Symbol-Anchoring in Cassie

Stuart C. Shapiro and Haythem O. Ismail*
Department of Computer Science and Engineering
and Center for Cognitive Science
201 Bell Hall
University at Buffalo, The State University of New York
Buffalo, NY 14260-2000
{shapiro | hismail}@cse.buffalo.edu

Abstract

We have been engaged in a series of projects in which Cassie, the SNePS cognitive agent, has been incorporated into a hardware- or software-simulated cognitive robot. In this paper, we present an informal summary of our approach to anchoring the abstract symbolic terms that denote Cassie's mental entities in the lower-level structures used by embodied-Cassie to operate in the real (or simulated) world. We discuss anchoring in the domains of: perceivable entities and properties, actions, time, and language.

Introduction

We have been engaged in a series of projects in which Cassie, the SNePS cognitive agent (Shapiro & Rapaport 1987; Shapiro 1989; Shapiro & Rapaport 1991; 1992; Shapiro 1998; Shapiro & The SNePS Implementation Group 1999; Shapiro, Ismail, & Santore 2000; Ismail & Shapiro 1999; 2000a; 2000b), has been incorporated into a hardware- or software-simulated cognitive robot. The capabilities of the embodied Cassie have included: input and output in fragments of English; reasoning; performance of primitive and composite acts; motion; and vision. In this paper, we discuss the ways in which we have anchored the symbols in Cassie's reasoning level.

We use **GLAIR** (Grounded Layered Architecture with Integrated Reasoning) (Hexmoor, Lammens, & Shapiro 1993; Hexmoor & Shapiro 1997) as the architecture of our cognitive robot. GLAIR consists of three levels: the Knowledge Level (KL), the Perceptuo-Motor Level (PML), and the Sensori-Actuator Level (SAL):

1. **The Knowledge Level (KL):** The level at which conscious reasoning takes place. The KL is implemented by the SNePS system (Shapiro & Rapaport 1987; 1992; Shapiro & The SNePS Implementation Group 1999), where SNeRE (the SNePS Rational Engine) (Kumar & Shapiro 1994a; 1994b; Kumar 1994; 1996) is used for initiating and controlling the execution of acts.
2. **The Perceptuo-Motor Level (PML):** The level at which routines for carrying out primitive acts are located. This is also the location for other subconscious activities that allow for Cassie's consciousness of its body and surroundings. The PML has been implemented in Common Lisp and other languages, depending on the implementation of the hardware or software-simulated robot.
3. **The Sensori-Actuator Level (SAL):** The level controlling the operation of sensors and actuators (being either hardware or simulated). The SAL has been implemented in C and other languages, depending on the implementation of the hardware or software-simulated robot.

In this paper, we shall not discuss the SAL, nor the division of responsibilities between the PML and the SAL, but will refer to the PML as our implementation of both the subsymbolic level and of the symbolic, but bodily and non-conscious level, the level containing what the call for this symposium refers to as the "physical-level representations of objects."

We refer to the KL as the "conscious" level, since that is the locus of symbols accessable to reasoning and to natural language interaction. It is the level containing what the call for this symposium refers to as the "abstract-level representations of objects."

Atomic symbols in the KL are terms of the SNePS logic (Shapiro 2000). Symbol structures in the KL are functional terms in the same logic (Shapiro 1993; 2000). All terms denote mental entities (Shapiro & Rapaport 1987). For example, Cassie's conceptualization of George W. Bush may differ in multiple respects from the real Bush, and also from Al Gore's conceptualization of Bush. The KL term for Bush denotes Cassie's conceptualization of Bush—the mental entity that is Bush for Cassie. In conversation, Cassie will mean by "Bush" her conceptualization of Bush, and will normally take someone else's use of expressions that refer to Bush to refer to her own conceptualization of Bush. Nevertheless, it is possible for Cassie to be sensitive enough to have a different term to denote her conceptualization of Gore's conceptualization of Bush, and to use that term to be the referent of expressions she hears

*Current address: Department of Engineering Mathematics and Physics, Cairo University, Giza, Egypt

Gore use.

