

Conceptual Primitive Decomposition for Knowledge Sharing via Natural Language

Jamie C. MACBETH^{a,1},

^a*Department of Electrical and Computer Systems Engineering,
Fairfield University, USA*

Abstract. Natural language is an ideal mode of interaction and knowledge sharing between intelligent computer systems and their human users. But a major problem that natural language interaction poses is linguistic variation, or the “paraphrase problem”: there are a variety of ways of referring to the same idea. This is a special problem for intelligent systems in domains such as information retrieval, where a query presented in natural language is matched against an ontology or knowledge base, particularly when its representation uses a vocabulary based in natural language. This paper proposes solutions to these problems in primitive decomposition methods that represent concepts in terms of structures reflecting low-level, embodied human cognition. We argue that this type of representation system engenders richer relations between natural language expressions and knowledge structures, enabling more effective interactive knowledge sharing.

Keywords. Natural language interaction, primitive decomposition, Conceptual Dependency, natural language understanding, knowledge sharing.

1. Introduction

Natural language promises to be the most effective and convenient method of transmitting and sharing of knowledge between humans and intelligent computing systems, and building computing systems that understand natural language as input is an essential part to having to natural language as a mode of interaction. While the application of machine learning and deep learning methods to large collections of natural language text have lead to major successes in natural language processing (NLP), many have noted that the space of applications that are targeted by these efforts misses out on the broad range of capabilities that humans exhibit in natural language [8,6].

Ontologies and knowledge bases will play multiple roles in creating interactive computing systems that work with natural language. Firstly, to approach applications like question answering, information retrieval, and narrative understanding, natural language understanding systems will need to go beyond surface levels of language to access and apply ontological world knowledge so that they can make inferences and deductions and draw conclusions similar to those that humans can draw [15,11]. Secondly, ontologies

¹Corresponding Author: Jamie C. Macbeth, Department of Electrical and Computer Systems Engineering, Fairfield University, Fairfield Connecticut, 06824, USA; E-mail: jmacbeth@fairfield.edu.

can be used to represent the knowledge that is shared interactively via the natural language understanding system.

But even when knowledge bases and ontological resources are available, a major challenge of knowledge sharing in the natural language understanding domain is the “paraphrase problem,” also known in linguistics by the term *linguistic variation*: the fact that there are many different ways of saying the same thing [14]. While this term applies to cross-linguistic variation, or variation between different languages, it also applies to cases of “free variation” in a single language, where differences in form do not correspond to any difference in meaning [14, pp. 7–8]. This is a crucial issue for environments and intelligent systems that share knowledge through interactions in natural language, because the knowledge about a particular thing may be referred to in many different ways.

Some ontologies and knowledge base systems (such as Cyc [7]) may also fall prone to these problems, because the knowledge about a particular thing is often indexed by lexical entries, words or phrases, related to that thing. This brief paper proposes research in a particular class of solutions to paraphrase and linguistic variation problems in natural language understanding and ontology representations. By formulating a representation that avoids lexical items and represents knowledge and meaning by decomposing them into structures that reflect primitive, low-level components of embodied human perception and cognition, it may enable rich relations between lexical items, varieties of linguistic expressions, and ontological structures that make for more effective and interactive natural language understanding systems.

2. The “Paraphrase Problem” and Linguistic Variation

While a significant literature addresses the challenge of generating paraphrases [9], detecting paraphrases in the face of linguistic variation serves to be a more challenging problem, because it requires understanding the language in each paraphrase. To illustrate, we take the following example problem in story matching and story understanding in which the system processes stories to make a determination about whether they describe similar events:

- (1) Bob bought a loaf of bread from the store for \$1.
- (2) The store sold a loaf of bread to Bob for \$1.

An information retrieval system or a story search and matching system might take (1) as input, and, given a collection of stories would need to select (2) from the story collection as a paraphrase of (1). Alternatively, the problem could be posed as a question answering problem, where the system is presented with the story (1) as input along with the question:

- (3) Who sold the loaf of bread to Bob?

