

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

A Computational Theory of Vocabulary Expansion

Permalink

<https://escholarship.org/uc/item/6t0254gr>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 19(0)

Authors

Ehrlich, Karen

Rapaport, William J.

Publication Date

1997

Peer reviewed

A Computational Theory of Vocabulary Expansion

Karen Ehrlich and William J. Rapaport

Department of Computer Science and Center for Cognitive Science
State University of New York at Buffalo, Buffalo, NY 14260
{ehrllich|rapaport}@cs.buffalo.edu
<http://www.cs.buffalo.edu/~snwiz>

Abstract

As part of an interdisciplinary project to develop a computational cognitive model of a reader of narrative text, we are developing a computational theory of how natural-language-understanding systems can automatically expand their vocabulary by determining from context the meaning of words that are unknown, misunderstood, or used in a new sense. ‘Context’ includes surrounding text, grammatical information, and background knowledge, but no external sources. Our thesis is that the meaning of such a word *can* be determined from context, can be *revised* upon further encounters with the word, “*converges*” to a dictionary-like definition if enough context has been provided and there have been enough exposures to the word, and eventually “*settles down*” to a “steady state” that is always subject to revision upon further encounters with the word. The system is being implemented in the SNePS knowledge-representation and reasoning system.

The Project and Its Significance

We are developing a computational theory of how NLU systems (including humans) can automatically expand their vocabulary by determining from context the meaning of words that are unknown to the system, familiar but misunderstood, or used in a new sense (Ehrlich 1995). ‘Context’ includes surrounding text, grammatical information, and background knowledge, but no access to a dictionary (Zadrozny & Jensen 1991) or other external sources of information (including a human).

We take the meaning of a word (as understood by a cognitive agent) to be its position in a network of words, propositions, and other concepts (Quillian 1968, 1969). In this (idiolectic) sense, the meaning of a word for a cognitive agent is determined by idiosyncratic experience with it. The contextual meaning described above includes a word’s relation to every concept in the agent’s mind. Thus, the extreme interpretation of “meaning as context” defines every word in terms of every other word an agent knows. This is circular and too unwieldy for use. In another sense, the meaning of a word is its dictionary definition, usually containing less information. Thus, we limit the connections used for the definition by selecting particular *kinds* of information. Not all concepts within a given subnetwork are equally salient to a dictionary-style definition of a word. People abstract certain conventional information about words to use as a definition.

We claim that a meaning for a word *can* be determined from any context, can be *revised* and refined upon further encounters with it, and “*converges*” to a dictionary-like definition given enough context and exposures to it. Each encounter

with it yields a definition—a hypothesis about meaning. Subsequent encounters provide opportunities for *unsupervised* revision of this hypothesis, with no (human) “trainers” or “error-correction” techniques. The hypothesized definitions are *not* guaranteed to converge to a “correct” meaning (if such exists) but to one stable with respect to further encounters. Finally, no *domain-specific* background information is required for developing the definition.

Evidence for this can be seen in the psychological literature (below) and in informal protocols taken from subjects who reasoned out loud about their definition-forming and revision procedures when shown passages containing unknown words (Ehrlich 1995). These same passages served as input to a computational system that develops and revises definitions in ways similar to the human subjects.

The vocabulary-expansion system is part of an interdisciplinary project developing a computational cognitive model of a reader of narrative text (Duchan et al. 1995). To fully model a reader, it is important to model the ability to learn from reading, in particular, to expand one’s vocabulary in a natural way while reading, without having to stop to ask someone or to consult a dictionary. A complete lexicon cannot be manually encoded, nor could it contain new words or new meanings (Zernik & Dyer 1987). Text-understanding, message-processing, and information-extraction systems need to be robust in the presence of unknown expressions, especially systems using unconstrained input text and operating independently of human intervention, such as “intelligent agents”. E.g., a system designed to locate “interesting” news items from an online information server should not be limited to keyword searches—if the user is interested in news items about dogs, and the filter detects items about “brachets” (a term not in its lexicon), it should deliver those items as soon as it figures out that a brachet is a kind of dog.

