Forthcoming in Proc. 6th World Multiconf. on Systemics, Cybernetics & Informatics (SCI 2002).

Contextual Vocabulary Acquisition: A Computational Theory and Educational Curriculum

William J. RAPAPORT[1]
Department of Computer Science & Engineering, Department of Philosophy, Center for Cognitive Science

and

Michael W. KIBBY[2]
Department of Learning and Instruction, Center for Literacy and Reading Instruction
State University of New York at Buffalo, Buffalo, NY 14260-2000

ABSTRACT

We discuss a research project that develops and applies algorithms for computational contextual vocabulary acquisition (CVA): learning the meaning of unknown words from context. We try to unify a disparate literature on the topic of CVA from psychology, first- and second-language acquisition, and reading science, in order to help develop these algorithms: We use the knowledge gained from the computational CVA system to build an educational curriculum for enhancing students' abilities to use CVA strategies in their reading of science texts at the middle-school and college undergraduate levels. The knowledge gained from case studies of students using our CVA techniques feeds back into further development of our computational theory. **Keywords:** artificial intelligence, knowledge representation, reading, reasoning, science education, vocabulary acquisition.

No doubt you have on occasion read some text containing an unfamiliar word, but you were unable or unwilling to find out from a dictionary or another person what it meant. Nevertheless, you might, consciously or not, have figured out a meaning for it. Suppose you didn't, or suppose your hypothesized meaning was wrong. If you never see the word again, it may not matter. But, if the text were from science, technology, engineering, or mathematics ("STEM", as NSF likes to call these), not understanding the unfamiliar term might seriously hinder your subsequent understanding of the text. If you do see the word again, you will have an opportunity to revise your hypothesis about its meaning. The more times you see it, the better your definition will become. And if your hypothesis development were deliberate, rather than "incidental", your command of the new word would be stronger.

We are conducting a research project to (a) extend and develop algorithms for computational contextual vocabulary acquisition (CVA): learning, from context, meanings for "hard" words (common and proper nouns, verbs, adjectives, and adverbs), (b) unify a disparate literature on CVA from psychology, first- and second-language acquisition, and reading science, in order to help develop these algorithms, and (c) use the knowledge gained from the computational CVA system to build and evaluate an educational curriculum for enhancing students' abilities to use deliberate (i.e., non-incidental) CVA strategies in their reading of STEM texts at the middle-school and college undergraduate levels. The knowledge gained from case studies of students using our CVA techniques will feed back into further development of our computational theory.

It is generally agreed among CVA researchers that "incidental" vocabulary acquisition does occur [18]: People know more words than they are explicitly taught, so they must have learned most of them as a by-product of reading or listening. Furthermore, at least some of this incidental acquisition was the result of conscious processes of guessing, inferring, etc., the meaning of unknown words from context.

It is also generally agreed that we don't know *how* readers do much of this. The psychology and first- and second-language-learning literature suggests various strategies [2,3,13,31,32,34]. But most are quite vague: e.g., step 1: "look at the word itself and its surroundings to decide on the part of speech"; step 2: "look at the immediate grammar context of the word, usually within a clause or sentence"; step 3: "look at the wider context of the word usually beyond the level of the clause and often over several sentences" (looking for causal, temporal, categorical information, etc.); step 4: "guess… the word and check… that the guess is correct" [3,19]. This is hardly a detailed algorithm that could easily be followed by a student! (Step 4 is reminiscent of the famous Sydney Harris cartoon showing a complicated mathematical formula, in the middle of which occurs the phrase, "here, a miracle occurs".) One reason for this vagueness in the educational literature is that it is not clear exactly how context oper-

[1]rapaport@cse.buffalo.edu; http://www.cse.buffalo.edu/∼rapaport
[2]mwkibby@acsu.buffalo.edu; http://www.gse.buffalo.edu/FAS/Kibby/

ates, in large part because of the lack of research on this topic. In turn, this means there is no generally accepted curriculum or set of strategies for teaching CVA. We need to know more about how context operates and how we can teach it strategically. With this knowledge, we could more effectively help students be more aware of context and know better how to use it.

