

**Contextual Vocabulary Acquisition
as Computational Philosophy
and as Philosophical Computation**

William J. Rapaport¹ and Michael W. Kibby²

**¹Department of Computer Science and Engineering,
Department of Philosophy, Department of Linguistics,
and Center for Cognitive Science**

`rapaport@cse.buffalo.edu`

`http://www.cse.buffalo.edu/~rapaport/`

**²Department of Learning and Instruction
and Center for Literacy and Reading Instruction**

`mwkibby@buffalo.edu`

`http://www.gse.buffalo.edu/FAS/Kibby/`

State University of New York at Buffalo, Buffalo, NY 14260

September 22, 2006

Abstract

Contextual vocabulary acquisition (CVA) is the active, deliberate acquisition of a meaning for an unknown word in a text by reasoning from textual clues, prior knowledge, and hypotheses developed from prior encounters with the word, but without external sources of help such as dictionaries or people. Published strategies for doing CVA vaguely and unhelpfully tell the reader to “guess”. AI algorithms for CVA can fill in the details that replace “guessing” by “computing”; these details can then be converted to a curriculum that can be taught to students to improve their reading comprehension. Such algorithms also suggest a way out of the Chinese Room and show how holistic semantics can withstand certain objections.

1 Computational Philosophy and Philosophical Computation

Computer science in general, and AI in particular, have a lot to give to philosophy, and vice versa, as Daniel Dennett once noted (1978: 126; cf. Rapaport 1986b). This essay discusses an interdisciplinary, applied cognitive-science research project that exhibits how philosophy can influence AI, how AI can influence philosophy, and how both can influence educational practice.

I¹ take “computational philosophy” to be the application of computational (i.e., algorithmic) solutions to philosophical problems.² An example from my own research would be the use of the SNePS knowledge-representation, reasoning, and belief-revision system (e.g., Shapiro 1979; Shapiro & Rapaport 1987, 1992; Martins & Shapiro 1988) to solve problems in the representation of fictional entities (Rapaport 1991, Rapaport & Shapiro 1995).

Not only can philosophers gain insights from computational techniques. The flow of information in the other direction gives rise to “philosophical computation”: the application of philosophy to problems in computer science. Another example from my own research is the application of Hector-Neri Castañeda’s theory of quasi-indexicals (1966, 1967, 1989) to solve problems in knowledge representation and reasoning. (The problem appears in Maida & Shapiro 1982; a resolution first appeared in Rapaport 1986b and then in corrected and generalized form in Rapaport, Shapiro, & Wiebe 1997.)

I will say more about computational philosophy and philosophical computation in §4. Now, I want to review an

¹In this essay, ‘I’, ‘my’, etc., refer to Rapaport; ‘we’, ‘our’, etc., usually refer to Rapaport and Kibby.

²This differs slightly from an earlier characterization of it as the study of which aspects of cognition are computable (Shapiro 1992, Rapaport 2003a); I leave for another time an examination of their connections.

interdisciplinary, cognitive-science project that can be classified as both. (For earlier presentations of some of this material, see, Ehrlich 1995; Ehrlich & Rapaport 1997; Rapaport & Ehrlich 2000; Rapaport 2003b, 2005.)

2 Contextual Vocabulary Acquisition

We characterize “contextual vocabulary acquisition” (CVA) as the active, deliberate acquisition (or learning) of a meaning for a word in a text by reasoning from “context”. (I place that last word in scare quotes, because I am using it in a non-standard sense that will be explicated in §2.4.)

CVA is what you do when you’re reading, you come to an unfamiliar word (or a familiar one whose meaning you’re unfamiliar with or that is being used in a new way), and you decide that understanding the word is important for understanding the passage that you’re reading. But there’s no one around whom you can ask, and a dictionary doesn’t help, either because there’s also no dictionary around that you can consult; or there is, but you’re too lazy to look the word up; or you look it up, but the word is not in your dictionary; or it is, but the definition is useless, either because it is too hard to understand or because it is inappropriate (cf. Miller 1985, 1986).

2.1 Dictionaries Don’t Help

Our favorite examples of useless dictionary definitions are these:

1. The first definition of ‘college’ in *Webster’s Ninth New Collegiate Dictionary* is “a body of clergy living together and supported by a foundation”. The *fifth* definition is “an independent institution of higher learning offering a course of general studies leading to a bachelor’s degree” (Mish 1983: 259). This would not be helpful to a reader (perhaps a non-native speaker) who wanted to understand the sentence “My daughter started college this Fall” and who did not know that *Webster’s 9th* orders its definitions historically rather than in order of commonality or frequency (Mish 1983: 19).
2. A student came across ‘infract’ and was told by her teacher to look it up on *Merriam-Webster OnLine*³, which defines it as a “transitive verb” associated with the noun ‘infracation’ and recommends its entry on ‘infringe’ for etymological information. Not understanding ‘infringe’ any better than ‘infract’, the student took the

³<http://www.m-w.com/dictionary/>

dictionary's suggestion and looked that up. The entry for 'infringe' suggests that a synonym is 'encroach'.⁴ Since she did not understand 'encroach' any better than 'infract' or 'infringe', she looked *that* up and found it defined as "to enter by gradual steps or by stealth into the possessions or rights of another", which is pretty far from any reasonable definition of 'infract'. Now, Schwartz 1988 observes that the definition that a friend or a dictionary might provide is just another context containing the word, so she might have been able to use CVA on some or all of these definitions. But then what would have been the purpose of looking 'infract' up in the first place? I submit that very little of this is any more useful than doing CVA directly and ignoring the dictionary.

Other dictionaries are more helpful, in particular, those aimed at non-native speakers. The *Collins COBUILD Dictionary* (Sinclair 1987) offers these more helpful definitions:

"A **college** is ... an institution where students study for qualifications or do training courses after they have left school." (p. 267.)

"If you **infringe** a law or an agreement, you break it." (p. 748; 'infract' is not in this dictionary.)

"To **encroach** on or upon something means to slowly take possession or control of it, so that someone else loses it bit by bit." (p. 464.)

In any case, the interesting situations requiring CVA are those in which you do not have such outside help as another person or a dictionary (online or hardcopy), as is usually the case with newly-coined terms. Thus, with or without such outside assistance, you still need to "figure out" a meaning for the word "from context". By 'figure out', I mean compute (or infer) a hypothesis about what the word might mean in that text.