Since KL terms denote mental entities, there is a 1-1 correspondence between KL terms and those mental entities Cassie has so far conceived of. For example, if Cassie believed that the Morning Star is the Evening Star, she would have one KL term to denote the Morning Star, another KL term to denote the Evening Star, and a third KL term to denote the proposition that the two have the same extension. When Cassie conceives of a mental entity she does not recognize (re-cognize) as one for which she already has a KL term, she creates a new KL term to denote it (Maida & Shapiro 1982).

The topic of this paper is our approach to anchoring the KL terms that denote Cassie's mental entities in the PML structures used by embodied-Cassie to operate in the real world. Briefly, our theoretical stance is that KL terms are accessible to natural language interaction and to reasoning—Cassie can discuss and reason about the entities they denote. PML structures, on the other hand, are accessable to sensors and effectors, but not to natural language interaction or reasoning. Anchoring is achieved by associating (we use the term "aligning") a KL term with a PML structure, thereby allowing Cassie to recognize entities and perform actions, but not to discuss or reason about the low-level recognition or performance. In the following sections, we discuss anchoring in the domains of: perceivable entities and properties, actions, time, and language. This paper has deliberately been kept as an informal summary of our approach. For more details, and more formal presentations, see the papers cited herein.

Perceivable Entities

There are KL terms for every mental entity Cassie has conceived of, including individual entities, categories of entities, colors, shapes, and other properties of entities.

We assume there are PML structures for features of the perceivable world that Cassie's perceptual apparatus can detect and distinguish. For example, each distinguishable color might be represented by a single number, by a symbol, or by a triple of RGB values, and each distinguishable shape might be represented by a single number or symbol. We further assume that each particular perceived object will be represented at this level by a collection of such structures. Without loss of generality (we believe), we will assume that these collections of PML structures can be considered to be feature vectors. (We have used feature vectors in versions of embodied-Cassie we have implemented.) A feature vector is a sequence of values, $(v_1, \ldots, v_n)$, where each v_i is a possible value of some dimension D_i. A feature vector can also be interpreted as a point in an $n-$dimensional feature space, where the i^{th} dimension is D_i. What dimensions are used depends on the perceptual apparatus of the robot.

Our approach to grounding KL terms for perceivable entities, categories, and properties is to align a KL term with a PML feature vector, possibly with unfilled (null) components. For example, one version of embodied Cassie used a two-dimensional feature vector in which the dimensions were color and shape. The KL term denoting Cassie's idea of blue was aligned with a feature vector whose color component was the PML structure the vision system used when it detected blue in the visual field, but whose shape component was null. The KL term denoting cubes was aligned with a feature vector whose shape component was the PML structure the vision system used when it detected a cube in the visual field, but whose color component was null. We have implemented alignment in various ways, including association lists, hash tables, and property lists.

Let us call a feature vector with some null components an "incomplete feature vector", and one with no null components a "complete feature vector".

KL terms denoting perceivable properties and KL terms denoting recognizable categories of entities are typically aligned with incomplete feature vectors. Examples include the terms for blue and for cubes mentioned above, and may also include terms for tall, fat, bearded, man, and woman. These terms may be combined into verbal descriptions, such as "a tall, fat, bearded man," that may be used to perceptually recognize the entity so described. In this paper, we will not use the term "description" to mean a verbal description that *cannot* be used for perceptual recognition, such as "a college-educated businessman who lives in Amherst, NY."

Cassie might have a KL term for an entity about which she knows no descriptive terms. For example, all she might believe about Fred is that he is a college-educated businessman who lives in Amherst, NY. Thus, she would be incapable of describing Fred (the way we are using "describe"). Nevertheless, it might be the case that Cassie's term denoting Fred is aligned with a complete feature vector. In this case, Cassie would be able to recognize Fred, though not describe him. We call such a complete feature vector aligned with an entity, the entity's "PML-description." To emphasize its completeness, we can use the term "complete PML-description."

A complete PML-description may be assembled for an entity by unifying the incomplete PML-descriptions of its known (conceived of) properties and categories. For example, if Cassie knows nothing about Harry, and we tell her that Harry is a tall, fat, bearded man, she would be able to assemble a PML-description of Harry and recognize him on the street. (Assuming that Cassie's terms for tall, fat, bearded, and man are aligned with incomplete feature vectors.) In some cases, this might result in a set of several complete PML-descriptions. For example, the PML-descriptions of some, but not a particular, red chair might include feature vectors with different shape components. Once a PML-description is assembled for an entity, it can be cached by aligning the entity directly with it. Afterwards, Cassie could recognize the entity without thinking about what it looks like.