There are also computational theories that implement various CVA methods, which do go into much more detail on how to use context to infer meaning [1,7–10,35]. But most of these assume the prior existence of a known concept that the unknown word is to be mapped to; this is especially the case for the large body of research on word-sense disambiguation [11]. As linguist Ellen Prince has suggested (in conversation), that makes the task more like a multiple-choice quiz, whereas CVA as our system does it is more like an essay test.

What is needed (and what we have been working on) is a *general* method that (a) shows how CVA *can* be done and (b) is explicit enough to be taught to human readers. Such a theory is best expressed algorithmically, for then the methods are made fully explicit and can be tested computationally. Admittedly, this does not solve the problem of how humans *actually* do CVA, though it does provide testable ideas of how they *might* do it. And it certainly provides ideas for how they *could* do it and, hence, how it might be taught.

The computational aspects of our project build upon our previous work on the development of a computational CVA theory and software agent called 'Cassie' [5,6,21]. Cassie consists of the SNePS semantic-network knowledge-representation and reasoning (KRR) system and a knowledge base (KB) of background information representing the knowledge that a reader (e.g., Cassie) brings to the text containing the unknown term. Cassie's input consists, in part, of information from the text being read, which is parsed and incorporated directly into the knowledge-representation (KR) formalism. Cassie's other input is questions asked about the material being read. The question "What does [word] mean?" triggers a deductive search of the KB, now consisting of background information plus information from the text, all marked with its "degree" of immunity from revision (roughly, a measure of the trustworthiness of the source of information). Output consists of a report of Cassie's current definition of the word in its context, or answers to other queries.

"Hard" words might be novel ('brachet'), familiar but misunderstood (does 'smite' mean to kill by hitting hard, or merely to hit hard?), or familiar but used in a new sense (might "dressing" a sword mean to clothe it?). Our theory is that the meaning of such a word (1) *can* be determined from context (including the surrounding text, grammatical information, and the reader's background knowledge, but no access to a dictionary or other external source of information (including a human)), (2) can be *revised* upon further encounters, (3) "*converges*" to a dictionary-like definition if enough context has been provided and there have been enough encounters, but (4) is always subject to further revision. Each encounter yields a definition (a hypothesis about meaning) and provides an opportunity to revise it in light of new evidence. The revision is *unsupervised*: There is no (human) "trainer" and no "error-correction" techniques. Finally, no *subject-matter- (or "domain")-specific* antecedent background information is required for the development and revision of the hypothesized definition (with the exception of the word's part of speech). The domain-independence of our system can make it more difficult to develop a good definition quickly, but is intended to model the typical reader of an arbitrary text. Clearly, the more background knowledge, including specialized knowledge, that the reader brings to the text, the more efficiently the unknown term can be learned.

The technology we employ is the SNePS-2.5 KRR system [26,27,29,30]. Each node in a SNePS network represents a concept or mental object (possibly built of other concepts), linked by labeled arcs. All information, including propositions, is represented by nodes, and propositions about propositions can be represented without limit. Arcs merely form the underlying syntactic structure of SNePS. Paths of arcs can be defined, allowing for path-based inference, including property inheritance within generalization hierarchies. There is a 1–1 correspondence between nodes and represented concepts. This uniqueness principle guarantees that nodes will be shared whenever possible and that nodes represent "intensional" objects, i.e., concepts, propositions, properties, and such objects of thought as fictional entities non-existents, and impossible objects [27,28]. This wide representational ability is especially appropriate for CVA from arbitrary texts, whose subject matter could range from factual science to science fiction. After all, objects about which we can think, speak, or write need not exist; this includes not only unicorns, but possibly black holes, bosons, or other theoretical entities of contemporary STEM.