We'll get to what I mean by 'context' in a moment, but I should say something here about the somewhat unusual phrase "*a meaning for a word*" (used in the first sentence of §2) vs. the more common "*the meaning of a word*". To talk about "the" meaning of a word suggest that there is a single, correct meaning, and to talk about the meaning "of" a word suggests (at least to my ears) that the meaning somehow belongs to the word. By contrast, I prefer to talk about "a" meaning for a word, because I believe that words have many possible meanings, depending on the textual context, on the reader's prior knowledge, etc. And I prefer to talk about a meaning "for" a word, because—especially in our project—the reader hypothesizes a meaning from "context" and *gives* it to the word.

⁴To be fair to *Merriam-Webster OnLine*, it only suggests that 'encroach' is a synonym for 'infringe' understood as an *intransitive* verb, but that subtlety was lost on the student.

2.2 Examples of CVA

Here are a few examples of CVA in action.

2.2.1 'brachet'

In *Le Morte Darthur* (Malory's 1470 version of the King Arthur legends), we find the following sequence of sentences:

1. There came a white hart running into the hall with a white **brachet** next to him, and thirty couples of black hounds came running after them. (p. 66; my boldface, here and below.)
2. As the hart went by the sideboard, the white **brachet** bit him. (p. 66.)
3. The knight arose, took up the **brachet** and rode away with the **brachet**. (p. 66.)
4. A lady came in and cried aloud to King Arthur, "Sire, the **brachet** is mine. (p. 66.)
10. There was the white **brachet** which bayed at him fast. (p. 72.)
18. The hart lay dead; a **brachet** was biting on his throat, and *other hounds* came behind. (p. 86; my italics.)

(Sentences (5)–(9) and (11)–(17) add no new information and are omitted here.) After reading sentence (1), most readers think that a brachet is an animal, though some (*mea culpa!*) think it is a piece of furniture in King Arthur's hall or a physical object attached to the hart. (We assume, by the way, that the reader knows the meaning of all other words in the sentence, such as 'hart', as well as understanding the sentence's grammar. Trying to figure out a meaning for one word in one sentence when two words are unknown is like trying to solve a single equation with two unknown variables; there's not enough information for more than a relative definition of one word in terms of the other.) Note, by the way, that if the reader thought that a brachet was a piece of furniture (which it isn't), *and if the reader never saw the word again*, then it probably wouldn't matter for understanding the story that the reader failed to figure out the intended meaning.

After reading (2), almost all readers decide that a brachet is an animal, because they have "prior knowledge" to the effect that only animals bite (and we assume that they know that a sideboard *is* a piece of furniture in King Arthur's hall). Note that "animal" is consistent with both sentences, whereas "furniture" is not. Alternatively, a reader might

stick to a belief that a brachet is a piece of furniture, and interpret ‘bite’ as metaphorically referring to a sharp corner on it that gouged the hart. (For a discussion of CVA and metaphor, see Budiu & Anderson 2001.)

After reading (3), about half of the people that we have verbal protocols from say that it gives them no new information—until I ask them how big a brachet is. Then, all readers agree that a brachet is small enough to be picked up. So far, then, a reader might have begun with a “wrong” hypothesis, no hypothesis, or a minimal hypothesis about what ‘brachet’ meant, and corrected or expanded it to “a small animal”, which is consistent with all three sentences read so far.

After reading (4), the new hypothesis is something like: A brachet is a small animal that is valuable, like a pet. After (10), readers with prior knowledge (here, perhaps prior “belief” would be a better term) that only hounds bay (e.g., at the moon) typically refine their hypothesis to: A brachet is a hound (or hunting dog). (It is important that each hypothesis be consistent with all occurrences of the word in any context.)

By this time, reading sentence (18) merely serves to confirm this hypothesis (‘other hounds’ only makes sense if the brachet is also a hound). However, this is a particularly helpful context, because, had it been the only sentence containing ‘brachet’, the reader could have determined that a brachet was a hound from its lexical and grammatical information alone.

2.2.2 ‘grudgingly’

Adverbs and adjectives are particularly hard to figure out using CVA techniques. (Is a *foobar* car a red car? a large car? a fast car? Without a great deal of helpful context, there’s little way to tell, although ‘foobar’ probably does *not* mean “salty”.) Nevertheless, there *are* inferences one can draw. Consider this example of an allegedly “misdirective” context (i.e., one that allegedly “direct[s] the . . . [reader] to an *incorrect* meaning for a target word”; Beck et al. 1983: 178; my emphasis):

Sandra had won the dance contest and the audience’s cheers brought her to the stage for an encore. “Every step she takes is so perfect and graceful,” Ginny said *grudgingly*, as she watched Sandra dance. (Beck et al. 1983: 178.)

First, the sentence can be rearranged to make the unfamiliar word its subject:

Grudgingly is how Ginny said “Every step she takes is so perfect and graceful,” as she watched Sandra dance.

Rather than jumping to the conclusion that ‘grudgingly’ means “admiringly”, as Beck et al. expect, the reader can decide merely that ‘grudgingly’ is a word that describes how someone might praise (or, for the cynical, apparently praise) someone. Next, the reader can list various ways of uttering such praise, e.g., admiringly, quickly, quietly, loudly, happily, in a frustrated manner, reluctantly, etc. Finally, the reader can then search the rest of the text for further clues as to which of these (if any) might be most appropriate. For example, knowing that Ginny was jealous of Sandra (as a larger textual context might indicate) might help us rule out all but the last two in that list. Note that it doesn’t matter if there are clues that lead us to think that ‘grudgingly’ might be closer in meaning to the first adverb on the list rather than the last. We are not after accuracy, merely understanding, and even a mistaken understanding that does not seriously harm the overall text-meaning suffices.

2.3 CVA as Computational Philosophy and as Cognitive Science

Before saying more about our project, let me review its genesis. In “How to Make the World Fit Our Language” (Rapaport 1981), I offered a neo-Meinongian theory of a word’s meaning for a person at a time as the set of contexts in which the person had heard or read the word before that time. I wondered if that notion could be made precise. When I began my study of AI, I realized that semantic-network theory offered a computational tool for doing just that.

Later, I learned that computational linguists, reading educators (see, e.g., Kibby 1995), second-language educators, psychologists, etc., were all interested in this topic (though very few had bothered to read the others’ literature!). It turned out to be a really multidisciplinary cognitive-science problem (for a bibliography, see Rapaport 2006a), and a computational implementation was developed by my student Karen Ehrlich (1995).