To find (come to be looking at) an entity, Cassie finds a PML-description of the entity that is as complete as possible, and directs her perceptual apparatus to do what is necessary to cause an object satisfying that feature vector to be in her visual field. (See the section on actions, below, for a description of how actions are grounded.)

If Cassie is looking at some object, she can recognize it if its feature vector is the PML-description of some entity she has already conceived of. If there is no such entity, Cassie can create a new KL term to denote this new entity, and believe of it that it has those properties and is a member of those categories whose partial PML-descriptions are parts of the PML-description of the new entity.

If there are multiple entities whose PML-descriptions match the feature vector, disambiguation is needed, or Cassie might simply not know which one of the entities she is looking at.

We are currently investigating the issue of when Cassie might decide that the object she is looking at is new, even though it looks exactly like another she has already conceived of.

We have not worked on the problem of recognizing an entity by context. For example a store clerk might be recognized as any person standing behind a cash register.[1] We speculate that this problem requires a combination of KL knowledge and KL-PML alignment. Knowing that a person standing behind a cash register is a clerk is KL knowledge. Recognizing a person, a cash register, and the "behind" relation requires KL-PML alignment.

Example

This subsection contains an example interaction with a simulated version of Cassie as a "Foveal Extra-Vehicular Activity Helper-Retriever (FEVAHR)" (Shapiro 1998). Cassie, the FEVAHR, was implemented on a commercial Nomad robot, including sonar, bumpers, and wheels, enhanced with a foveal vision system consisting of a pair of cameras with associated hardware and software. The simulated version allows interactions to be shown in a paper.

In this simulation, both the PML and the simulated world are implemented in Common Lisp. The PML feature vector has two dimensions, called "color" and "shape". There are seven objects in the simulated world. The Common Lisp symbols that represent these objects and their PML-descriptions are shown in the following table:

Object	**Color**	**Shape**
`WORLD:BILL`	`BLUE`	`SQUARE`
`WORLD:STU`	`YELLOW`	`SQUARE`
`WORLD:CASSIE`	`CYAN`	`CIRCLE`
`WORLD:GREENIE`	`GREEN`	`CIRCLE`
`WORLD:REDROB-1`	`RED`	`CIRCLE`
`WORLD:REDROB-2`	`RED`	`CIRCLE`
`WORLD:REDROB-3`	`RED`	`CIRCLE`

The KL terms that are aligned with feature vectors are shown in the following table:

KL Term	**Color**	**Shape**
`B1`	`CYAN`	`CIRCLE`
`B5`	`YELLOW`	`SQUARE`
`B6`	`BLUE`	`SQUARE`
`green`	`GREEN`	`NIL`
`red`	`RED`	`NIL`
`person`	`NIL`	`SQUARE`
`robot`	`NIL`	`CIRCLE`

Notice that `B1`, `B5`, and `B6` are aligned with complete feature vectors, while `green`, `red`, `person`, and `robot` are aligned with incomplete feature vectors. `B1`, `B5`, and `B6` denote individuals. `green` and `red` denote properties. `person` and `robot` denote categories.

Cassie's relevant beliefs about the entities denoted by these terms may be glossed as:

B1's name is Cassie.
B5's name is Stu.
B6's name is Bill.
Cassie is a FEVAHR.
FEVAHRs are robots.
Bill and Stu are people.
Robbie is a green robot.
B8, B9, and B10 are red robots.
People and robots are things.

The only descriptive terms Cassie has for Bill and Stu are that they are people, and the only descriptive term she has for herself is that she is a robot. Nevertheless, Bill, Stu, and Cassie are aligned with complete feature vectors, so she has PML-descriptions for them, and, therefore, can recognize them. On the other hand, Robbie, `B8`, `B9`, and `B10` are not aligned with any feature vectors, but PML-descriptions can be assembled for them from their properties and categories.

Following is an interaction with Cassie about these entities. Sentences preceded by "`:`" are human inputs. Sentences starting with "`The FEVAHR is`" are reports of simulated actions and perceptions, and are not output by the hardware version. Sentences beginning with "`I`" are generated by Cassie. At the beginning of the interaction, Cassie is looking at listening to, and talking to Stu.

