

Metacognition in SNePS

Stuart C. Shapiro, William J. Rapaport,
Michael Kandefer, Frances L. Johnson, and Albert Goldfain

October 21, 2006

Abstract

The SNePS knowledge representation, reasoning, and acting system has several features that facilitate metacognition in SNePS-based agents. The most prominent is the fact that propositions are represented in SNePS as terms rather than as sentences, so that propositions can occur as arguments of propositions and other expressions without leaving first-order logic. The SNePS acting subsystem is integrated with the SNePS reasoning subsystem in such a way that: there are acts that affect what an agent believes; there are acts that specify knowledge-contingent acts and lack-of-knowledge acts; there are policies that serve as “daemons”, triggering acts when certain propositions are believed or wondered about. The GLAIR agent architecture supports metacognition by specifying a location for the source of self-awareness, and of a sense of situatedness in the world. Several SNePS-based agents have taken advantage of these facilities to engage in self-awareness and metacognition.

Keywords: metacognition, knowledge representation, SNePS, agent architectures, self-awareness, embodied agents, cognitive robotics, commonsense reasoning.

Introduction

The terms metacognition, metarepresentation, and metareasoning have been defined in various ways, varying from general to specific: “Broadly defined *metacognition* is any knowledge or cognitive process that refers to, monitors, or controls any aspect of cognition” (Moses and Baird, 1999) (*italics in the original*). “Cognitive systems are characterized by their ability to construct and process representations of objects and states of affairs. Mental representations and public representations such as linguistic utterances are themselves objects in the world, and therefore potential objects of second-order representations, or ‘metarepresentations’” (Sperber, 1999). “Metareasoning is reasoning about reasoning—in its broadest sense, any computational process concerned with the operation of some other computational process within the same entity” (Russell, 1999). “A metacognitive reasoner is a system that *reasons* specifically about itself (its knowledge, beliefs, and its reasoning process), not one that simply *uses* such knowledge” (Cox, 2005) (*italics in original*).

The SNePS¹ knowledge representation, reasoning, and acting system, and its predecessor systems were designed from their beginnings to support metarepresentations (Shapiro, 1971; Shapiro, 1979) (*cf.* (Shapiro and Rapaport, 1992)). More recently, SNePS-based cognitive agents, usually called Cassie, have taken advantage of these facilities to have models of themselves, and to have explicit acting plans that control their own beliefs (Shapiro and Rapaport, 1987; Shapiro, 1989; Shapiro, 1998; Santore and Shapiro, 2003; Shapiro and Kandefer, 2005). This is further facilitated by the GLAIR (Grounded Layered Architecture with Integrated Reasoning) agent architecture, that specifies how an agent acquires beliefs about its own actions and percepts, and how it has a sense of its own situatedness in the world.

In this article, we will report on various aspects of SNePS, GLAIR, and Cassie that facilitate metacognition, and give examples of metacognition in SNePS-based systems.

The GLAIR Architecture

The Grounded Layered Architecture with Integrated Reasoning (GLAIR) (Hexmoor et al., 1993; Hexmoor and Shapiro, 1997; Shapiro and Ismail, 2003), illustrated in figure 1, is a cognitive agent architecture with five layers.

¹SNePS is available for free downloading from <http://www.cse.buffalo.edu/sneps/Downloads/>.

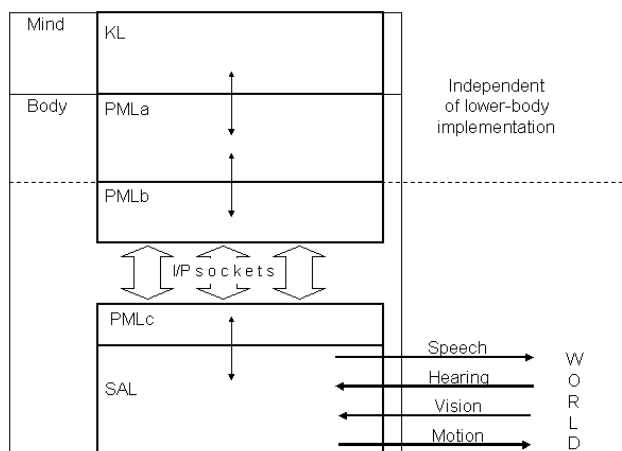


Figure 1: The GLAIR architecture: The KL layer is the agent’s mind, the PML and SAL layers are its brain/body. The KL and PMLa layers are independent of whether the agent’s body is implemented in software, virtual reality, or hardware. If the PMLc and SAL run on a different computer from the KL, PMLa, and PMLb, then the PMLb and PMLc communicate over I/P sockets, one for each modality. The SAL controls the sensors and effectors.

The Knowledge Layer (KL) is the layer at which “conscious” reasoning takes place. The KL is implemented in SNePS and its subsystem SNeRE (the SNePS Rational Engine) (Shapiro and Rapaport, 1992; Kumar, 1996; Kumar and Shapiro, 1994a; Kumar and Shapiro, 1994b; Shapiro, 2000b; Shapiro and The SNePS Implementation Group, 2004). SNePS, in turn, is implemented in Common Lisp. In the KL, terms of the SNePS logical language represent the mental entities conceived of and reasoned about by the agent.