Two features of our system mesh nicely with these desiderata, summarized as the advantages of learning over being told: (1) Being told requires human intervention. *Our system operates independently of a human teacher or trainer* (with one eliminable exception). (2) Learning is necessary, since one can’t predict all information needed to understand unconstrained, domain-independent text. *Our system does not constrain the subject matter (“domain”) of the text*. Although we are primarily concerned with narrative text, our techniques are general. Given an appropriate grammar, our algorithms produce domain-independent definitions, albeit ones dependent on the system’s background knowledge: The more background knowledge it has, the better its definitions will be, and

the more quickly they will “converge”. The system does *not* develop “correct” definitions, but *dictionary-like* definitions enabling her to continue understanding the text.

Our system, “Cassie”, consists of the SNePS-2.1 knowledge-representation and reasoning system (Shapiro 1979, Shapiro & Rapaport 1992), and a knowledge base (KB) representing Cassie’s background knowledge. Currently, the KB is hand-coded, since *how* she acquired this knowledge is irrelevant. Although we begin with a “toy” KB, each of our tests includes all previous information, so the KB grows as we test more words. Cassie’s input consists of (1) information from the text being read and (2) questions (e.g., “What does (word) mean?”) that trigger a deductive search of the KB. Output consists of a report of Cassie’s current definition of the word, or answers to other queries.

SNePS has an English lexicon, morphological analyzer/synthesizer, and a generalized ATN parser-generator that translates English input directly into a propositional semantic network without building an intermediate parse tree (Shapiro 1982, 1989; Rapaport 1988; Shapiro & Rapaport 1995). All information, including propositions, is represented by SNePS nodes; propositions about propositions can also be represented. Labeled arcs form the underlying syntactic structure of SNePS, embodied in the restriction that one cannot add an arc between two existing nodes, which would be tantamount to a proposition not represented by a node. Arc-paths can be defined for path-based inference, including property inheritance. Nodes and represented concepts are in 1–1 correspondence; this uniqueness principle guarantees that nodes are shared whenever possible and that nodes represent intensional objects (e.g., concepts, propositions, properties, and objects of thought including fictional entities, non-existents, and impossible objects; Shapiro & Rapaport 1991).

SNePS’s inference package accepts rules for deductive and default reasoning, allowing Cassie to infer “probable” conclusions in the absence of contrary information. When combinations of asserted propositions lead to a contradiction, the SNeBR belief-revision package allows Cassie to remove from the inconsistent context one or more of those propositions (Martins & Shapiro 1988). Once a premise is no longer asserted, the conclusions that depended on it are no longer asserted in that context. Cassie uses SNeBR to revise beliefs about the meanings of words. We have developed algorithms for partially automating the identification and removal or modification of a premise, based on SNePSwD, a default belief-revision system that enables automatic revision (Cravo & Martins 1993; Martins & Cravo 1991). We are exploring techniques for full automation.

Psychological Evidence

Psychological research on how humans store and access word meanings and expand their vocabularies once the basics of language have been acquired informs and supports our theory.

Johnson-Laird’s (1987) theory about the mental representation of words asserts that understanding a word (and a sentence containing a word) does not imply that one has a readily accessible definition of that word stored in one’s mind. Various aspects of a word’s meaning may be called to mind by sentences for which those aspects are relevant, without calling the entire definition to mind. He describes experiments

showing that responses to questions about specific aspects of a word’s meaning come faster following a priming sentence that calls that aspect to mind than to questions where the preceding sentence uses the word, but does not call that aspect to mind. The same speed-up occurs when the priming is a result of factual inference about a word’s referent as when it results from selectional restriction on the word’s sense.

Linguistic context has at least three effects on the interpretation of words. It can enable selection of an appropriate sense of a truly ambiguous word. It can lead to inference of a more specific referent than is strictly warranted by the meaning of an unambiguous word. It can call to mind particular aspects of a word’s meaning at the expense of other aspects. In each case, the mental representation of real or imaginary referents is important to the understanding of the word as used.

Some aspects of a word’s meaning are more salient to understanding its use than others. Since the evidence indicates that we do not retrieve definitions in their entirety when understanding sentences, we may not notice gaps in our lexical knowledge, so long as we can retrieve the aspects of meaning necessary to understanding the sentence. Such gaps point to the importance of the acquisition process. One can learn a word’s meaning by being told that meaning, or one can infer its meaning from encountering it in use.

Even fairly early in childhood, learning word meanings may be as much a matter of making inferences from linguistic context as of simple association. Johnson-Laird (1987) reports on experiments in which 3- and 4-year-old children listened to stories involving a novel verb. From hearing the verb used in sentences that contained other words they already understood, the children were able to perform selectional restriction on the arguments to the new verb. Children have also been shown to be able to learn aspects of the meanings of nonsense nouns from hearing them used in the context of familiar verbs.