Determining that (1) describes the same event or set of events as (2) depends crucially on the system understanding that the same sequence or collection of acts occurred. An understanding system based naively on analyzing the sentence syntactically and determining that the verb in (1) is “buy” and the verb in (2) is “sell” might determine that the sentences did not describe the same events because none of the verbs matched.

A more advanced natural language understanding system might search a knowledge base or an ontology for propositions stored under the verb “buy” and for relations between the word “buy” and the word “sell”. If the knowledge base is represented as an ontology or a semantic network, there may be one or more links or a path consisting of a series of links in the network in between buying and selling concepts to represent that they are related.

Continuing the example, an ontology system such as Text2Onto [2] could be used to extract ontological concepts and relations from textual and linguistic resources to build a story understanding system. Text2Onto could be run on a lexical database like WordNet [3], which is frequently used for text analysis and AI applications. WordNet organizes an English lexicon primarily into *synonym sets* or *synsets*—groupings of words “that are interchangeable in some context.” These relations include those such as synonymy/antonymy, hypernymy/hyponymy, meronymy/holonymy and entailment. WordNet and lexical resources like it should have a synset relation between “buy” and “sell”.

But synset relations alone do not appear to capture the relationships of the actors and both direct and indirect objects to the events in (1) and (2). One assumes that a synonym relation between words implies that one can be replaced with the other in the correct context. Using the synonym relation to test the story-similarity of (1) and (2) by replacing “sell” with “buy” in (2) yields:

- (4) The store bought a loaf of bread to Bob for \$1.

which is obviously not a paraphrase for (1) or (2) and does not lead to the conclusion that (1) is a paraphrase of (2).

One might think that an antonym relation is more accurate. Since an antonym relation implies that “buy” is the opposite act of “sell”, testing that (2) is a paraphrase of (1) would require a more complex transformation of the sentence, replacing “sell” with “buy” and switching the actor and the indirect object, and switching the preposition “to” to “from”. This more complex transformation appears to work in mapping (1) to (2). However, a third story:

- (5) Bob gave the store \$1, and the store gave Bob a loaf of bread.

appears to be a paraphrase of both (1) and (2), but our synset method will definitely not work. While it is clear that the act of “giving” is an integral part of buying and selling, none of the commonly used linguistic synset relations seems adequate in capturing the relationship. For example, if we try to represent that giving is synonymous with selling, a transformation of the selling sentence might produce a reasonable-sounding paraphrase such as

- (6) The store gave a loaf of bread to Bob for \$1.

But it would also produce a number of clearly incorrect paraphrases for matching and understanding such as these:

- (7) Bob gave a loaf of bread from the store for \$1.
- (8) Bob bought the store \$1, and the store bought Bob a loaf of bread.
- (9) Bob sold the store \$1, and the store sold Bob a loaf of bread.
- (10) Bob bought the store \$1, and the store sold Bob a loaf of bread.

(11) Bob sold the store \$1, and the store bought Bob a loaf of bread.

FrameNet [1] is another resource that is frequently used in natural language understanding systems which, at a glance, may appear to solve this problem. FrameNet links word senses to each other in *frames* [4] that are meant to represent the complex relationships between nouns and verbs in natural language expressions describing a particular social situation. FrameNet has a `Commercial_transaction` frame which has semantic and syntactic valence relationships between verbs such as “buy”, “sell”, “charge”, “spend”, “pay”, and “cost”, and nouns such as “goods” and “money”, as well as relationships between actors and a `Commercial_transaction` event, such as “buyer” and “seller”.