2.4 What Is the “Context” for CVA?

By ‘context’, I do not mean merely *textual* context, i.e., the surrounding words, or what is sometimes called the ‘co-text’ of the unfamiliar word (Brown & Yule 1983: 46–50, citing Halliday; Haastrup 1991). Instead, I mean what I shall sometimes call “wide” context: the “internalized” co-text “integrated” via belief revision with the reader’s “prior knowledge”. The scare-quoted terms need to be explicated.

By ‘internalized co-text’, I mean the reader’s interpretive mental model of the textual co-text. (Here, I am not using “mental model” as a brand name (Johnson-Laird 1983) but merely as a description of a conceptual representation “in” a reader’s mind.) By ‘integrated via belief revision’, I mean that the reader infers new beliefs from the internalized co-text together with his or her prior knowledge, removing inconsistencies. And “prior knowledge” includes “world” knowledge, commonsense knowledge, language knowledge, and previous hypotheses about the word’s meaning.

Furthermore, in the interesting cases, wide context *excludes* external sources such as a human or a dictionary. Thus, *the “context” for CVA is in the reader’s mind, not in the text.* Typically, a reader comes to a text with some prior knowledge; suppose for the sake of example, that a reader has several beliefs among his or her prior knowledge: PK1, PK2, The reader reads the first sentence (T1) in the text. The reader then “internalizes” it, i.e., creates in his or her mind a mental representation or interpretation of that sentence; denote it as I(T1). I(T1) might retain the exact wording of T1, or it might merely retain the propositional content of T1, or it might be a mental image of that propositional content, etc. (It might even be a misinterpretation of T1; cf. Rapaport 2003b, 2005.) The reader might draw an inference at this point (or later) using I(T1) and, say, PK1 as premises; call this conclusion P1. The reader reads the second sentence, T2; internalizes it as I(T2); and perhaps draws a further inference, P2, from (say) I(T2) and P1. The reader reads T3, internalizes it as I(T3), and perhaps notes that it is inconsistent with PK4. (Perhaps the text is a historical novel in which Abraham Lincoln is re-elected for a third term, and PK4 is the mental proposition that Lincoln died during his second term. For the sake of the novel, the reader decides to make believe that PK4 is false. For the details of this, see Rapaport & Shapiro 1995.) Finally, perhaps sentence T4 contains the unknown word ‘brachet’. The reader internalizes T4, and “searches” his or her mind for clues about a meaning for ‘brachet’, computing a hypothesis that a brachet is an animal. And so on. The point is that *all “contextual” reasoning is done in the reader’s mind*, which can be thought of as a belief-revised, integrated knowledge base (Rapaport 2003b).

2.5 Overview of the CVA Project

There are actually two kinds of CVA: incidental and deliberate. Incidental CVA is possibly the best explanation of how we learn most of our vocabulary. It has been estimated that the typical high-school graduate knows about 45,000 words. This figure is at the low end of published estimates. Even so, this means that, in the roughly 18 years it takes to learn them, about 2,500 words would have to be learned each year *on average*. However, only about 10% of this

number is explicitly taught over 12 years of schooling (statistics from Nagy & Anderson 1984). Therefore, about 90% of the words a high-school graduate knows must have been learned by other means, primarily just by hearing or reading them in context, i.e., by incidental CVA.

One of the other means of learning them is by deliberate CVA, the topic of our research. This is often taught in schools, but poorly in our estimation. To give you a taste of the state of the art circa the 1980s (and nothing has improved in the intervening 20 years; cf. Honig et al. 2000: 15.24–15.25, Graves et al. 2007: 220–222), let me describe two such programs.

Mueser & Mueser 1984, used as a text in at least one nationwide, for-profit, remedial-education center, claims to teach CVA. The student is first given a multiple-choice pre-test of several vocabulary words. Then, each word is presented in each sentence of a 4- or 5-sentence paragraph, one sentence of which always contains a definition of the word! A post-test (identical to the pre-test) unsurprisingly shows improvement.

Clarke & Nation 1980, writing from the perspective of English as a Second Language, offer a more useful strategy:

1. Determine the part of speech of the unfamiliar word.
2. Look at the grammatical context; who is doing what to whom?
3. Look at the surrounding textual context; search for clues.
4. *Guess the word*; check your guess.

The first three tasks are fine (and are similar to our algorithms; see §2.7). When I present this to computer scientists, they almost always laugh out loud at step (4). (When I present it to teachers, they think that it *all* sounds reasonable.) This is because “guessing” is too high-level an instruction; the student needs to be told in more detail *how* to guess. Of course, a computer program that implemented such a strategy would need all that detail.

And that is what Ehrlich 1995 implemented in a computational theory of CVA. We are currently developing an improved, teachable curriculum based on our algorithms. Before outlining our computational theory and curriculum, it is important to see *why* it is teachable. Although other computational linguists have developed algorithms similar to ours (e.g., Hastings & Lytinen 1994ab, Widdows & Dorow 2002), most work in computational linguistics on CVA falls into one of two (related) camps: There are statistical algorithms that examine large corpora to determine word meanings (e.g., Landauer et al. 1998, Gildea & Jurafsky 2002, Thompson & Mooney 2003; for a more complete

bibliography, see Rapaport 2006a), and there is work on “word-sense disambiguation” (WSD; e.g., Ide & Veronis 1998). The WSD research, most of which is also statistical in nature, differs from CVA as we do it in the following way: The WSD task is to determine which of a small number of given possible meanings a particular word has in a particular context. As Ellen Prince once described it to me, WSD is a multiple-choice test, but CVA as we do it is an open-ended essay question. The important point for teachability is that statistical techniques cannot be converted to a teachable curriculum. We cannot ask a student to read hundreds of texts containing an unfamiliar word and do a statistical analysis to compute a meaning. But good old-fashioned, symbolic, rule-based AI *can* be adapted for human use. (After all, the ur-theory of symbolic computation (Turing 1936) was explicitly modeled on human computing behavior. Here, we are merely reversing the process.)

