972, 1975), which are called PTRANS, MOVE, GRASP, PROPEL, INGEST, and EXPEL. Conceptual Dependency has several other primitives that deal with transfer of information, thought, and abstract transfer of possession of objects; CD also has several mechanisms for connecting primitives into larger conceptual structures. However, the physical primitives were a sufficiently large and complementary set and gave us a simplicity and focus that we desired in this formative study.

The descriptions of the primitives that we presented to subjects are in Table 1. We made our descriptions of the primitives as general as possible. Descriptions of the primitives did not include the names of the primitives; this agrees with the spirit of CD as a language-free representation, and made it so that subjects would not be biased to making a selection based on primitive names. The descriptions of the primitives were numbered 1 through 6, and subjects matched a primitive to each sentence by mapping each sentence to the numbered category. However, for clarity in describing the task, we have shown both the numbers and the names of the primitives in Table 1.

After reading the category descriptions, subjects read a sequence of 12 sentences, and were asked to place each sentence in one of the categories. We constructed the 12 sentences to be reminiscent of those in Schank (1975), and for each of the 6 physical primitives, two sentences corresponding to that primitive were included in the set. The sentences were presented in the randomized order shown in Table 2.

Having a pair of sentences for each primitive gave us the opportunity to test the primitive in a broad generality of surface expressions, which occurred in the following ways: The majority of actors and objects would appear to be human based on the common stereotypical male and female first names, but "the gecko" and "a car" were non-human actors; the two sentences for INGEST described the typical act of eating associated with INGEST, as well as inhalation; one of the two sentences for EXPEL described a person vomiting or regurgitating food, while the other described a person bleeding; one of the PTRANS sentences used the verb "to return" instead of a typical verb such as "to go"; and "the gecko stuck to the wall" used an "attach" meaning of GRASP. We made sure that the sentences used different action words than those in the primitive descriptions, so that subjects would not simply perform a word matching.

The study was administered via Amazon Mechanical Turk[1], an Internet croudsourcing marketplace. We recruited 50 subjects for the study whom we assumed had no knowledge of or experience with primitive decomposition systems. Each subject was paid $1, and we estimated that each subject would take 15 minutes to complete the task.

1. http://www.mturk.com

Table 1. The descriptions of the conceptual primitives that subjects in our study were given, described to subjects as "categories". So as not to bias the subjects, the names of the primitives were not given, only numbered descriptions. However, primitive names are shown here for clarity. Our subjects were asked to give numbers of categories in their responses.

(1) PTRANS A person, object, or thing changes physical position or location.	**(2) MOVE** A person, object, or thing moves a part of its body or part of itself.
(3) GRASP A person, object, or thing grabs hold of another person, object, or thing, or becomes attached to another person, object, or thing.	**(4) PROPEL** A person, object, or thing applies a force to another person, object, or thing, or a moving person, object, or thing strikes or impacts another person, object, or thing.
(5) EXPEL A person, object, or thing, is taken from or comes from inside another person, object, or thing and is forced out.	**(6) INGEST** A person, object, or thing is forced (or forces itself) to go inside of another person, object, or thing.

5. Results

5.1 Agreement and Inter-rater Reliability

We were interested to know, in quantitative terms, whether the primitives that our subjects conceptualized when reading the sentences were the same primitives we conceptualized when crafting them. For each of the 12 target sentences, we calculated the percentage of the 50 subjects who answered with the primitive we expected. The percentage agreement with our expectations for each sentence is shown in Table 2. The overall agreement with our expectations for these 600 items was 91.3% (548 out of 600).

Table 2. The 12 target sentences, corresponding expected primitives, and subjects' percentage agreement with the expected primitive. Sentences are shown in the randomized order in which they appeared in the study.