Johnson-Laird suggests that lexical learning involves a sort of “bootstrapping” in which, once a fragment of language has been mapped onto a child’s internal representation of states of affairs in the world, other words are acquired either indirectly from context or directly from explicit definition. Words may also be of mixed acquisition; different individuals will acquire a given word differently. Some words can be completely lexically specified; others (e.g., natural-kind terms) cannot, but instead rely in part on a schema of default information based on a prototypical exemplar.

Johnson-Laird (1987: 579) summarizes parts of his theory of lexical meanings as follows: (1) Comprehension requires the listener to construct a model of the state of affairs described by the discourse. (2) There is a mental dictionary that contains entries in which the senses of words are represented. (But cf. Johnson-Laird 1987: 563.) (3) A lexical entry may be incomplete as a result of ignorance or because the word is a theoretical term with an intrinsically incomplete sense. (4) The senses of words can be acquired from definitions or from encountering instances of the word in use. (5) Corresponding to the method of acquisition, elements of a lexical representation can consist of (a) relations to other words, which could be represented by a mechanism akin to a semantic network, and (b) ineffable primitives that are used in constructing and manipulating mental models of the world.

Cassie understands narrative input by building mental rep-

representations of the information contained in the narrative; forming concepts of individuals, propositions, and events described; and connecting them with her prior knowledge. Her understanding of a concept in narrative is precisely that concept's connections to the rest of the narrative, together with its connections (if any) to previously acquired knowledge.

We adopt the idea that lexical entries have aspects of meaning connected to them in a semantic network, but do not have compiled, dictionary-style definitions permanently attached. Cassie selects salient features from her knowledge of a word when asked to define it, but does not permanently store those features as a definition. Our semantic network, however, includes meaning postulates (represented as rules) that Cassie can use as part (but not all) of her knowledge for producing definitions. Some information (e.g., about natural kinds) is expressed as default information. Thus, our approach is compatible with Johnson-Laird's theory and experiments.

Elshout-Mohr & van Daalen-Kapteijns (1987) treat verbal comprehension as a mental skill involving procedural and propositional knowledge. They hold that it is useful to have a "meaning unit" that is stable across contexts, so that new contexts may provide new information at the same time that an established understanding of a word allows interpretation in a new context.

In one experiment, they chose students with high and low verbal skills as measured by standard tests. Subjects were asked to think aloud as they tried to determine from context the meanings of invented words. The neologisms filled lexical gaps, so that they would refer to things with which the subjects were familiar, but would not have direct synonyms that could be substituted for them. (E.g., 'kolper': a window that transmits little light because of something outside it.) A series of sentences were presented, one per page, in which the new word was used. Subjects were aware of the need to construct a definition rather than search for a synonym. They were asked to report on new information gained from each sentence without reviewing previous pages. It was thought that this would tax their working memory, which was considered important, because most words learned from context are not learned in situations where word acquisition is the primary focus of cognition. Rather, one usually reads for the content of a text and acquires new vocabulary incidentally (if at all).

Most subjects tried to compare the unknown word with at least one familiar word. Those with high verbal skills used the model analytically, as a group of separable components that could be individually compared with further information. New information compatible with all facets of the model could be added; conflicting facets of the model could be replaced with new information. Those with lower verbal skills tended to use the model holistically, with new compatible information broadening or restricting the domain (Kiersey 1982) and incompatible information causing the rejection of the model (and perhaps the adoption of a different model).

According to the authors, in the task of word learning, a model: (1) provides a plan for knowledge retrieval; all aspects of the semantic unit of the model are accessible; (2) provides a frame-like structure for the meaning to be acquired, with certain slots to fill; (3) allows conventional aspects of definitions to steer abstraction toward similar conventions in the new definition; and (4) provides default information to fill

slots in the anticipated structure. The meaning of the new word is presumed to inherit aspects of meaning connected with the model unless otherwise specified.

Cassie uses her prior knowledge and the network that represents her understanding of the story (up to and including the sentence containing that word) to establish a definitional framework for the target word. This framework is not an analogical model in Elshout-Mohr & van Daalen-Kapteijns's sense, but does provide a plan for knowledge retrieval and a structure with certain slots to fill. The framework for nouns includes slots for synonyms and for hypernyms from which the target word can inherit aspects of meaning.