[1]This example was suggested by one of the anonymous reviewers.

```
 : Find a robot.
The FEVAHR is looking for something
    that's a CIRCLE.
The FEVAHR found WORLD:REDROB-1.
 I found a red robot.
The FEVAHR is looking at WORLD:REDROB-1.
 I am looking at a red robot.

 : Find a person.
The FEVAHR is looking for something
    that's a SQUARE.
The FEVAHR found WORLD:STU.
 I found you, Stu.
The FEVAHR is looking at WORLD:STU.
 I am looking at you.

 : Find a green thing.
The FEVAHR is looking for something
    that's GREEN.
The FEVAHR found WORLD:GREENIE.
 I found Robbie.
The FEVAHR is looking at WORLD:GREENIE.
 I am looking at Robbie.

 : Find Bill.
The FEVAHR is looking for something
    that's BLUE and a SQUARE.
The FEVAHR found WORLD:BILL.
 I found Bill.
The FEVAHR is looking at WORLD:BILL.
 I am looking at Bill.
```

Attentional Structures

Cassie's attentional apparatus is modeled by a set of PML registers (variables), each of which can hold one or more KL terms or PML structures. Some of these registers derive from the theory of the Deictic Center (Duchan, Bruder, & Hewitt 1995), and include: `I`, the register that holds the KL term denoting Cassie herself, and `NOW`, the register that holds the KL term denoting the current time. Others, termed "modality registers", model Cassie's knowledge of what she is doing. For example, if one of Cassie's modalities were speech, and she were currently talking to Stu, her `SPEECH` register would contain the KL term denoting the state of Cassie's talking to Stu. In many cases, a single modality of an agent can be occupied by only one activity at a time. In that case the modality register for that modality would be constrained to contain only one term at a time.

One of the modality registers we have used is one to keep track of what Cassie is looking at. When she recognizes an object in her visual field, the KL term denoting the state of looking at the recognized entity is placed in the register, and is removed when the object is no longer in the visual field. If one assumes that Cassie can be looking at several objects at once, this register would be allowed to contain several terms.

If asked to look at or find something that is already in her visual field, Cassie recognizes that fact, and doesn't need to do anything. The following interaction continues from the previous one:

```
 : Look at Robbie.
The FEVAHR is looking for something
    that's GREEN and a CIRCLE.
The FEVAHR found WORLD:GREENIE.
 I found Robbie.
The FEVAHR is looking at WORLD:GREENIE.
 I am looking at Robbie.

 : Find a robot.
 I am looking at Robbie.
```

Comparing Cassie's response to the second request with her response to the previous requests, one can see that she realized that she was already looking at a robot, and so didn't need to do anything to find one.

Actions

Some KL terms denote primitive actions that Cassie can perform. We term an action, along with the entity or entities it is performed on, to be an "act". For example, the act of going to Bill consists of the action of going and the object Bill. Acts are denoted by KL functional terms.

Each KL action term that denotes a primitive action that Cassie can perform is aligned with a procedure in the PML. The procedure takes as arguments the KL terms for the arguments of the act Cassie is to perform. For example, if Cassie is to perform the act of going to Bill, the PML going-procedure would be called on the KL Bill-term. It would then find the PML-description of Bill, and cause the robot hardware to go to an object in the world that satisfies that description (or cause the robot simulation to simulate that behavior). The PML going-procedure would also insert the KL term denoting the state of Cassie's going to Bill into the relevant modality register(s), which, when `NOW` moves (see below), would cause an appropriate proposition to be inserted into Cassie's belief space.

Acts whose actions are primitive are considered to be primitive acts. Composite acts are composed of primitive "control actions", and their arguments, which, themselves are primitive or composite acts. Control actions include sequence, selection, iteration, and nondeterministic choice (Kumar & Shapiro 1994a; 1994b; Kumar 1994; 1996; Shapiro & The SNePS Implementation Group 1999; Ismail & Shapiro 1999). There are also propositions for act preconditions, goals, effects, and for plans (what some call recipes) for carrying out non-primitive acts.