Sentence	Expected Primitive	Percentage Agreement
"Jim held on to the railing."	GRASP	98%
"Lisa kicked the ball."	PROPEL	98%
"Matthew flew home from Los Angeles."	PTRANS	92%
"Michelle threw up her lunch."	EXPEL	92%
"The gecko stuck to the wall."	GRASP	92%
"Robert returned home from downtown."	PTRANS	96%
"Charles ate a hamburger."	INGEST	96%
"Bill was hit by a car."	PROPEL	100%
"Amy took a deep breath."	INGEST	76%
"Stephanie bled from a cut on her leg"	EXPEL	92%
"Kevin crossed his arms."	MOVE	92%
"Joe swung his fist at David."	MOVE	72%
All sentences		91.3%

We also calculated inter-rater reliability measures to assess whether the subjects agreed with each other regardless of whether they agreed with our preconception of the "right" primitive for a sentence. Fleiss' kappa (κ) is a statistical measure of inter-rater reliability when more than two raters or annotators are rating multiple data items (Fleiss, 1971). We calculated Fleiss' κ for the 50 subjects as raters, and the sentences as 12 data items with the primitives as ratings. Fleiss' κ came to 0.818, which is generally accepted as "almost perfect" agreement (Landis & Koch, 1977). These results indicate that, at least for the sentences we created, there is a strong coherence to the "conceptual signal" corresponding to the CD physical primitives.

5.2 Explanations and Expected Answers

We were fortunate to have subjects explain their choice of primitive category for each sentence. One reason we made this a requirement was that it would force human subjects to provide some evidence that they had not simply rushed through the task, answering with primitive categories randomly. Because providing a one-sentence explanation is part of the task, and because we informed subjects that their Mechanical Turk submissions would be rejected if they do not complete all parts of the task, we believed subjects would think carefully about their answers. More importantly, however, forcing subjects to explain their answers, in addition to telling us that they were thinking at all,

Table 3. Examples of subjects' answer pairs—conceptual primitives and explanations of how the conceptual primitive fit the target sentence—organized by primitive.

Primitive	Example Explanations
GRASP	"Jim's hand makes firm contact with the railing."
	"The foot pads on the gecko had a grip on the wall."
EXPEL	"The contents of Michelle's stomach came from within her."
	"Stephanie's blood is moving from inside her leg to outside her leg."
INGEST	"Charles eating a hamburger is taking something that is outside his body and putting it inside of him."
	"Breathing in would be forcing something inside, since it is forcing air in."
MOVE	"Joe moved a body part, in this case his hand and arm and fist."
	"This describes that Kevin has moved a part of his body by crossing his arms."
PROPEL	"Lisa used the force of her foot and leg to impact the ball."
	"A person (Bill) was impacted by an object (the car)."
PTRANS	"Matthew flying home is him changing his location."
	"Robert changed location to his home from downtown."

gave us insights into how they were thinking, and whether or not they grasped the primitives and connected them to their conceptualizations of the acts described in the sentences.

Examples of the explanations that subjects gave for the sentences are shown in Table 3, where they are organized by CD conceptual primitive. We wanted to verify that subjects achieved a high level of agreement with our expectations by examining their explanations. Also, we realized that subjects did not always answer the way we expected, and their explanations would reveal why. In particular, we recognized that the sentences described complex events, and many of our subjects would have reasonably conceptualized other primitive acts as the main act in their conceptualization of the target sentence.

We sought to determine whether subjects' categorizations of sentences were reasonable with respect to their explanations, so we coded the explanations as being reasonable or unreasonable independently of whether they gave the expected primitive as an answer. Two researchers read each primitive category answer and each explanation to see if they made sense together—checking to see if the subjects' explanation linked the primitive act description, an actor, and an object in a way that the researchers could conceptualize as being part of the complex act described in the sentence. The researchers coded these answer pairs as being reasonable or unreasonable based on these criteria. Subjects often gave a primitive answer that disagreed with the expected primitive, but gave a sensible explanation that showed that that primitive was part of a reasonable conceptualization of the target sentence.