2.6 Computational CVA

The computational CVA theory is implemented in SNePS, an intensional, propositional semantic-network, knowledge-representation, reasoning, and acting system that indexes information by nodes: From the “point of view” of any node, the rest of the network can be described. A SNePS network serves as a model of the mind of a reader. The system begins with a knowledge base consisting of SNePS representations of the reader’s prior knowledge, and it takes as input a SNePS representation of the unfamiliar word in its co-text. The processing simulates or models reading (or *is* reading, if you take the perspective of strong AI!). It uses logical inference, generalized inheritance, and belief revision to reason about the internalized text integrated with the reader’s prior knowledge. When queried about a meaning for the unfamiliar word, noun- or verb-definition algorithms deductively search this belief-revised, integrated knowledge base (the context) for slot fillers for a definition frame. The output is the definition frame, whose slots (or features) are things like class-membership information, structural information, actions, properties, etc., and whose fillers (or values) are the information gleaned from the context (the integrated knowledge base), with some redundancies removed.

For example, for the system to learn a meaning for ‘brachet’, we give it background information about harts, animals, King Arthur, etc., but no information about brachets. The input is a SNePS version of a simplified English version of Malory’s 15th-century English. So, e.g., ‘brachet’-sentence (1) from §2.2.1, above, becomes: “A hart runs into King Arthur’s hall. A white brachet is next to the hart.” The first of these is represented roughly as:

In the story, b1.2 is a hart.

In the story, b13 is a hall.

In the story, b13 is King Arthur's.

In the story, b12 runs into b13.

(where b12 and b13 are labels of nodes in the network). The second sentence is represented as:

In the story, b14 is a brachet.

In the story, b14 has the property "white".

While reading, SNePS infers that brachets are physical objects, because it has prior knowledge that white is a color and that only physical objects have color. When asked what 'brachet' might mean, it outputs this frame:

Definition of brachet:

Class Inclusions: phys obj,

Possible Properties: white,

Possibly Similar Items: animal, mammal, deer, horse, pony, dog,

I.e., a brachet is a physical object that can be white and that might be like an animal, mammal, deer, horse, pony, or dog. The "possibly similar items" are things that the knowledge base lists as being physical objects that can be white; this list will differ for different knowledge bases. A "possible property" is a property that at least one instance has, though there is not enough information to know whether all items of that kind have the property.

When the next sentence ("The brachet bit the hart's buttock.") is read in the context of prior knowledge to the effect that only animals bite, the definition frame is updated to include a "possible action" (and to delete a few "possibly similar items" that are no longer deemed to be similar):

Definition of brachet:

Class Inclusions: animal,

Possible Actions: bite buttock,

Possible Properties: white,

Possibly Similar Items: mammal, pony,

When "The knight picks up the brachet. The knight carries the brachet." is read, the definition frame is updated by adding small to the list of possible properties because of prior knowledge that only small things can be picked up or

carried. After reading “The lady says that she wants the brachet.”, `valuable` is added as a possible property because of prior knowledge that only valuable things are wanted. Finally, after reading “The brachet bays at Sir Tor.”, in the context of prior knowledge that only hunting dogs bay, the definition frame is updated to:

Definition of brachet:

Class Inclusions: hound, dog,

Possible Actions: bite buttock, bay, hunt,

Possible Properties: valuable, small, white,

I.e., a brachet is a hound (a kind of dog) that can bite, bay, and hunt, and that may be valuable, small, and white. Although the system still “believes” that a brachet is an animal, that information is redundant in light of the more specific, basic-level category it now lists.

Note that the system’s behavior is similar to the protocols taken from humans reading the passage (see §2.2.1, above), and that the definition that the system (as well as most humans) figures out is very close to that in the *Oxford English Dictionary Online*:⁵ A brachet is “a kind of hound which hunts by scent”. Our system missed the “hunting by scent” aspect of the definition, but, first, there was no information in the passage that would have generated that, and, second, it is not needed to understand the passage.

2.7 Noun and Verb Definition Algorithms

The noun and verb algorithms behave roughly as follows:

Noun Algorithm:

1. Generate an initial hypothesis by “syntactic manipulation”.

Just as in algebra one can solve an equation for an unknown value x by syntactically rearranging it (Rapaport 1986a), so one can “solve” a sentence for an unknown word x by syntactic rearrangement it: “A white brachet (x) is next to the hart” can be “solved for x ” to become “ x (a brachet) is something that is next to the hart and that can be white” (cf. Higginbotham 1989). In

⁵<http://dictionary.oed.com/>

other words, we initially “define” the node representing the unfamiliar word in terms of the nodes immediately connected to it.

2. Find in the wide context, or infer from it:

- (a) Basic-level class memberships (e.g., “dog”, rather than “animal”)
 - Else most-specific-level class memberships
 - Else names of individuals
- (b) Properties of *xs* (else of individual *xs*, etc.) (e.g., size, color, etc.)
- (c) Structure of *xs* (else . . .) (e.g., part-whole information, physical structure, etc.)
- (d) Acts that *xs* perform (else . . .) or that can be done to or with *xs*
- (e) Agents that do things to or with *xs*, or to whom things can be done with *xs*, or that own *xs*
- (f) Possible synonyms and antonyms

In other words, “define” the word *x* in terms of some (but not all) of the other nodes that it is connected to.

Verb Algorithm:

1. Generate initial hypothesis by syntactic manipulation.

2. Find in the wide context, or infer from it:

- (a) Class membership (e.g., such as is given by Conceptual Dependency classification (Schank & Rieger 1974) or Levin’s (1993) theory of verb classes)
 - What kind of act is *xing*? (e.g., walking (*x*) is a kind of moving)
 - What acts are kinds of *xings*? (e.g., sauntering is a kind of walking (*x*))
- (b) Properties or manners of *xing* (e.g., moving by foot, walking slowly)
- (c) Transitivity or subcategorization information (find class membership of agent, object, indirect object, instrument, etc.)
- (d) Possible synonyms and antonyms
- (e) Causes and effects of *xing*.

We have done some very preliminary work on an adjective algorithm, but, as noted, adjectives are very hard to figure out from context (see, however, Garver 2002, Lammert 2002, and §3, below). Finally, belief revision is used to revise incorrect definition hypotheses or the definitions of polysemous words (see Rapaport & Ehrlich 2000 for examples).