SNePS's inference package allows rules for both deductive and default reasoning. In the presence of a contradiction, the SNeBR belief-revision package allows for the removal of one or more of the propositions from which the contradiction was derived, as well as the conclusions that depended on it [15]. This mechanism is used to revise definitions that are inconsistent with a word's current use. We have developed algorithms for *partially* automating the identification and removal or modification of the offending premise, based on SNePSwD, a default belief-revision system that enables automatic revision [4,16]. SNePS also has an English lexicon, morphological analyzer/synthesizer, and a generalized augmented-transition-network parser-generator that, rather than building an intermediate parse tree, translates the input English directly into the propositional semantic-network KR system [23,24,30].

"Cassie", our computational CVA agent, consists of SNePS-2.5 (including SNeBR and the parser-generator), SNePSwD, and a KB of background information. Currently, the KB is hand-coded, because it represents Cassie's antecedent knowledge; *how* she acquired this knowledge is irrelevant. We begin with what some might call a "toy" KB, but each of our tests so far has included all previous information, so the KB grows as we test more words. Cassie's input consists, in part, of information from the text being read. Currently, this, too, is coded directly in the SNePS.

We take the meaning of a word (as understood by a cognitive agent) to be the position of that word in a highly interconnected network of words, propositions, and other concepts. In this (idiolectic) sense, the meaning of a word for a cognitive agent is determined by idiosyncratic experience with it. But a word's dictionary definition usually contains less information than that. Contextual meaning as thus described includes a word's relation to every concept in the agent's mind, which is too unwieldy to be of much use. To limit the connections used to provide the definition, we select for particular *kinds* of information. Not all concepts within a given subnetwork need be equally salient to a dictionary-style definition of a word. In the attempt to understand and be understood, people abstract certain conventional information about words and accept this information as a definition.

When a new word is encountered, people begin to hypothesize a definition. Applying the fundamental principle that the meaning of a term is its location in the network of background information and story information, our algorithms for hypothesizing a definition operate by deductively searching the network for information appropriate to a dictionary-like definition (see [21] for the algorithms themselves). We assume that our grammar has been able to identify the unknown word as a noun or a verb. We have developed algorithms for hypothesizing and revising meanings for nouns and verbs that are unknown, mistaken, or being used in a new way.

Cassie was provided with background information for understanding the King Arthur stories [14]. In one test, when presented with passages involving the unknown noun 'brachet', Cassie was able to develop a theory that a brachet was a dog whose function is to hunt and that can bay and bite. (*Webster's Second* [33] defines it as "a hound that hunts by the scent".) However, based on the first context in which the term appeared (viz., "Right so as they sat, 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 with a great cry."), the initial hypothesis was merely that a brachet was a physical object that may be white.

Each time the term appeared, Cassie was asked to define it. To do so, she deductively searched her background KB, together with the information she had read in the narrative to that point, for information concerning (1) direct class inclusions (especially in a basic-level category), (2) general functions of brachets (in preference to those of individuals), (3) the general structure of brachets (if appropriate, and in preference to those of individuals), (4) acts that brachets perform (partially ordered in terms of universality: probable actions in preference to possible actions, actions attributed to brachets in general in preference to actions of individuals, etc.), (5) possible ownership of brachets, (6) part/whole relationships to other objects, (7) other properties of brachets (when structural and functional description is possible, the less salient "other properties" of particular brachets are not reported, although we do report any properties that apply to brachets in general), and (8) possible synonyms for 'brachet' (based on similarity of the above attributes). Some of these are based on psycholinguistic studies of CVA [31,32]. In the absence of some or all of this information, or in the presence of potentially inconsistent information (e.g., if the text says that one brachet hunts and another does not), Cassie either leaves certain "slots" in her definitional framework empty, or includes information about particular brachets. Such information is filled in or replaced upon further encounters with the term.

To define a verb (*V*), we currently report its predicate structure, a categorization of its arguments, and any causal or enablement information we can find. Clearly, our verb-definition algorithm is not as elaborate as our noun-defining algorithm. We are endeavoring to remedy this.