The two researchers coded all 600 data items, and only disagreed on nine of them, achieving 98.5% agreement. There were only three items (0.5 percent of cases) where one or the other of the researchers coded the explanation as unreasonable; therefore 99.5% of the time subjects were giving "reasonable" answers, an even higher percentage than the cases where the primitive categorization agreed with expectations.

The researchers who coded the explanations only found three data items that were unreasonable. In one case, for the sentence "Robert returned home from downtown," we expected PTRANS, but one subject answered EXPEL and explained:

> "Robert is taken from downtown and forced to go home."

The researchers disagreed with the idea that downtown was an object that Robert would be taken from the inside to the outside of, and that home was "the outside". In another case, the subjects explanations failed to give enough information to indicate their conceptualization of the sentence in connection to the primitive. For example one subject gave the explanation, "I think that fits very well," which contained no descriptive value relevant to the conceptualization of the meaning of the sentence.

5.3 Implications of Unexpected Answers

In cases where subjects did not answer with the expected categorization, we have qualitative data that tells us that they were still "making sense" and demonstrating useful commonsense knowledge about the situation. Subjects disagreed with our expectations in 52 cases out of 600 (8.6%). While coding the explanations, we learned that occasionally a subject gave an explanation that agreed strongly with the expected category but unwittingly marked a different category. There were 13 items of this kind. If we disregard 1 item with an unexpected primitive for which both researchers marked the explanation as "unreasonable", this left 38 items where subjects gave well-reasoned explanations that a primitive that we did not expect was part of their conceptualization. These unexpected answers, which were often given by multiple subjects and formed patterns, demonstrated novel applications of the primitives to parts of the conceptualization of a complex act, or demonstrated how subjects were drawing inferences about events which were not explicitly stated in the sentence. We describe examples of these cases in this section.

5.3.1 "Matthew..."

Occasionally we found that subjects' disagreeable answers opened our eyes to applications of a primitive to new situations. The INGEST primitive of CD has traditionally been used when some object, either animate or inanimate, goes inside an animate object. Typical examples of the INGEST primitive are used to describe humans or other animate objects and their acts of consuming food and drink, breathing, or inhaling smoke (Schank, 1975). For the sentence "Matthew flew home from Los

Angeles," we expected the PTRANS primitive. However, two subjects answered with the INGEST primitive. They explained:

> "Matthew forces himself inside an airplane to get home."
>
> "Matthew got into the plane that flew him home from LA."

In their explanations of these answers, the subjects recognized that Matthew would get inside the plane, effectively arguing that the plane would be INGESTing Matthew, generalizing the meaning of INGEST to an animate object going inside of an *inanimate* object. This generalization was partially in our making, since our description of INGEST is not specific about who or what is causing the movement to the inside: "A person, object, or thing is forced (or forces itself)...". Our description is also not specific about whether the destination of the movement is an animate object or not, stating "...to go inside of another person, object, or thing."

5.3.2 "Amy..."

For the sentence "Amy took a deep breath," we expected subjects to answer INGEST, but many answered MOVE. Here are examples of explanations from subjects who answered this way:

> "She moved her chest muscles to inflate her lungs."
>
> "Amy taking a deep breath is moving her body so as to breathe."
>
> "Taking a deep breath caused her body to move."

It appears that when subjects decomposed the act, many of them focused on one primitive and many of them focused on another. For this sentence, 76% (38 out of 50) of subjects focused on the fact that the act of taking a breath involved taking what people conceptualized as a substance, air, inside the body, and coded the sentence as an INGEST. But 20% (10 out of 50) focused on the fact that, as part of breathing, Amy moves part of her body—her diaphragm, or her lungs—and categorized the main act as a MOVE. This was a clear case where conceptualizing the complex act in the target sentence in terms of the primitives required multiple primitives to be involved in the conceptualization.