In the interactions shown above, sentences starting with "`The FEVAHR is`" were printed by the simulated action function which was called by the PML procedure aligned with the KL term for finding something. When Cassie was asked to look at Robbie, she did so by finding Robbie, because there is a KL belief that the plan for carrying out the non-primitive act of looking at something is to find that thing.

Time

As mentioned above, the attentional NOW register always contains the KL term denoting the current time (Shapiro 1998; Ismail & Shapiro 2000b; 2001). Actually, since "now" is vague (it could mean this minute, this day, this year, this century, etc.), NOW is considered to include the entire semi-lattice of times that include the smallest current now-interval Cassie has conceived of, as well as all other times containing that interval.

NOW moves whenever Cassie becomes aware of a new state. Some of the circumstances that cause her to become aware of a new state are: she acts; she observes a state holding; she is informed of a state that holds. NOW moves by Cassie's conceiving of a new smallest current now-interval (a new KL term is introduced with that denotation), and NOW is changed to contain that time. The other times in the old NOW are defeasibly extended into the new one by adding propositions asserting that the new NOW is a subinterval of them.

Whenever Cassie acts, the modality registers change (see above), and NOW moves. The times of the state(s) newly added to the modality registers are included in the new NOW semi-lattice, and the times of the state(s) deleted from the modality registers are placed into the past by adding propositions that assert that they precede the new NOW.

The following interaction, following the ones shown above, shows an action of Cassie's passing from the present into the past:

```
 : Who have you talked to?
 I am talking to you.

 : Talk to Bill.
The FEVAHR is starting to talk to Bill.
 I am talking to you, Bill.

 : Who have you talked to?
 I talked to Stu
and I am talking to you.
```

Temporal Durations

To give Cassie a "feel" for the amount of time that has passed, she has a COUNT register acting as an internal pacemaker. The value of COUNT is a non-negative integer, incremented at regular intervals. Whenever NOW moves, the following happens:

1. the old now-interval t_o is aligned with the current value of COUNT, grounding it in a PML-measure of its duration;
2. the value of COUNT is quantized into a value δ which is the nearest half-order of magnitude (Hobbs 2000) to COUNT, providing an equivalence class of PML-measures that are not noticeably different;
3. a KL term d, aligned with δ, is found or created, providing a mental entity denoting each class of durations;
4. a belief is introduced into the KL that the duration of t_o is d, so that Cassie can have beliefs that two different states occurred for about the same length of time;
5. COUNT is reset to 0, to prepare for measuring the new now-interval.

Language

Cassie interacts with humans in a fragment of English. Although it is possible to represent all her linguistic knowledge in the KL, use reasoning to analyze input utterances (Shapiro & Neal 1982; Neal & Shapiro 1985; 1987a; 1987b), and use the acting system to generate utterances (Haller 1996; 1999), we do not currently do this. Instead, the parsing and generation grammars, as well as the lexicon, are at the PML. (See, e.g. (Shapiro 1982; Shapiro & Rapaport 1995; Rapaport, Shapiro, & Wiebe 1997).) There are KL terms for lexemes, and these are aligned with lexemes in the PML lexicon. We most frequently use a KL unary functional term to denote the concept expressed by a given lexeme, but this does not allow for polysemy, so we have occasionally used binary propositions that assert that some concept may be expressed by some lexeme.

This facility was used for Cassie to understand the human inputs shown in the example interactions in this paper, and for her to generate her responses (the sentences beginning with "I"). We can also use the low level surface function to see the NL expression Cassie would use to express the denotation of various SNePS terms[2]:

```
* (surface B1)
 Cassie

* (surface B5)
 Stu

* (surface B6)
 you
```

(Remember, Cassie is currently talking to Bill.)

Summary

We have given an informal summary of our approach to connecting the abstract-level representations to the physical-level representations of Cassie, our cognitive robot. The abstract-level representations are terms of SNePS logic contained in the Knowledge Level (KL) of our GLAIR agent architecture, while the physical-level representations are feature vectors, procedures, and other symbol structures contained at the Perceptuo-Motor Level (PML) of the architecture.

KL terms denoting perceivable entities, perceivable properties, and recognizable categories are aligned with PML feature vectors. Primitive actions are aligned with

[2]The prompt for this Lispish interaction level is "*".