In another test, Cassie was told that 'to smite' meant "to kill by hitting hard" (a mistaken belief actually held by one of us before reading [14]). Passages in which various characters were smitten but then continued to act triggered SNeBR, which identified several possible "culprit" propositions in the KB to remove in order to block inconsistencies. The reader then decides which belief to revise. Although the decision about which proposition (representing an incorrect definition) to withdraw and which *new* proposition (representing a revised definition) to add has been partially automated (using SNePSwD), this remains the one area still occasionally requiring human intervention. Automating this is a major focus of our research.

A third case is exemplified by 'to dress', which Cassie antecedently understood to mean "to put clothes on something". This is a well-entrenched meaning, which should *not* be rejected. However, upon reading that King Arthur "dressed" his sword, SNeBR detects an inconsistency. Rather than *rejecting* the prior definition, we *add* to it. In this case, Cassie decides that to dress is *either* to put clothes on *or* to prepare for battle.

It might be objected that it would be easier for the reader simply to go to a dictionary to look up the meaning of the unknown word. But not all words are in dictionaries, nor are dictionaries always readily available. In addition, many researchers have pointed out that dictionary definitions are neither always correctly understood by readers nor are they always useful [12,17]. Further, upwards of 90% of all the words we know are learned from context while reading or listening [18]. There is no intention here of demeaning the value of the dictionary, but it is simply not the case that all or most new words are learned by consulting one. This view is not compatible with the research on vocabulary acquisition between ages 0 and 18; the dictionary simply is not the major source of learning word meanings in elementary, middle, and high schools. Our intent, speaking broadly, is to find ways to facilitate students' natural CVA by developing a more rigorous theory of how context operates and creating a more systematic and viable curriculum for teaching students to use CVA strategies.

Another objection might be that teaching humans how to learn is not the same thing as teaching computers how to learn. To respond to this, we begin with some comments on the nature of Artificial Intelligence (AI). AI can be viewed in at least three ways [25]: (1) as a branch of engineering whose goal is to advance the field of computer science; this, however, is neither our immediate goal nor our methodology; (2) as "computational psychology", where the goal is to study human cognition using computational techniques; a good computational-psychology computer program will simulate some human cognitive task in a way that is faithful to human performance, with the same failures as well as successes—AI as cognitive psychology can tell us something about the human mind; or (3) as "computational philosophy", where the goal is to learn which aspects of cognition in general are computable; a good computational-philosophy computer program will simulate some cognitive task but not necessarily in the way that humans would do it—AI as computational philosophy can tell us something about the limits and scope of cognition in general, but not necessarily about human cognition in particular.

The present project falls under the category of computational psychology (and to a lesser extent under the category of computational philosophy). Our goal is not to teach people to "think like computers". Rather, our goal is to explicate methods for inferring the meanings of unknown words from context. The "vague" strategy mentioned above is *not* a caricature; it is the actual recommendation of one writer in the field of vocabulary acquisition! But neither is it an algorithm—i.e., an explicit, step-by-step procedure for solving a problem correctly.

Our goal is to "teach" (i.e., program) a computer to do the "educated" guessing—or inferencing—that is left vague in the strategy above. To do that, we must determine what information is needed and what inference methods must be supplied, and we must spell this all out in enough detail so that "even" a computer could do it [22]. But that is not all: For once we have such a method, we can then actually teach it to people, rather than leave them wondering what to do with all the contextual information that they might have found in steps 1–3 of the above vague strategy—we can teach them what information to look for and what to do with it. This is our final goal.

Our overall goals are to (1) make our existing computational CVA system [5] more robust: to improve or create algorithms for inferring the meanings of unknown nouns, verbs, adjectives, and adverbs, utilizing grammatical, morphological, and etymological information; (2) develop and test educational curricula at the secondary and post-secondary levels for teaching CVA methods, and (3) integrate these two tasks by using the computational theory and the educational curriculum to help develop each other. The focus of our current research is to concentrate on (3) in order to facilitate the eventual transfer between, and mutual interaction of, (1) and (2).