5.3.3 "Joe..."

For the target sentence "Joe swung his fist at David," only 72% (36 out of 50) of our subjects answered MOVE, the primitive we expected. 13 subjects (26%) answered PROPEL, and even though it was not explicitly stated that Joe actually struck David, PROPELing David was part of their conceptualizations. Some of the subjects' explanations illustrate conceiving Joe's intent to hit David:

> "Joe *tried* to apply force with his fist to David."
>
> "Swinging your fist is applying force with *intent* to injure."

While some explanations of PROPEL made it sound as if Joe definitely hit David:

"Joe applied force to David."

"Joe's hand impacted David."

It is somewhat unsurprising that, as part of their conceptualizations, subjects made inferences about events that occurred in these sentences, similar to the inferences they might make in reading stories. Subjects inferred Joe's intent, or that Joe actually hit David, and their explanations were reasonable to the researchers when they coded the primitive-explanation answer pairs.

6. Discussion

Subjects in our study achieved a high level of agreement with the researchers on matching the primitives with the sentences, as well as high inter-rater reliability, and reasonableness of their explanations. This was clear evidence that human subjects were able to comprehend the primitives in ways we expected, and that they were generally capable of performing the task of matching primitives with sentences. Although the primitives may not completely capture our human subjects' thought processes, this is evidence that the conceptual primitives are at least a strong reflection of the way that our human subjects think when comprehending the sentences. That we ran our study entirely through Amazon Mechanical Turk, and that we were successful—in spite of subjects' lack of previous knowledge of Conceptual Dependency or other primitive decomposition systems—together indicate the feasibility of using crowdsourcing or other forms of human computation to build knowledge bases of this form.

We recognized from the outset of our study that having subjects decompose their conceptualizations about complex acts would naturally require multiple primitive acts to be part of the conceptualization. This was certainly true for the sentence "Amy took a deep breath," where Amy likely MOVEs her chest muscles to expand her lungs, causing air to be INGESTed into her lungs. But it is also true for "Joe swung his fist at David," where Joe MOVEs his fist as part of his intent to PROPEL David, and for "Matthew flew home from Los Angeles," where subjects recognized that Matthew would PTRANS and change physical location, but also that he would have to get inside the plane (which was categorized as INGEST). In these and in other cases the researchers found subjects' unexpected primitives to be reasonable because they would have combined with the expected primitives in the conceptualization of the complex act broadly. These unexpected primitives were either integral to the explicit meaning of the sentence, or they were primitive acts that were inferred to have happened, as was the case with "Joe swung his fist at David," where subjects inferred that striking or PROPELing David was a goal. This leads us to believe that subjects' answers may have been even more consistent had they not been constrained to answering with only a single conceptual primitive.

One of the stated goals of Conceptual Dependency, which differentiates it from other primitive decomposition systems, was to establish a small set of conceptual primitives that could represent a large number of concepts by combining with each other in many complex ways (Schank & Riesbeck, 1982). Our study gathered direct evidence from our human subjects that may lead to simplifications or improvements to CD that align with these goals. We saw that for "Matthew flew home from Los Angeles," subjects generalized the INGEST primitive to apply to objects entering other objects, whether animate or inanimate. For the case where a subject answered EXPEL (the opposite of

INGEST) for the sentence "Robert returned home from downtown," we recognize a potential generalization of INGEST and EXPEL which goes even further, meaning that an object changes location and crosses a boundary between one conceptualizable space, "downtown," to another, "home". Further study and experimentation will shed light on whether such an extension of INGEST/EXPEL (for which INGEST and EXPEL cease to be good descriptive names) may be useful, or if it is an overgeneralization.