Because we wish our system to model an individual of high verbal ability (for general usefulness and specifically for use as a lexicographer's assistant), it does use its selected framework analytically, although new information may cause the system to select a different framework than that chosen after the first encounter with the target word.

Elshout-Mohr & van Daalen-Kapteijns suggest that children's verbal skills be developed by instruction in distinguishing between idiosyncratic experiences with a word and the more general experiences associated therewith, and in constraining the richness of individual experience by selecting a limited number of aspects of meaning. Especially important is the ability to select structural and functional aspects of a concept that are "in common with" and "in distinction of" other known concepts. Developing theory and technique for such movement from idiosyncratic understanding to conventional dictionary definitions is a central portion of our research.

Sternberg (1987) holds that three processes are applied in acquiring words from context: (1) distinguishing irrelevant from relevant information, (2) selectively combining relevant clues, and (3) comparing what has been selected and combined with previous knowledge. These processes operate on a basis of several types of cues: (1) temporal cues regarding the duration or frequency of X, or when X can occur; (2) spatial cues regarding the location of X or possible locations where X can sometimes be found; (3) value cues regarding the worth or desirability of X, or the kinds of affects X arouses; (4) stative descriptive cues regarding the properties of X (size, colour, etc.); (5) functional descriptive cues regarding possible purposes of X, actions X can perform, or potential uses of X; (6) cues regarding the possible causes of, or enabling conditions for, X; (7) cues regarding one or more classes to which X belongs, or other members of one or more classes of which X is a member; (8) equivalence cues regarding the meaning of X, or contrast (e.g., antonymy) to the meaning of X.

Sternberg's experiments show that readers who are trained in his three processes of recognizing relevant cues, combining cues, and comparing with known terms do better at defining new words than those who receive training only in the eight types of cues available, although any training in cue types produces better results than word memorization or no training.

Our system is designed to select certain information as most relevant, if it is present. For example, we follow Johnson-Laird and Elshout-Mohr & van Daalen-Kapteijns in emphasizing the structural and functional information about physical objects. If, however, such information is lacking, our system uses such other cue types as may be available. We also combine relevant information, and compare new words

with known words in certain specific ways, such as possible synonymy or hyponymy.

Algorithms

Our algorithms hypothesize and revise meanings for nouns and verbs that are unknown, misunderstood, or being used in a new way. Applying the principle that the meaning of a term is its location in the network of background and story information, our algorithms deductively search the network for information appropriate to a dictionary-like definition, assuming our grammar has identified the unknown word as a noun or a verb. The algorithms are presented in detail in Ehrlich 1995, and are illustrated by example here.

Cassie was provided with background information for understanding King Arthur stories (Malory 1470). When presented with a sequence of passages containing the unknown noun 'brachet', Cassie developed a theory that a brachet was a dog whose function is to hunt and that can bay and bite. (*Webster's* (1937) defines it as "a hound that hunts by the scent".) However, based on the first context in which the term appeared ("... there came a white hart running into the hall with a white brachet next to him, ..."), her initial hypothesis was that a brachet was a physical object that may be white. Each time 'brachet' 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, less salient "other properties" of particular brachets are not reported, although we do report 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 the psycholinguistic studies of the sort of vocabulary expansion we are modeling (discussed above). In the absence of any of this information, or in the presence of potentially inconsistent information (e.g., if the text says that one brachet hunts and another doesn't), 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.

In another test, Cassie was told that 'to smite' meant "to kill by hitting hard" (a mistaken belief held by the second author before reading Malory 1470). Passages in which various characters were smitten but then continued to act triggered SNeBR, which asks Cassie which of several possible "culprit" propositions in the KB to remove in order to block inconsistencies. Ideally, Cassie then decides which belief to revise. When humans encounter a discrepancy between a word's use and their previous understanding, they either assume the word is used incorrectly or decide that their previous understanding requires revision. When Cassie encounters a contradiction

derived from combining story information with background knowledge, she must decide which premises leading to the contradiction should be revised. To facilitate this, each assertion in the KB and story is tagged with a "knowledge category", ordered by certainty of belief. In case of contradiction, Cassie selects, from among the conflicting propositions, a proposition of greatest uncertainty as a candidate for revision. A set of rules for replacing discarded definitions with revised definitions is being developed. E.g., suppose the culprit were: If x smites y , then x hits y & y is dead & x hitting y causes y to be dead. On first encountering a smitee who survives, substitute these rules: (1) If x smites y , then x hits y & possibly y is dead; (2) If x smites y and y is dead, then x hitting y causes y to be dead. If asked for a definition of 'smite' now, Cassie will report that the result of smiting is that x hits y and possibly y is dead. The only human intervention is to tell Cassie to order her beliefs (she does the ordering, based on the knowledge categories, but has to be nudged to "do it now") and to tell Cassie to *assert* the revised belief she has *already automatically* generated from the lowest-ranked belief in the conflict set.