The computational stream of our research has as its main goal the development and implementation of a *computational theory* of CVA [21]; the educational stream has as its main goal the development and implementation of an *educational curriculum* in CVA. Although these two streams still have independent goals and, to some extent, independent methodologies, their full development must be intimately and synergistically integrated. Thus, the development of the bridge between the two research streams is its focus. Accordingly, in the computational stream, less time is being spent on developing new algorithms, and more on developing and revising the current ones for use by the educational stream. Similarly, the educational stream is not spending much time on *testing* new curricula, but on *using the computational system* to begin the *development* of new curricula, as well as *providing feedback* to the computational stream based on students' actual CVA techniques, in order to improve and further develop the algorithms. Thus, we are both using Cassie to teach humans *and* humans to teach Cassie.

This bridge-building is being accomplished as follows: (1) We are identifying common texts (not necessarily text*books*) that we will both work on. This is the focus of the "synergy" between the two streams. We are starting with the texts that the computational stream has already developed algorithms for [5], and we are looking for real texts that the students in our study might be reading. In addition to STEM *textbooks*, we are also looking at popular science writing such as is found in *Scientific American* or daily newspapers. The ability to read and understand such texts is an important aspect of STEM literacy, and such writing is more likely to require CVA skills than STEM text*books* (which are often quite detailed and specific in giving definitions of terms).

(2) The computational research stream is (a) developing computational grammars for those texts, (b) developing knowledge representations of both the texts and the background knowledge necessary for understanding and reasoning about them, (c) testing the current system on the unknown words in them, and (d) developing new algorithms, as necessary, for CVA on them, as well as ones based on student protocols: In the educational stream, we are eliciting student protocols (or "thinking-out-loud" records) of students' attempts to figure out the meanings of unknown words; the researchers in the computational stream will then try to formalize and implement any (successful!) methods that students actually use but that we have not (yet) implemented.

Another task we might tackle is to develop an *explanation facility* for Cassie that can be used by the students when they are stuck. That is, if students have trouble figuring out the meaning of an unknown word (or merely want to check to see if their answer is acceptable), they could ask Cassie for an explanation of how "she" figured it out (or, perhaps, for guidance on how the *student* could figure it out).

(3) The educational stream is (a) identifying students who will participate in the experiments, (b) having them read the chosen texts, (c) having them figure out the meaning of the unknown words in those texts, (d) eliciting protocols of their thought processes while doing this (which will be used to modify Cassie), and (e) beginning to develop educational curricula to teach them Cassie's (successful) techniques.

Why is our research important? As noted, it is a fact that most meaning vocabulary is learned from context; teachers have too little time for directly teaching an extensive list of meaning vocabulary. Also as noted, although the use of the dictionary is extremely important in learning words, it is the case that dictionaries are not always available, that dictionary definitions are not always decipherable, and that sometimes, humans just do not bother to go "look it up." Learning words from context is simply required if a student is to try to learn the many new terms and words that must be known to learn STEM topics [18,31].

Further, newly revised educational standards for language arts, science, social studies, and mathematics all call for students to have a greater command of concepts and the words that signify those concepts. Since these concepts and their words and terms cannot all be taught directly in the classroom, it is important that not only do we devote more instructional time in school to teaching CVA, but also gain more knowledge about what context is and how it operates.

Learning when and how to use CVA strategies has broader implications than just the classroom learning and learning standards, however. Students learn a great deal of STEM from reading trade books (i.e., books that are not textbooks), articles in general-interest children's magazines (e.g., *Highlights*, *Cricket*, *Spider*, *Weekly Reader*), and children's magazines devoted especially to science (e.g., *Spinner*, *TimeKids*, *Science Scholastic*, *Quantum*). If students are better able to use surrounding context to help determine the meaning of unknown words or terms, then more STEM will be learned when students are independently engaged in reading these materials. It is usually in reading magazines, trade books, and websites such as these that students first encounter articles on STEM. If schools are more effective in teaching CVA, and if the writers and editors of these articles structure their texts to accommodate CVA, then students will gain more knowledge and heighten their interest and motivation in STEM.