However, when we perform a close examination of the FrameNet frames² for `Commerce_buy`, `Commerce_sell`, and `Giving`, the `Commercial_transaction` subframes that one might naturally consult in the context of our story matching problem, we encounter similar issues when we try to generate a relationship between several different ways of describing the same set of events. The `Giving` frame in FrameNet describes the situation using information about the syntactic and semantic relations between words like `Donor`, `Recipient`, and `Theme`, while the `Commerce_buy` and `Commerce_sell` frames use a different set of terms: `Buyer`, `Seller`, `Money`, and `Goods`. Although FrameNet does define an inheritance relationship between these three frames in which both the `Commerce_buy` and `Commerce_sell` frames inherit from the `Giving` frame, no other annotations in FrameNet exist to indicate, for example, that buying or selling something involves two acts of giving, or that the donor in the giving act where money was the theme is the buyer.

In building automated systems to understand and relate natural language expressions about similar events through FrameNet frames, we will encounter problems similar to those we encountered using WordNet synsets. Ultimately, this appears to be because the content of FrameNet frames is composed of linguistic forms and lexical items which are subject to the same paraphrasing and linguistic variation that are inherent in natural language generally.

3. Solutions in Cognitive Primitive Decomposition

The preceding examples demonstrate a number of issues with building natural language understanding systems based on lexical relations and ontologies built from text resources, given the “paraphrase problem”. In cases where linguistic knowledge bases have established relations such as synonymy or antonymy, they may also establish rules for transforming, for example, stories about buying into stories about selling. However, scaling these systems may require a large number of transformation rules. In the case above, rules or relations are needed for switching instances of the pair of verbs “buy” and “sell”, for exchanging the subject with the indirect object. Even more rules are needed for all of the various permutations involving “give”, “take” and other similar verbs.

Similarly, while typical ontological relations—such as *is-a*, *instance-of*, *part-whole*, and *equivalence*—are very relevant to relations between certain types of objects and concepts, they do not appear to capture the complex conceptual relationships in ques-

²An index of FrameNet frames can be found at <https://framenet.icsi.berkeley.edu/fndrupal/frameIndex>

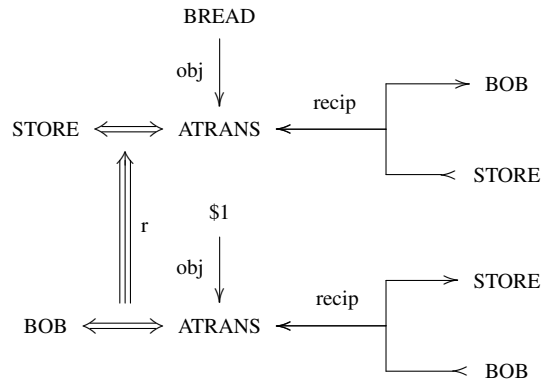


Figure 1. A Conceptual Dependency diagram representing the conceptualization for the three paraphrases: “Bob bought a loaf of bread from the store for \$1,” “the store sold a loaf of bread to Bob for \$1”, and “Bob gave the store a \$1, and the store gave Bob a loaf of bread.” Double arrows represent the relationship between the animate actor and the ATRANS primitive act, while single arrows marked “obj”, “recip”, indicate the object and the recipient case of the ATRANS act, respectively. The triple arrow marked “r” represents “result causation”, indicating that Bob’s act of ATRANSing \$1 to the store caused the store to ATRANS a loaf of bread to Bob.

tion here. Although one could devise unique labels for links between “buy”, “sell”, and “give”, or place intermediate nodes between them in an ontological network, this appears to proliferate labels of relations, just as it would transformation rules in a lexical database.

A potential solution to these problems is to work to represent meaning in a way that is not only non-linguistic, but reflects imagery, image schemas, or mental models [12,10,5] composed of the primitive perceptual and cognitive experiences that people use when processing and understanding language. Figure 1 shows a conceptual diagram meant to represent the concept behind the three paraphrases (1), (2), and (5) simultaneously. The conceptual diagram is composed in a system called Conceptual Dependency [13], which attempts to decompose the concept behind the paraphrases into a complex combination of language-free conceptual primitives; in this case using a conceptual primitive called ATRANS, which represents an event that changes an abstract relationship between human actors and an inanimate object. While the top ATRANS primitive act in the diagram represents a conceptualization of the store transferring possession of the bread to John, the bottom ATRANS primitive act represents bob transferring possession of the \$1 to the store.