A third case is exemplified by 'to dress', which Cassie antecedently understood to mean "to put clothes on (something)", a well-entrenched meaning that 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. Cassie decides that to dress is *either* to put clothes on *or* to prepare for battle.

An Example

Here, we sketch Cassie's handling of 'smite', with this background information in the KB: *There is a king named King Arthur. There is a king named King Lot. There is a sword named Excalibur. Excalibur is King Arthur's sword. Horses are animals. Kings are persons. Knights are persons. Dukes are persons. "Person" is a basic-level category. "Horse" is a basic-level category. "Before" and "after" are transitive relations. If x is dead at time t , x can perform no actions at t or at any subsequent time. If x belongs to a subclass of person, x is a person. If a person acts, the act performed is an action. If an agent acts on an object, and there is an indirect object of the action, then the action is bitransitive. If an agent acts on an object, then the action is transitive. If an agent acts on itself, then the action is reflexive. If x is hurt at time t , then x is not dead at t . If x is not dead at time t , then x was not dead at any prior time. If x smites y at time t , then x hits y at t , and y is dead at t , and the hitting caused the death.* (Note that the last is the only information about 'smite' in the knowledge base.)

Cassie is then given a sequence of passages containing 'smite' (adapted from Malory 1470: 13ff) interspersed with questions and requests for definitions:

Passage1: King Arthur turned himself and his horse. He smote before and behind. His horse was slain. King Lot smote down King Arthur.

Definition1: A person can smite a person. If x smites y at time t , then x hits y at t , and y is dead at t .

Question1: What properties does King Arthur have?

Reply1: King Arthur is dead.

Passage2: King Arthur's knights rescued him. They sat him

on a horse. He drew Excalibur.

Question2: When did King Arthur draw [Excalibur]?

The inference required to reply to **Question2** triggers SNeBR, which reports that King Arthur's drawing (i.e., acting) is inconsistent with his being dead. Cassie automatically removes the proposition reporting her belief that smiting entails killing, which is replaced with two beliefs: that although smiting entails hitting, it only possibly entails killing, and that if smiting results in a death, then the hitting is the cause of death. These rules are *not* built in; they are *inferred* by revision rules.

Definition2: A person can smite a person. If x smites y at time t , then x hits y at t and possibly y is dead at t .

Passage3: Two of King Claudas's knights rode toward a passage. Sir Ulfyas and Sir Brastias rode ahead. Sir Ulfyas smote down one of King Claudas's two knights. Sir Brastias smote down the other knight. Sir Ulfyas and Sir Brastias rode ahead. Sir Ulfyas fought and unhorsed another of Claudas's knights. Sir Brastias fought and unhorsed the last of Claudas's knights. Sir Ulfyas and Sir Brastias laid King Claudas's last two knights on the ground. All of King Claudas's knights were hurt and bruised.

The information that the knights were hurt was added in forward-chaining mode to allow Cassie to notice that that they were still alive at the time they were hurt and therefore could not have died earlier at the time they were smitten. Cassie has now heard of two cases in a row (King Arthur, and the two knights) where a smitee has survived being smitten, with no intervening cases of death by smiting, yielding:

Definition3: A person can smite a person. If x smites y at time t , then x hits y at t .

Further encounters with 'smite' cause no further revisions. The definition has stabilized ("converged") in a manner similar to our human protocols.

Related Work

Classic work along these lines includes Granger 1977, Berwick 1983, Haas & Hendrix 1983. Here, we review more recent work.