There are also considerations from a broader science-education perspective: One of the goals of education should be to instill in students the knowledge—and the confidence and life-long desire to use that knowledge—of how to learn on one's own. Most often, there are no ultimate authorities or experts to consult when one has a problem to solve or a question to answer [20]. This is just as true in the world of, say, particle physics (where there are no "answers in the back of the book") as it is when one comes across an unknown word while reading (and there is no dictionary or glossary at hand, or no other person who knows the word).

The skills required for CVA are not only useful for helping one read (and hence learn) on one's own, but are also among those most useful in science and mathematics: finding clues or evidence (among the context surrounding the unknown word), integrating them with one's background knowledge, and using both to infer (whether by deduction, induction, or abduction) the meaning of the unknown word. CVA is a wonderful model "in the small" of the scientific method of hypothesis formation, testing, and revision, as well as a useful tool for learning on one's own.[3]

[3] This research is supported by NSF grant #REC-0106338.

REFERENCES

[1] Barg, P., & Walther, M. (1998), "Processing Unknown Words in HPSG", in *Proc. COLING-ACL '98* (San Francisco: Morgan Kaufmann), 1: 91–95.

[2] Blachowicz, C., & Fisher, P. (1996), *Teaching Vocabulary in All Classrooms* (Englewood Cliffs, NJ: Prentice-Hall): 15–34.

[3] Clarke, D.F., & Nation, I.S.P. (1980), "Guessing the Meanings of Words from Context", *System* 8: 211–220.

[4] Cravo, M.R., & Martins, J.P. (1993), "SNePSwD: A Newcomer to the SNePS Family", *J. Exp. & Theor. Artif. Intel.* 5: 135–148.

[5] Ehrlich, K. (1995), "Automatic Vocabulary Expansion through Narrative Context", *Tech. Rep. 95-09* (Buffalo: SUNY Buffalo Computer Science Dept.).

[6] Ehrlich, K., & Rapaport, W.J. (1997), "A Computational Theory of Vocabulary Expansion", *Proc. 19th Annual Conf. Cog. Sci. Soc.* (Mahwah, NJ: Erlbaum): 205–210.

[7] Granger, R.H. (1977), "Foul-Up: a Program that Figures Out Meanings of Words from Context", *Proc. IJCAI-77* (Los Altos, CA: William Kaufmann): 67–68.

[8] Hahn, U., & Schnattinger, K. (1998), "A Text Understander that Learns", in *Proc. COLING-ACL '98* (San Francisco: Morgan Kaufmann), 1: 476–482.

[9] Hastings, P.M., & Lytinen, S.L. (1994a), "The Ups and Downs of Lexical Acquisition", *AAAI-94* (Menlo Park, CA: AAAI/MIT): 754–759.

[10] Hastings, P.M., & Lytinen, S.L. (1994b), "Objects, Actions, Nouns, and Verbs", *Proc. 16th Annual Conf. Cog. Sci. Soc.* (Hillsdale, NJ: Erlbaum): 397–402.

[11] Ide, N.M., & Veronis, J. (eds.) (1998), Special Issue on Word Sense Disambiguation, *Comp. Ling.* 24(1).

[12] Johnson-Laird, P.N. (1987), "The Mental Representation of the Meanings of Words", *Cogn.* 25: 189–211.

[13] Kibby, M.W. (1995), "The Organization and Teaching of Things and the Words that Signify Them", *J. Adolescent & Adult Literacy* 39(3): 208–223.

[14] Malory, Sir T. (1470), *Le Morte Darthur*, ed. by R.M. Lumiansky (New York: Collier, 1982).

[15] Martins, J., & Shapiro, S.C. (1988), "A Model for Belief Revision", *Artif. Intel.* 35: 25–79.

[16] Martins, J., & Cravo, M.R. (1991), "How to Change Your Mind", *Noûs* 25: 537–551.

[17] Miller, G.A., & Gildea, P.M. (1987, September), "How Children Learn Words", *Sci. Amer.*: 94–99, 120.