2.8 A Computational Theory of CVA

To sum up, our theory holds that:

1. A word does not have a unique meaning.
2. A word does not have a “correct” meaning, nor does the author’s intended meaning for a word have to be known by the reader in order for the reader to understand the word in context. Even familiar or well-known words can acquire new meanings in new contexts; neologisms are usually learnable only from context.
3. Every co-text can give some clue to a meaning for a word, by generating an initial hypothesis by syntactic manipulation.
4. But the co-text must be integrated with the reader’s prior knowledge; the larger the co-text and the more prior knowledge the reader has, the more clues the reader will have to compute a meaning. In fact, the more occurrences of the word that a reader encounters, the more likely the reader can asymptotically approach a stable meaning hypothesis.
5. CVA is computable. It is open-ended, hypothesis generation; it is not mere guessing of a missing or unknown word nor is it WSD. Some words will be easier to compute meanings for than others (nouns are easier than verbs, verbs are easier than modifiers).
6. CVA can improve general reading comprehension (through active reasoning).
7. Finally, CVA can, and should, be taught in schools.

(For defense of some of these claims, see Rapaport 2005.)

3 Outline of an Algorithm-Based Curriculum

But how should it be taught? As exemplified in Clarke & Nation’s strategy, as well as in work by Sternberg et al. 1983 and Fukkink 2005, what’s missing from most suggested curricula are the details about how to figure out (hypothesize, infer, compute) a meaning. The curriculum we are devising aims to use our CVA algorithms to help fill that gap. But converting an algorithm implemented as a computer program so that it can be understood by humans, in particular, by middle-school students (ages 11–15, approximately) requires some adjustments. For instance, we need to allow for something that humans are good at but that computers are not (yet) as good at: making hunches.⁶ We also need to allow for a human following such an algorithm to go “out of order”, and we don’t want the instructions to be so bogged down in what software engineers would consider “good software design” that they are unreadable by our target population. But, equally, we need to provide a fairly detailed sequence of steps to help a human who would be lost, or feel inadequate, if merely told to “guess”.

Our overall curriculum (cf. an earlier presentation in Kibby et al. (forthcoming)) begins with the teacher modeling CVA on a prepared text, possibly followed by having students “dare” the teacher to do CVA on a text that *they* choose. This would be followed by the teacher modeling CVA with student assistance, then students modeling it with teacher assistance, students doing CVA in small groups (working jointly), and finally, students doing CVA on their own.

The basic algorithm is a generate-and-test strategy:

1. Become aware of the unfamiliar word (call it x) and of the need to understand it. (Many readers simply ignore such words unless they have seen them multiple times or are otherwise motivated to figure out a definition for them.)
2. Repeat until you have a plausible meaning for x in its co-text (call the co-text $C(x)$):
 - (a) Generate a hypothesis about what x might mean in $C(x)$ (call it H).
 - (b) Test H .

To test H , do:

⁶I believe that for a computer to make a good hunch, it would probably need to have a statistically-based or neural-network program, rather than a symbolic one. Thus, I don’t think that making hunches is uncomputable, but it is beyond the scope of our current CVA program.

1. Replace all occurrences of x in $C(x)$ by H .
2. If the sentences containing x make sense, then proceed with your reading; else, generate a new hypothesis.

It is important to test H on *all* occurrences of x , to ensure that H is not merely “local” to a single occurrence but makes sense of the passage globally.

Our major contribution concerns the refinement of the “generate” step: To generate H , do:

1. Make an “intuitive” guess or “hunch”. (You may have an inexpressible “feeling” about what x might mean; this step is intended to allow you to test it out.)
2. If you can’t make any hunches, or if your hunch fails the test, then do any or all of the following three steps (in any order):
 - (a) Have you read or heard x before? If so, and if you can you remember its context or what you thought it meant, test that.
 - (b) Does x ’s morphological structure help? I.e., use your prior knowledge about the meanings of prefixes, suffixes, roots, etc., to generate H .
 - (c) Create an environment in which you might be able to make an “educated” guess:
 - i. Read and re-read the sentence (call it $S(x)$) containing x slowly and actively (i.e., think about it as you read).⁷
 - ii. Determine x ’s part of speech.
 - iii. Try to summarize what the entire text says so far.
 - iv. Activate (i.e., think about) your prior knowledge about the topic of the passage.
 - v. “Put 2 and 2 together”; i.e., what can you infer from combining the information in the text so far with what you already know about the topic?

3. Now comes the heart of our algorithmic theory of CVA:

- (a) Solve for x .

⁷For more on slow and active reading, see Rapaport 2006b.

(b) Search $C(x)$ for clues.

(c) Create H .

These last three steps can be elaborated as follows:

To solve for x , do:

1. Syntactically manipulate $S(x)$ so that x is the subject of the sentence (as discussed above).
2. List possible synonyms for x as “hypotheses-in-waiting” (as we did with ‘grudgingly’, above).

Use your prior knowledge to search $C(x)$ for clues:

1. If x is a noun, then look for clues that will give you information about:
 - (a) class membership (What kind of thing is an x ? What kinds of things are x s?)
 - (b) properties (What is x 's size, color, etc.?)
 - (c) structure (What parts do x s have? What wholes are x s part of? What is x 's physical structure?)
 - (d) acts (What can x s do? What can be done to or with x s?)
 - (e) agents (Who can do things to or with x s? To whom can things be done with x s? Who can own x s?)
 - (f) comparisons (Look for, or think of, possible synonyms.)
 - (g) contrasts (Look for, or think of, possible antonyms.)
2. If x is a verb, then look for clues that will give you information about:
 - (a) class membership (What kind of act is x ing? What kinds of acts are x ings?)
 - (b) properties of x ing (How can it be done?)
 - (c) transitivity or subcategorization information (Can you say “Someone can x ”, or “Someone can x something”, or “Someone can x something to someone else”? Look for agents and objects of x ing.)
 - (d) comparisons and contrasts (as for nouns)

If x is a adjective or adverb, then look for clues that will give you information about:

- (a) class membership (Does the modifier tell you something about color, size, shape, manner, etc.?)

- (b) contrasts (E.g., if you read, “He did it *xly* instead of *yly*”, where you know what *y* means, then you can hypothesize that *x* might be an opposite or complement of *y*.)
- (c) parallels (E.g., if you read, “He did it *xly*, *yly*, and *zly*”, and if you know what *y* and *z* are, and—if you’re lucky—it turns out that *y* and *z* are near-synonyms, then you can hypothesize that *x* means something very similar.)

Note that it doesn’t matter if you’re not sure if some clue that you find should be listed, say, as a “property” or “structural” information. What matters is finding the clues, not what to call them!