Zernik & Dyer (1987) compile definitions of words and figurative phrases from conversation into a hierarchical lexicon, using a pattern constructor that analyzes parsing failures to modify its patterns and a concept constructor that selects a strategy according to background information about the goals and plans a person is likely to have in various situations. If the first interpretation of a phrase is inconsistent with that information, the system queries the user, suggesting various interpretations more consistent with the goals of persons in the story until the user confirms that a correct interpretation has been reached. However, since we focus on literary, not conversational, discourse, Cassie is not informed by a human user when she misunderstands. As long as the misunderstanding is compatible with further encounters with the word and with her general knowledge, there is no reason for Cassie to revise her understanding. If further reading leads to the conclusion that a previous definition was wrong, Cassie revises her understanding without explicit instruction.

Hastings (1994; Hastings & Lytinen 1994ab) presents several versions of a system, Camille, that uses knowledge of a

given domain to infer a word's meaning. Hastings's approach is like ours: The goal is to allow the system to read and acquire word meanings without the intervention of a human tutor. His approach differs, however, in the types of information sought as the meaning of a word, and in the nature of the KB. For each domain, the initial KB consisted of a taxonomy of relevant objects and actions. Camille attempts to place the unknown word appropriately in the domain hierarchy. To this basic system, Hastings has added: a mutual exclusivity constraint; a script applier allowing Camille to match the unknown word with a known word that has filled the same slot in a particular script, or, for a verb, with a known word whose arguments match those of the target word; and an ability to recognize the existence of multiple senses for a word. In most instances, the meaning sought appears to be synonymy with a known word, unlike Cassie, which can create new concepts (defined in terms of preexisting ones). In one version, however, Camille is given the capacity to create a new node, and insert it into the domain hierarchy. This, however, is only available for unknown nouns. Verbs can be "defined" only by mapping them to their closest synonyms. Hastings's evaluation of Camille's performance is given in terms of "correctness" of word meaning. His focus is on the comparative precision and accuracy of the various versions of Camille as they attempt to map unknown terms onto known nodes. For us, such a notion of "correctness" does not apply.

Siskind's (1996) focus is on first-language acquisition during childhood, whereas ours is on mature cognitive agents who already know a large part of their language and are (merely) expanding their vocabulary. Siskind takes as given (1) an utterance, (2) a simultaneous visual perception, (3) a mental representation of the situation perceived, which is caused by it, and (4) an assumption that the utterance means that mental representation. His algorithms assign parts of the mental-representation meaning to parts (words) of the utterance. Given "multiple situations," these algorithms "converge" to "correct word-to-meaning mappings". Although we also assume an utterance and a mental representation that the utterance means, Cassie does not use visual perception to produce the mental representation. In most cases of reading, any mental representation (including imagery) would be produced by the text, so visual perception of a real-world situation does not arise, except for illustrated texts. Although Cassie does not use illustrations, she could in principle (Srihari 1991, Srihari & Rapaport 1989). Siskind's system begins with a mapping between a whole meaning and a whole utterance, and infers mappings between their parts. Cassie already has both of those mappings and seeks to infer definition-style relations between the unknown word and the rest of the KB. Moreover, it does not make sense in our theory to speak of "correct word-to-meaning mappings". Finally, Siskind claims that his theory provides evidence for "semantic bootstrapping"—using semantics to aid in learning syntax. In contrast, Cassie uses *syntactic* bootstrapping (using syntax to aid in learning semantics (Gleitman 1990, 1994)), which seems more reasonable for our situation.

References

Berwick, R.C. (1983), "Learning Word Meanings from Examples," *IJCAI-83*: 459–461.