[18] Nagy, W.E.; Anderson, R.C.; & Herman, P.A. (1987), "Learning Word Meanings from Context during Normal Reading", *Amer. Ed. Res. J.* 24(2, Summer): 237–270.

[19] Nation, P., & Coady, J. (1988), "Vocabulary and Reading", in R. Carter & M. McCarthy (eds.), *Vocabulary and Language Teaching* (London: Longman): 97–110.

[20] Perry, W.G., Jr. (1970), *Forms of Intellectual and Ethical Development in the College Years* (New York: Holt, Rinehart, & Winston).

[21] Rapaport, W.J., & Ehrlich, K. (2000), "A Computational Theory of Vocabulary Acquisition", in L. Iwańska & S.C. Shapiro (eds.), *Natural Language Processing and Knowledge Representation* (Menlo Park, CA: AAAI): 347–375.

[22] Schagrin, M.L.; Rapaport, W.J.; & Dipert, R.R. (1985), *Logic: A Computer Approach* (New York: McGraw-Hill).

[23] Shapiro, S.C. (1982), "Generalized Augmented Transition Network Grammars for Generation from Semantic Networks," *Amer. J. Comp. Ling.* 8: 12–25.

[24] Shapiro, S.C. (1989), "The CASSIE Projects: An Approach to Natural Language Competence," in J.P. Martins & E.M. Morgado (eds.), *Proc. 4th Portuguese Conf. Artif. Intel.*, *Lecture Notes Artif. Intel.* 390 (Berlin: Springer-Verlag): 362–380.

[25] Shapiro, S.C. (1992), "Artificial Intelligence", in S.C. Shapiro (ed.), *Ency. Artif. Intel.* (New York: Wiley): 54–57.

[26] Shapiro, S.C. (2000), "SNePS: A Logic for Natural Language Understanding and Commonsense Reasoning", in L. Iwańska & S.C. Shapiro (eds.), *Natural Language Processing and Knowledge Representation* (Menlo Park, CA/Cambridge, MA: AAAI/MIT): 175–195.

[27] Shapiro, S.C., & Rapaport, W.J. (1987), "SNePS Considered as a Fully Intensional Propositional Semantic Network", in N. Cercone & G. McCalla (eds.), *The Knowledge Frontier* (New York: Springer-Verlag): 262–315.

[28] Shapiro, S.C., & Rapaport, W.J. (1991), "Models and Minds: Knowledge Representation for Natural-Language Competence", in R. Cummins & J. Pollock (eds.), *Philosophy and AI* (Cambridge, MA: MIT): 215–259.

[29] Shapiro, S.C., & Rapaport, W.J. (1992), "The SNePS Family", *Comp. & Math. with Applications* 23: 243–275.

[30] Shapiro, S.C., & Rapaport, W.J. (1995), "An Introduction to a Computational Reader of Narratives", in J.F. Duchan et al. (eds.), *Deixis in Narrative* (Hillsdale, NJ: Erlbaum): 79–105.

[31] Sternberg, R.J.; Powell, J.S.; & Kaye, D.B. (1983), "Teaching Vocabulary-Building Skills: A Contextual Approach", in A.C. Wilkinson (ed.), *Classroom Computers and Cognitive Science* (New York: Academic): 121–143.

[32] van Daalen-Kapteijns, M.M., & Elshout-Mohr, M. (1981), "The Acquisition of Word Meanings as a Cognitive Learning Process", *J. Verbal Learning & Verbal Behavior* 20: 386–399.

[33] *Webster's New Int'l. Dict. Engl. Lang.* (Merriam, 1937).

[34] Wesche, M., & Paribakht, T.S. (eds.) (1999), "Incidental L2 Vocabulary Acquisition: Theory, Current Research, and Instructional Implications", *Studies in Second Language Acquisition* 21(2).

[35] Zernik, U., & Dyer, M.G. (1987), "The Self-Extending Phrasal Lexicon", *Comp. Ling.* 13: 308–327.