PML procedures. Deictic and modality registers hold KL terms for individuals and states that Cassie is currently aware of, including states of her own body. They are updated by the PML procedures. The `NOW` register is used to give Cassie a personal sense of time, including keeping track of current and past states. KL terms denoting times and temporal durations are aligned with PML numeric measures of durations created by the PML pacemaker. Lexemes are represented by KL terms that are aligned with PML lexicon entries used by the parsing and generation grammars, which, like PML procedures, mediate between Cassie and the outside world, in this case, humans with which she communicates.

Acknowledgments

SNePS and Cassie are products of SNeRG, the SNePS Research Group, Department of Computer Science and Engineering, University at Buffalo. Many group members, past and present, have contributed to these efforts. We appreciate the valuable comments on previous drafts of this paper made by William J. Rapaport and other current members of SNeRG, and by two anonymous reviewers.

References

Brachman, R. J., and Levesque, H. J., eds. 1985. *Readings in Knowledge Representation.* San Mateo, CA: Morgan Kaufmann.

Duchan, J. F.; Bruder, G. A.; and Hewitt, L. E., eds. 1995. *Deixis in Narrative: A Cognitive Science Perspective.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Haller, S. 1996. Planning text about plans interactively. *International Journal of Expert Systems* 9(1):85–112.

Haller, S. 1999. An introduction to interactive discourse processing from the perspective of plan recognition and text planning. *Artificial Intelligence Review* 13(4):259–333.

Hexmoor, H., and Shapiro, S. C. 1997. Integrating skill and knowledge in expert agents. In Feltovich, P. J.; Ford, K. M.; and Hoffman, R. R., eds., *Expertise in Context.* Cambridge, MA: AAAI Press/MIT Press. 383–404.

Hexmoor, H.; Lammens, J.; and Shapiro, S. C. 1993. Embodiment in GLAIR: a grounded layered architecture with integrated reasoning for autonomous agents. In Dankel II, D. D., and Stewman, J., eds., *Proceedings of The Sixth Florida AI Research Symposium (FLAIRS 93).* The Florida AI Research Society. 325–329.

Hobbs, J. R. 2000. Half orders of magnitude. In Obrst, L., and Mani, I., eds., *Papers from the Workshop on Semantic Approximation, Granularity, and Vagueness*, 28–38. A Workshop of the Seventh International Conference on Principles of Knowldege Representation and Reasoning, Breckenridge, CO.

Ismail, H. O., and Shapiro, S. C. 1999. Cascaded acts: Conscious sequential acting for embodied agents. Technical Report 99-10, Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY. Submitted for journal publication.

Ismail, H. O., and Shapiro, S. C. 2000a. Conscious error recovery and interrupt handling. In Arabnia, H. R., ed., *Proceedings of the International Conference on Artificial Intelligence (IC-AI'2000)*, 633–639. Las Vegas, NV: CSREA Press.

Ismail, H. O., and Shapiro, S. C. 2000b. Two problems with reasoning and acting in time. In Cohn, A. G.; Giunchiglia, F.; and Selman, B., eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Seventh International Conference (KR 2000)*, 355–365. San Francisco: Morgan Kaufmann.

Ismail, H. O., and Shapiro, S. C. 2001. The cognitive clock: A formal investigation of the epistemology of time. Technical Report 2001-08, Department of Computer Science and Engineering, University at Buffalo, Buffalo, NY. Submitted for journal publication.

Kumar, D., and Shapiro, S. C. 1994a. Acting in service of inference (and *vice versa*). In Dankel II, D. D., ed., *Proceedings of The Seventh Florida AI Research Symposium (FLAIRS 94).* The Florida AI Research Society. 207–211.

Kumar, D., and Shapiro, S. C. 1994b. The OK BDI architecture. *International Journal on Artificial Intelligence Tools* 3(3):349–366.

Kumar, D. 1994. *From Beliefs and goals to Intentions and Actions: An Amalgamated Model of Inference and Acting.* Ph.d. dissertation, technical report 94-04, State University of New York at Buffalo, Buffalo, NY.

Kumar, D. 1996. The SNePS BDI architecture. *Decision Support Systems* 16(1):3–19.

Lehmann, F., ed. 1992. *Semantic Networks in Artificial Intelligence.* Oxford: Pergamon Press.