- Cravo, M.R., & Martins, J.P. (1993), "SNePSwD: A New-comer to the SNePS Family", *JETA1* 5: 135–148.
- Duchan, J.F.; Bruder, G.A.; & Hewitt, L.E. (eds.) (1995), *Deixis in Narrative: A Cognitive Science Perspective* (Erlbaum).
- Ehrlich, K. (1995), "Automatic Vocabulary Expansion through Narrative Context," *TR 95-09* (Buffalo: SUNY Buffalo Dept. of Computer Science).
- Elshout-Mohr, M., & van Daalen-Kapteijns, M.M. (1987), "Cognitive Processes in Learning Word Meanings," in M. G. McKeown & M. E. Curtis (eds.), *The Nature of Vocabulary Acquisition* (Erlbaum): 53–71.
- Gleitman, L. (1990), "The Structural Sources of Verb Meanings," *Lang. Acquisition* 1: 1–55.
- Gleitman, L. (1994), "A Picture is Worth a Thousand Words—But That's the Problem" (talk presented at the 1st Int'l. Summer Inst. in Cog. Sci. (SUNY Buffalo, July 1994); abstract in *Proc. 16th Annual Conf. Cog. Sci. Soc.*: 965).
- Granger, R.H. (1977), "Foul-Up: A Program that Figures Out Meanings of Words from Context," *IJCAI-77*: 67–68.
- Haas, N., & Hendrix, G. (1983), "Learning by Being Told: Acquiring Knowledge for Information Management," in Michalski et al. (eds.), *Machine Learning* (Tioga): 405–428.
- Hastings, P.M. (1994), "Automatic Acquisition of Word Meaning from Context," Ph.D. diss. (Comp. Sci. & Eng'g., Univ. of Michigan).
- Hastings, P.M., & Lytinen, S.L. (1994a), "The Ups and Downs of Lexical Acquisition," *AAAI-94*: 754–759.
- Hastings, P.M., & Lytinen, S.L. (1994b), "Objects, Actions, Nouns, and Verbs," *Proc. 16th Annual Conf. Cog. Sci. Soc.*: 397–402.
- Johnson-Laird, P.N. (1987), "The Mental Representation of the Meanings of Words," in A.I. Goldman (ed.), *Readings in Philosophy and Cognitive Science* (MIT, 1993): 561–583.
- Kiersey, D.M. (1982), "Word Learning with Hierarchy Guided Inference," *AAAI-82*: 172–178.
- Malory, Sir T. (1470), *Le Morte Darthur*, ed. R.M. Lumiansky (Collier, 1982).
- Martins, J., & Cravo M.R. (1991), "How to Change Your Mind," *Noûs* 25: 537–551.
- Martins, J., & Shapiro, S.C. (1988), "A Model for Belief Revision," *Artif. Intell.* 35: 25–79.
- Quillian, M.R. (1968), "Semantic Memory," in M. Minsky (ed.) *Semantic Information Processing* (MIT): 227–270.
- Quillian, M.R. (1969), "The Teachable Language Comprehender: A Simulation Program and Theory of Language," *Comm. ACM* 12: 459–476.
- Rapaport, W.J. (1988), "Syntactic Semantics: Foundations of Computational Natural-Language Understanding," in J.H. Fetzer (ed.), *Aspects of Artificial Intelligence* (Kluwer): 81–131.
- Shapiro, S.C. (1979), "The SNePS Semantic Network Processing System," in N. Findler (ed.), *Associative Networks* (Academic): 179–203.
- Shapiro, S.C. (1982), "Generalized Augmented Transition Network Grammars for Generation from Semantic Networks," *American J. Comp. Ling.* 8: 12–25.
- Shapiro, S.C. (1989), "The CASSIE Projects: An Approach to Natural Language Competence," in J.P. Martins & E.M. Morgado (eds.), *Proc. 4th Portugese Conf. A.I. (EPIA-89)*, Lecture Notes in A.I. 390 (Springer-Verlag): 362–380.
- 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* (MIT): 215–259.
- Shapiro, S.C., & Rapaport, W.J. (1992), "The SNePS Family," *Computers and Mathematics with Applications* 23: 243–275.
- Shapiro, S.C., & Rapaport, W.J. (1995), "An Introduction to a Computational Reader of Narrative," in Duchan et al. 1995: 79–105.
- Siskind, J.M. (1996), "A Computational Study of Cross-Situational Techniques for Learning Word-to-Meaning Mappings," *Cognition* 61: 39–91.
- Srihari, R.K. (1991), "PICTION: A System that Uses Captions to Label Human Faces in Newspaper Photographs," *AAAI-91*: 80–85.
- Srihari, R.K. & Rapaport, W.J. (1989), "Extracting Visual Information from Text: Using Captions to Label Human Faces in Newspaper Photographs," *Proc. 11th Annual Conf. Cog. Sci. Soc.*: 364–371.
- Sternberg, R.J. (1987), "Most Vocabulary is Learned from Context," in M.G. McKeown & M.E. Curtis (eds.), *The Nature of Vocabulary Acquisition* (Erlbaum): 89–105.
- Webster's New Int'l. Dictionary of the English Language (2/e, Unabridged)* (Merriam, 1937).
- Zadrozny, W., & Jensen, K. (1991), "Semantics of Paragraphs," *Comp. Ling.* 17: 171–209.
- Zernik, U., & Dyer, M.G. (1987), "The Self-Extending Phrasal Lexicon," *Comp. Ling.* 13: 308–327.