A story understanding system tasked with relating the given paraphrases can attempt to decompose each the different sentences into this “conceptual base” form, and perform a comparison of the conceptual diagrams to determine that they describe the same set of events. If the representation system uses a small number of conceptual primitives and connectives, but combines them in complex ways to correspond to the variation of natural

language expressions, the matching occurs through relatively simple processes of graph isomorphism and structure mapping.

If the problem with synset relations and standard ontology relations is that they will proliferate a complex array of named relations between words or between concepts, we foresee that methods of decomposing words into complex combinations of primitives creates a substrate to capture the complex relations between words and other words, between words and concepts, and between concepts and other concepts. Ultimately, decomposing concepts in this way allows the decomposition into cognitive primitives itself to define the system of relations. This appears to be a more realistic representation of human cognition behind the relations between concepts than a simple system of labeled edges.

We propose further research in supplementing or enhancing available ontologies, knowledge bases, and lexical databases with resources that present knowledge in a form that is decomposed into structured combinations of conceptual primitives. We argue that this is essential to achieving richer and more in-depth forms of computer-based natural language understanding. This, in turn, may enable greater knowledge sharing between people and all systems that represent knowledge in the form of natural language.

References

- [1] Collin F Baker, Charles J Fillmore, and John B Lowe. The berkeley FrameNet project. In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics-Volume 1*, pages 86–90. Association for Computational Linguistics, 1998.
- [2] Philipp Cimiano and Johanna Völker. A framework for ontology learning and data-driven change discovery. In *Proceedings of the 10th International Conference on Applications of Natural Language to Information Systems (NLDB)*, volume 3513 of *Lecture Notes in Computer Science*, pages 227–238. Springer, 2005.
- [3] Christiane Fellbaum. *WordNet: An Electronic Lexical Database*. MIT Press, 1998.
- [4] Charles J Fillmore and Beryl T Atkins. Toward a frame-based lexicon: The semantics of RISK and its neighbors. In *Frames, Fields, and Contrasts: New Essays in Semantics and Lexical Organization*, pages 75–102. Lawrence Erlbaum Associates, Hillsdale, NJ, 1992.
- [5] Philip Nicholas Johnson-Laird. *Mental models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Harvard University Press, 1983.
- [6] Pat Langley. Intelligent behavior in humans and machines. *Advances in Cognitive Systems*, 2:3–12, 2012.
- [7] D. B. Lenat and R. V. Guha. *Building Large Knowledge-Based Systems*. Addison-Wesley, Reading, MA, 1990.
- [8] Hector J Levesque. On our best behaviour. *Artificial Intelligence*, 212:27–35, 2014.
- [9] Nitin Madnani and Bonnie J. Dorr. Generating phrasal and sentential paraphrases: A survey of data-driven methods. *Comput. Linguist.*, 36(3):341–387, September 2010.
- [10] Jean M Mandler and Cristóbal Pagán Cánovas. On defining image schemas. *Language and Cognition*, 6(4):510–532, 2014.
- [11] Vincent Ng. Machine learning for entity coreference resolution: A retrospective look at two decades of research. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 4877–4884. AAAI, 2017.
- [12] M. Sadoski and A. Paivio. *Imagery and Text: A Dual Coding Theory of Reading and Writing*. Taylor & Francis, New York, 2013.
- [13] Roger C Schank. *Conceptual Information Processing*. Elsevier, New York, NY, 1975.
- [14] James A Walker. *Variation in Linguistic Systems*. Routledge, New York, NY, 2010.
- [15] Patrick Henry Winston. The genesis story understanding and story telling system: A 21st century step toward artificial intelligence. Technical report, Center for Brains, Minds and Machines (CBMM), 2014.