7. Conclusion and Future Work

In this paper we presented an empirical study with human subjects that we conducted to address a number of questions and challenges that are raised in building primitive decomposition systems for knowledge representation. We found that our human subjects could work within a language-free primitive decomposition representation; they demonstrated a facile understanding of the physical primitives from Conceptual Dependency, matching them reliably to sentences in ways that agreed with our expectations. Our results also showed that the conceptual primitives resembled real human conceptualizations of natural language, and that the primitives that we used were complementary and coherent. We also provided strong evidence that crowdsourcing may provide a vast and inexpensive source of conceptual structures based on primitive decomposition.

Future work will examine human subjects working with more complex sentences, and with other, more abstract conceptual primitives that are relevant to events involving human perception, thought, communication, planning, and social interactions. Subjects recruited via crowdsourcing platforms may also build more sophisticated conceptual structures involving the actors and objects of primitive acts, as well as complex combinations of primitives—for example, where one primitive act causes another, or when one primitive act is instrumental to another. Looking forward, we wish to collect these structures in large numbers to act as a semantic memories for understanding systems, and to develop methods to augment conventional knowledge bases with them.

References

Barbella, D., & Forbus, K. D. (2013). Analogical word sense disambiguation. *Advances in Cognitive Systems*, *2*, 297–315.

Fleiss, J. L. (1971). Measuring nominal scale agreement among many raters. *Psychological Bulletin*, *76*, 378.

Fodor, J. A. (1975). *The Language of Thought*. Cambridge, MA: Harvard University Press.

Havasi, C., Speer, R., & Pustejovsky, J. (2010). Coarse word-sense disambiguation using common sense. *AAAI Fall Symposium: Commonsense Knowledge*.

Jackendoff, R. S. (1983). *Semantics and Cognition*, volume 8 of *Current Studies in Linguistics*. Cambridge, MA: MIT Press.

Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics*, *33*, 159–174.

Lenat, D. B., & Guha, R. V. (1990). *Building Large Knowledge-Based Systems*. Reading, MA: Addison-Wesley.

Levesque, H., Davis, E., & Morgenstern, L. (2012). The winograd schema challenge. *13th International Conference on the Principles of Knowledge Representation and Reasoning, KR 2012* (pp. 552–561).

Lytinen, S. L. (1992). Conceptual dependency and its descendants. *Computers & Mathematics with Applications*, *23*, 51–73.

Miller, G. A., & Johnson-Laird, P. N. (1976). *Language and Perception*. Cambridge, MA: Belknap Press.

Rosch, E. (1975). Cognitive representations of semantic categories. *Journal of Experimental Psychology: General*, *104*, 192.

Schank, R. C. (1972). Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology*, *3*, 552–631.

Schank, R. C. (1975). *Conceptual Information Processing*. New York, NY: Elsevier.

Schank, R. C., & Abelson, R. P. (1977). *Scripts, plans, goals and understanding : an inquiry into human knowledge structures*. Hillsdale, N.J.: L. Erlbaum Associates.

Schank, R. C., & Riesbeck, C. K. (1982). *Inside Computer Understanding: Five Programs Plus Miniatures*. Hillsdale, NJ: L. Erlbaum Associates Inc.

Speer, R., & Havasi, C. (2013). Conceptnet 5: A large semantic network for relational knowledge. In *The PeopleâĂŹs Web Meets NLP*, (pp. 161–176). Springer.

Wason, P. C., & Johnson-Laird, P. N. (1972). *Psychology of Reasoning: Structure and Content*, volume 86. Harvard University Press.

Wilks, Y. (1996). *Good and Bad Arguments for Semantic Primitives*. Technical report, Computing Research Laboratory, New Mexico State University, Las Cruces, New Mexico,.

Wilks, Y., & Fass, D. (1992). The preference semantics family. *Computers & Mathematics with Applications*, *23*, 205–221.

Winograd, T. (1978). On primitives, prototypes, and other semantic anomalies. *Proceedings of the 1978 Workshop on Theoretical Issues in Natural Language Processing* (pp. 25–32). Stroudsburg, PA, USA: Association for Computational Linguistics.