Maida, A. S., and Shapiro, S. C. 1982. Intensional concepts in propositional semantic networks. *Cognitive Science* 6(4):291–330. Reprinted in (Brachman & Levesque 1985, pp. 170–189).

Neal, J. G., and Shapiro, S. C. 1985. Parsing as a form of inference in a multiprocessing environment. In *Proceedings of the Conference on Intelligent Systems and Machines*, 19–24. Rochester, Michigan: Oakland University.

Neal, J. G., and Shapiro, S. C. 1987a. Knowledge representation for reasoning about language. In Boudreaux, J. C.; Hamill, B. W.; and Jernigan, R., eds., *The Role of Language in Problem Solving 2.* Elsevier Science Publishers. 27–46.

Neal, J. G., and Shapiro, S. C. 1987b. Knowledge-based parsing. In Bolc, L., ed., *Natural Language Parsing Systems.* Berlin: Springer-Verlag. 49–92.

Orilia, F., and Rapaport, W. J., eds. 1998. *Thought, Language, and Ontology: Essays in Memory of Hector-Neri Castañeda*. Dordrecht: Kluwer Academic Publishers.

Rapaport, W. J.; Shapiro, S. C.; and Wiebe, J. M. 1997. Quasi-indexicals and knowledge reports. *Cognitive Science* 21(1):63–107. Reprinted in (Orilia & Rapaport 1998, pp. 235–294).

Shapiro, S. C., and Neal, J. G. 1982. A knowledge engineering approach to natural language understanding. In *Proceedings of the 20th Annual Meeting of the Association for Computational Linguistics*. Menlo Park, CA: ACL. 136–144.

Shapiro, S. C., and Rapaport, W. J. 1987. SNePS considered as a fully intensional propositional semantic network. In Cercone, N., and McCalla, G., eds., *The Knowledge Frontier*. New York: Springer-Verlag. 263–315.

Shapiro, S. C., and Rapaport, W. J. 1991. Models and minds: Knowledge representation for natural-language competence. In Cummins, R., and Pollock, J., eds., *Philosophy and AI: Essays at the Interface*. Cambridge, MA: MIT Press. 215–259.

Shapiro, S. C., and Rapaport, W. J. 1992. The SNePS family. *Computers & Mathematics with Applications* 23(2–5):243–275. Reprinted in (Lehmann 1992, pp. 243–275).

Shapiro, S. C., and Rapaport, W. J. 1995. An introduction to a computational reader of narratives. In Duchan, J. F.; Bruder, G. A.; and Hewitt, L. E., eds., *Deixis in Narrative: A Cognitive Science Perspective*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc. 79–105.

Shapiro, S. C., and The SNePS Implementation Group. 1999. *SNePS 2.5 User's Manual*. Department of Computer Science and Engineering, State University of New York at Buffalo, Buffalo, NY.

Shapiro, S. C.; Ismail, H. O.; and Santore, J. F. 2000. Our dinner with cassie. In *Working Notes for the AAAI 2000 Spring Symposium on Natural Dialogues with Practical Robotic Devices*, 57–61. Menlo Park, CA: AAAI.

Shapiro, S. C. 1982. Generalized augmented transition network grammars for generation from semantic networks. *The American Journal of Computational Linguistics* 8(1):12–25.

Shapiro, S. C. 1989. The CASSIE projects: An approach to natural language competence. In Martins, J. P., and Morgado, E. M., eds., *EPIA 89: 4th Portugese Conference on Artificial Intelligence Proceedings, Lecture Notes in Artificial Intelligence 390*. Berlin: Springer-Verlag. 362–380.

Shapiro, S. C. 1993. Belief spaces as sets of propositions. *Journal of Experimental and Theoretical Artificial Intelligence (JETAI)* 5(2&3):225–235.

Shapiro, S. C. 1998. Embodied Cassie. In *Cognitive Robotics: Papers from the 1998 AAAI Fall Symposium, Technical Report FS-98-02*. Menlo Park, California: AAAI Press. 136–143.

Shapiro, S. C. 2000. SNePS: A logic for natural language understanding and commonsense reasoning. In Iwańska, Ł., and Shapiro, S. C., eds., *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language*. Menlo Park, CA: AAAI Press/The MIT Press. 175–195.