Finally, what do you do with all these clues? You need to create *H*, a tentative definition: your current hypothesis about what *x* means (or might mean, at least in the current context). It’s the analogue of both a scientific theory of *x* and of a detective’s hypothesis about “who done it”. So, how do you create a definition? Classically (going back to Aristotle), a definition consists of a “genus” and a “difference”: You say what kind of thing *x* is (its genus) and how it differs from other things of that kind. Schwartz & Raphael 1985 suggest creating a definition “map”—a semantic network that diagrammatically represents answers to the following questions:

What is it? (This gives you its genus or superclass.)

What is it like? (This also gives you its differences from other things in that genus.)

What are some examples?

Then try to create a single sentence (like a *Collins COBUILD* definition) that tells you what *x* is.

4 CVA as Computational Philosophy and Philosophical Computation

4.1 CVA and Holistic Semantic Theories

It is often said, with respect to semantic networks, that the “meaning” of a node is its location in the entire network (see, e.g., Quillian 1967: 413 (1985: 101)). This is also the rallying cry of holistic semantic theories: The meaning of a word is its relationships to all other words in the language. Jerry Fodor and Ernest Lepore (1992) have argued that holism entails the following unacceptable consequences:

- No two people can ever share a belief.

- No two people can ever mean the same thing.
- No single person ever means the same thing at different times.
- No one can ever change his or her mind.
- Nothing can be contradicted.
- Nothing can be translated.

I have argued elsewhere that these consequences either don't follow or are actually acceptable (Rapaport 2002, 2003a). Here, I wish to make the point that CVA offers a principled way to restrict the "entire network" to a useful subnetwork that can be shared across people, individuals, languages, etc., and that also can account for language and conceptual change.

4.1.1 CVA and the Chinese Room

John Searle's (1980) Chinese Room Argument postulates an individual (Searle) in a sealed-off room accompanied only by an instruction book containing an algorithm for manipulating "squiggles" that are input to it and which, after manipulation, are output from it. A native speaker of Chinese outside the room is inputting a story in Chinese and reading-comprehension questions in Chinese, and receiving as output appropriate answers in fluent Chinese. Does Searle-in-the-room understand Chinese? Without tackling that here, let us ask instead how Searle-in-the-room might be able to figure out a meaning for an unknown squiggle, relying solely on a co-text of other squiggles together with prior knowledge (perhaps stored in the Room's instruction book). My answer is: by CVA techniques!

Searle's Chinese Room Argument from semantics is this:

1. Computer programs are purely syntactic.
2. Cognition is semantic.
3. Syntax alone does not suffice for semantics.
4. Therefore, no purely syntactic computer program can exhibit semantic cognition.

But my theory of “syntactic semantics” (e.g., Rapaport 1986b, 1988, 2000, 2002, 2003a, and forthcoming) says that syntax does suffice for the kind of semantics needed for natural-language understanding in the Chinese Room: All input (linguistic, perceptual, etc.) is encoded in a single network (in fact, in a single, real—not artificial—neural network, namely, the brain), and all relations (including semantic ones) among the nodes of such a network are manipulated syntactically, hence computationally. CVA helps make this precise.

5 Summary

Our CVA project is both computational philosophy and philosophical computation, with applications to computational linguistics and to reading education. We have an explicit, good-old-fashioned, symbolic AI theory of how to do CVA, which is, therefore, teachable. Our goal is not to teach people to “think like computers” but to explicate computable and teachable methods of hypothesizing word meanings from context. It is also an example of “computational psychology” (Shapiro 1992, Rapaport 2003a), since we are devising computer programs that faithfully simulate (human) cognition and can tell us something about the (human) mind. We are teaching a machine, to see if what we learn in teaching *it* can help us teach *students* better.⁸

Acknowledgments

The computational theory was developed by Karen Ehrlich under Rapaport’s direction. The curriculum is joint work by Rapaport and Kibby. We are grateful to Stuart C. Shapiro and the members of the SNePS Research Group, and to Tanya Christ, Debra Dechert, and Karen Wieland for discussion and assistance with curricular development.

References

- Beck, Isabel L.; McKeown, Margaret G.; & McCaslin, Ellen S. (1983), “Vocabulary Development: All Contexts Are Not Created Equal”, *Elementary School Journal* 83(3): 177–181.
- Brown, Gillian, & Yule, George (1983), *Discourse Analysis* (Cambridge, UK: Cambridge University Press).

⁸This essay is based in part on a talk given by Rapaport at the North American Computing and Philosophy Conference (NA-CAP 2006) at Rensselaer Polytechnic University, August 2006.

- Budiu, Raluca, & Anderson, John R. (2001), "Word Learning in Context: Metaphors and Neologisms", *Technical Report CMU-CS-01-147* (Pittsburgh: Carnegie Mellon University, School of Computer Science)
- [http://act-r.psy.cmu.edu/papers/158/rb_jra_2001_a.pdf].
- Castañeda, Hector-Neri (1966), "'He': A Study in the Logic of Self-Consciousness", *Ratio* 8: 130–157.
- Castañeda, Hector-Neri (1967), "Indicators and Quasi-Indicators", *American Philosophical Quarterly* 4: 85–100.
- Castañeda, Hector-Neri (1989), *Thinking, Language, and Experience* (Minneapolis: University of Minnesota Press).
- Clarke, D.F., & Nation, I.S.P. (1980), "Guessing the Meanings of Words from Context: Strategy and Techniques", *System* 8: 211–220.
- Dennett, Daniel C. (1978), "Artificial Intelligence as Philosophy and as Psychology", in D.C. Dennett (ed.), *Brainstorms: Philosophical Essays on Mind and Psychology* (Montgomery, VT: Bradford Books): 109–126.
- Ehrlich, Karen (1995), "Automatic Vocabulary Expansion through Narrative Context", *Technical Report 95-09* (Buffalo: SUNY Buffalo Department of Computer Science).
- Ehrlich, Karen, & Rapaport, William J. (1997), "A Computational Theory of Vocabulary Expansion", *Proceedings of the 19th Annual Conference of the Cognitive Science Society (Stanford University)* (Mahwah, NJ: Lawrence Erlbaum Associates): 205–210.
- Fodor, Jerry, & Lepore, Ernest (1992), *Holism: A Shopper's Guide* (Cambridge, MA: Basil Blackwell).
- Fukkink, R.G. (2005), "Deriving Word Meaning from Written Context: A Process Analysis", *Learning and Instruction* 15: 23–43.
- Garver, Christopher (2002), "Adjective Representation in Contextual Vocabulary Acquisition"
- [<http://www.cse.buffalo.edu/~rapaport/CVA/Taciturn/garver.finalreport.doc>].
- Gildea, Daniel, & Jurafsky, Daniel (2002), "Automatic Labeling of Semantic Roles", *Computational Linguistics* 28(3): 245–288.
- Graves, Michael F.; Juel, Connie; & Graves, Bonnie B. (2007), *Teaching Reading in the 21st Century, 4th Edition* (Boston: Pearson Allyn & Bacon).
- Haastруп, Kirsten (1991), *Lexical Inferencing Procedures, or Talking about Words* (Tübingen: Gunter Narr Verlag Tübingen).
- Hastings, Peter M., & Lytinen, Steven L. (1994a), "The Ups and Downs of Lexical Acquisition", *Proceedings of the 12th National Conference on Artificial Intelligence (AAAI-94, Seattle)* (Menlo Park, CA: AAAI Press/MIT Press): 754–759.
- Hastings, Peter M., & Lytinen, Steven L. (1994b), "Objects, Actions, Nouns, and Verbs", *Proceedings of the 16th Annual Conference of the Cognitive Science Society* (Hillsdale, NJ: Lawrence Erlbaum Associates): 397–402.
- Higginbotham, James (1989), "Elucidations of Meaning", *Linguistics and Philosophy* 12: 465–517.
- Honig, Bill; Diamond, Linda; & Gutlohn, Linda (2000), *Teaching Reading Sourcebook for Kindergarten through Eighth Grade*

- (Novato, CA: Arena Press).
- Ide, Nancy M., & Veronis, J. (eds.) (1998), Special Issue on Word Sense Disambiguation, *Computational Linguistics* 24(1).
- Johnson-Laird, Philip N. (1983), *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness* (Cambridge, MA: Harvard University Press).
- Kibby, Michael W. (1995), "The Organization and Teaching of Things and the Words that Signify Them", *Journal of Adolescent and Adult Literacy* 39(3): 208–223.
- Kibby, Michael W.; Rapaport, William J.; Wieland, Karen M.; & Dechert, Debra A., "CSI: Contextual Semantic Investigation for Word Meaning", in Lawrence A. Baines (ed.), *Multisensory Learning* (tentative title) [<http://www.cse.buffalo.edu/~rapaport/CVA/CSI.pdf>].
- Lammert, Adam (2002), "Defining Adjectives through Contextual Vocabulary Acquisition" [http://www.cse.buffalo.edu/~rapaport/CVA/Taciturn/lammert.full_report_ms.doc].
- Landauer, Thomas K.; Foltz, Peter W.; & Laham, Darrell (1998), "An Introduction to Latent Semantic Analysis", *Discourse Processes* 25: 259–284.
- Levin, Beth (1993), *English Verb Classes and Alternations: A Preliminary Investigation* (Chicago: University of Chicago Press).
- Maida, Anthony S., & Shapiro, Stuart C. (1982), "Intensional Concepts in Propositional Semantic Networks", *Cognitive Science* 6: 291–330; reprinted in Ronald J. Brachman & Hector J. Levesque (eds.), *Readings in Knowledge Representation* (Los Altos, CA: Morgan Kaufmann, 1985): 169–189.
- Malory, Sir Thomas (1470), *Le Morte Darthur*, ed. by R.M. Lumiansky (New York: Collier Books, 1982).
- Martins, João, & Shapiro, Stuart C. (1988), "A Model for Belief Revision", *Artificial Intelligence* 35: 25–79.
- Miller, George A. (1985), "Dictionaries of the Mind", *Proceedings of the 23rd Annual Meeting of the Association for Computational Linguistics (University of Chicago)* (Morristown, NJ: Association for Computational Linguistics): 305–314.
- Miller, George A. (1986), "Dictionaries in the Mind", *Language and Cognitive Processes* 1(3): 171–185.
- Mueser, Anne Marie, & Mueser, John Alan (1984), *Practicing Vocabulary in Context* (SRA/McGraw-Hill).
- Mish, Frederick C. (1983), *Webster's Ninth New Collegiate Dictionary* (Springfield, MA: Merriam-Webster).
- Nagy, William E., & Anderson, Richard C. (1984), "How Many Words Are There in Printed School English?", *Reading Research Quarterly* 19(3, Spring): 304–330.
- Quillian, M. Ross (1967), "Word Concepts: A Theory and Simulation of Some Basic Semantic Capabilities", *Behavioral Science* 12: 410–430; reprinted in R.J. Brachman & H.J. Levesque (eds.), *Readings in Knowledge Representation* (Los Altos, CA: Morgan Kaufmann, 1985): 97–118.

- Rapaport, William J. (1981), "How to Make the World Fit Our Language: An Essay in Meinongian Semantics", *Grazer Philosophische Studien* 14: 1–21.
- Rapaport, William J. (1986a), "Searle's Experiments with Thought", *Philosophy of Science* 53: 271–279.
- Rapaport, William J. (1986b), "Logical Foundations for Belief Representation", *Cognitive Science* 10: 371–422.
- Rapaport, William J. (1988), "Syntactic Semantics: Foundations of Computational Natural-Language Understanding", in James H. Fetzer (ed.), *Aspects of Artificial Intelligence* (Dordrecht, Holland: Kluwer Academic Publishers): 81–131 (errata, [<http://www.cse.buffalo.edu/~rapaport/Papers/synsem.original.errata.pdf>]); reprinted in Eric Dietrich (ed.) (1994), *Thinking Computers and Virtual Persons: Essays on the Intentionality of Machines* (San Diego: Academic Press): 225–273.
- Rapaport, William J. (1991), "Predication, Fiction, and Artificial Intelligence", *Topoi* 10: 79–111.
- Rapaport, William J. (2000), "How to Pass a Turing Test: Syntactic Semantics, Natural-Language Understanding, and First-Person Cognition", *Journal of Logic, Language, and Information* 9(4): 467–490; reprinted in James H. Moor (ed.), *The Turing Test: The Elusive Standard of Artificial Intelligence* (Dordrecht, the Netherlands: Kluwer Academic Publishers, 2003): 161–184.
- Rapaport, William J. (2002), "Holism, Conceptual-Role Semantics, and Syntactic Semantics", *Minds and Machines* 12(1): 3–59.
- Rapaport, William J. (2003a), "What Did You Mean by That? Misunderstanding, Negotiation, and Syntactic Semantics", *Minds and Machines* 13(3): 397–427.
- Rapaport, William J. (2003b), "What Is the 'Context' for Contextual Vocabulary Acquisition?", in Peter P. Slezak (ed.), *Proceedings of the 4th International Conference on Cognitive Science/7th Australasian Society for Cognitive Science Conference (ICCS/ASCS-2003; Sydney, Australia)* (Sydney: University of New South Wales), Vol. 2, pp. 547–552.
- Rapaport, William J. (2005), "In Defense of Contextual Vocabulary Acquisition: How to Do Things with Words in Context", in Anind Dey, Boicho Kokinov, David Leake, & Roy Turner (eds.), *Modeling and Using Context: 5th International and Interdisciplinary Conference, CONTEXT 05, Paris, France, July 2005, Proceedings* (Berlin: Springer-Verlag Lecture Notes in Artificial Intelligence 3554): 396–409.
- Rapaport, William J. (compiler) (2006a), "A (Partial) Bibliography (in Chronological Order) of (Computational) Theories of Contextual Vocabulary Acquisition" [<http://www.cse.buffalo.edu/~rapaport/refs-vocab.html>].
- Rapaport, William J. (2006), "How to Study" [<http://www.cse.buffalo.edu/~rapaport/howtostudy.html>].
- Rapaport, William J. (forthcoming), "How Helen Keller Used Syntactic Semantics to Escape from a Chinese Room" [<http://www.cse.buffalo.edu/~rapaport/Papers/helenkeller-MM-revised.pdf>].
- Rapaport, William J., & Ehrlich, Karen (2000), "A Computational Theory of Vocabulary Acquisition", in Łucja M. Iwańska & Stuart C. Shapiro (eds.), *Natural Language Processing and Knowledge Representation: Language for Knowledge*

- and Knowledge for Language* (Menlo Park, CA/Cambridge, MA: AAAI Press/MIT Press): 347–375; errata online at: [<http://www.cse.buffalo.edu/~rapaport/Papers/krnlp.errata.pdf>].
- Rapaport, William J., & Shapiro, Stuart C. (1995), “Cognition and Fiction”, in Judith Felson Duchan, Gail A. Bruder, & Lynne E. Hewitt (eds.), *Deixis in Narrative: A Cognitive Science Perspective* (Hillsdale, NJ: Lawrence Erlbaum Associates): 107–128; abridged and slightly edited as “Cognition and Fiction: An Introduction”, in Ashwin Ram & Kenneth Moorman (eds.), *Understanding Language Understanding: Computational Models of Reading* (Cambridge, MA: MIT Press, 1999): 11–25.
- Rapaport, William J.; Shapiro, Stuart C.; & Wiebe, Janyce M. (1997), “Quasi-Indexicals and Knowledge Reports”, *Cognitive Science* 21: 63–107; reprinted in Francesco Orilia & William J. Rapaport (eds.), *Thought, Language, and Ontology: Essays in Memory of Hector-Neri Castañeda* (Dordrecht: Kluwer Academic Publishers, 1998): 235–294.
- Schank, Roger C., & Rieger, Charles J. (1974), “Inference and the Computer Understanding of Natural Language”, *Artificial Intelligence* 5: 373–412.
- Schwartz, Robert M. (1988), “Learning to Learn Vocabulary in Content Area Textbooks”, *Journal of Reading*: 32 (November): 108–118.
- Schwartz, Robert M., & Raphael, Taffy E. (1985), “Concept of Definition: A Key to Improving Students’ Vocabulary”, *The Reading Teacher* 39: 198–203.
- Searle, John R. (1980), “Minds, Brains, and Programs,” *Behavioral and Brain Sciences* 3: 417–457.
- Shapiro, Stuart C. (1979), “The SNePS Semantic Network Processing System”, in Nicholas Findler (ed.), *Associative Networks* (New York: Academic Press): 179–203.
- Shapiro, Stuart C. (1992), “Artificial Intelligence”, in Stuart C. Shapiro (ed.), *Encyclopedia of Artificial Intelligence, second edition* (New York: John Wiley & Sons): 54–57.
- Shapiro, Stuart C., & Rapaport, William J. (1987), “SNePS Considered as a Fully Intensional Propositional Semantic Network”, in Nick Cercone & Gordon McCalla (eds.), *The Knowledge Frontier: Essays in the Representation of Knowledge* (New York: Springer-Verlag): 262–315; shorter version appeared in *Proceedings of the 5th National Conference on Artificial Intelligence (AAAI-86, Philadelphia)* (Los Altos, CA: Morgan Kaufmann): 278–283; a revised shorter version appears as “A Fully Intensional Propositional Semantic Network”, in Leslie Burkholder (ed.), *Philosophy and the Computer* (Boulder, CO: Westview Press, 1992): 75–91.
- Shapiro, Stuart C., & Rapaport, William J. (1992), “The SNePS Family”, *Computers and Mathematics with Applications* 23: 243–275; reprinted in Fritz Lehmann (ed.), *Semantic Networks in Artificial Intelligence* (Oxford: Pergamon Press, 1992): 243–275.

- Sinclair, John (ed.) (1987), *Collins COBUILD English Language Dictionary* (London: Collins).
- Sternberg, Robert J.; Powell, Janet S.; & Kaye, Daniel B. (1983), "Teaching Vocabulary-Building Skills: A Contextual Approach",
in Alex Cherry Wilkinson (ed.), *Classroom Computers and Cognitive Science* (New York: Academic Press): 121–143.
- Thompson, Cynthia A., & Mooney, Raymond J. (2003), "Acquiring Word-Meaning Mappings for Natural Language Interfaces",
Journal of Artificial Intelligence Research 18: 1–44.
- Turing, Alan M. (1936), "On Computable Numbers, with an Application to the Entscheidungsproblem", *Proceedings of the London
Mathematical Society*, Ser. 2, Vol. 42: 230–265.
- Widdows, Dominic, & Dorow, Beate (2002), "A Graph Model for Unsupervised Lexical Acquisition", *Proceedings, COLING-2002*.