The Perceptuo-Motor Layer, Sublayer a (PMLa) contains the Common Lisp implementation of the actions that are primitive at the KL, *i.e.*, the routine behaviors that can be carried out without thinking about each step. PMLa is implemented in Common Lisp in a way that is independent of the implementation of the lower layers.

The Perceptuo-Motor Layer, Sublayer b (PMLb) implements the functions of PMLa taking into account the particular implementation of the agent’s body. PMLb is implemented in Common Lisp using, when necessary, its facilities for interfacing with programs written in other languages.

The Perceptuo-Motor Layer, Sublayer c (PMLc) contains the implementations of the PMLb functions for the particular hardware or software agent body being used. PMLc has been implemented variously in C, Java, and the Ygdrasil (Pape et al., 2003) virtual reality authoring system.

The Sensori-Actuator Layer (SAL) contains the sensor and effector controllers of the agent body.

SNePS

Basic Representation: Propositions, Acts, Actions, Policies

We view the contents of the SNePS knowledge base (KB), appropriately for its use as the agent’s KL, to be the contents of the agent’s mind. SNePS terms denote the mental entities of the agent, entities the agent has conceived of (Maida and Shapiro, 1982; Shapiro and Rapaport, 1987). Mental entities must include everything an agent can possibly conceive of, including particular individuals, classes, relations, acts, and, most significantly for this paper, propositions (Shapiro, 1993). The root of the SNePS ontology is thus entity. The next level, developed by considering what a SNePS-based agent can do with each entity, is currently: proposition, act, policy, and thing. A *proposition* is an entity such that the agent can believe it or its negation. An *act* is an entity that the agent can perform. A *policy* is an entity, connecting propositions and acts, which the agent can adopt. A *thing* an entity that is neither a proposition, act, nor policy² Note that propositions, acts, policies, and things are exhaustive and mutually exclusive subclasses of entity.

²In previous papers, “things” were called “individuals”, but that was probably confusing, since this kind of entity includes classes and relations.

A subset of the propositions represented in the KB are “asserted” and thus believed by the agent. The implemented SNePS rules of inference specify how the assertion status can be inherited by propositions from other propositions (Shapiro, 1993; Chalupsky and Shapiro, 1994). Open terms (terms containing unbound variables) do not exist in the latest version of SNePS, SNePS 3 (Shapiro, 2000a), which is based on the logic \mathcal{L}_A (Shapiro, 2004), and is still being implemented. They also did not exist in the variant of SNePS discussed in (Ali and Shapiro, 1993).

An act may be performed by an agent, and is composed of an **action** and zero or more arguments. For example, for the Fevahr³ version of Cassie (Shapiro, 1998), the term `find(Bill)` denotes the act of finding Bill (by looking around in a room for him), composed of the action `find` and the object `Bill`.⁴ That is, an action is represented by an act-valued function symbol.

Three policy-forming function symbols are built into SNePS, each of which take as arguments a proposition p and an act a : `ifdo(p, a)` is the policy that if the agent wants to know whether to believe p , it should perform a ; `whendo(p, a)` is the policy that when the agent believes p , it should perform a ; `wheneverdo(p, a)` is the policy that whenever the agent believes p , it should perform a .

A blocks-world example of `ifdo` is “If the agent wants to know whether block A is red, it should look at it”: `ifdo(ColorOf(A, red), lookAt(A))` (Kumar and Shapiro, 1994a).⁵

In the case of both `whendo` and `wheneverdo`, if the policy has been adopted, the agent performs a when forward inference causes p to be believed. Also, a is performed if p is already believed when the policy is adopted with forward inference. The difference is that a `whendo` policy is unadopted after firing once, but a `wheneverdo` remains adopted until explicitly unadopted.

Things are all entities other than propositions, acts, and policies. Things mentioned so far in this section are the person denoted by `Bill`, the block denoted by `A`, the color denoted by `red`, the actions denoted by `find` and `lookAt`, and the relation denoted by `ColorOf`. Fevahr Cassie’s things also include the category of robots, and the sentence “Neither a borrower nor a lender be”, denoted by a functional term.

Bidirectional Inference

SNePS performs inference in a bidirectional fashion (Shapiro et al., 1982). A rule of the form $\text{all}(x)(P(x) \Rightarrow Q(x))$ can be used for backward chaining if an instance of $Q(x)$ is a goal or subgoal, or for forward chaining if an instance of $P(x)$ is asserted into the KB in an appropriate way. (When a belief is asserted into the KB, it may be done in a way that triggers forward chaining or not.) If both $\text{all}(x)(P(x) \Rightarrow Q(x))$ and $\text{all}(x)(P(x) \Rightarrow R(x))$ are asserted, and $P(a)$ is not asserted, and $Q(a)$ is issued as a goal or subgoal, and then $P(a)$ is asserted with forward chaining, bidirectional inference focuses the forward chaining so that $Q(a)$ is inferred, but $R(a)$ is not.

Metaknowledge

Metaknowledge is knowledge about knowledge. More accurately, we are concerned with metabeliefs—beliefs about beliefs. Since propositions are represented by terms in the SNePS logic (Shapiro, 1993; Shapiro, 2000b), metabeliefs—which are propositions about propositions—are easily represented without leaving first-order logic. One example of a proposition represented by a term is from the Wumpus-World-agent version of Cassie (Shapiro and Kandefter, 2005). There, `Have` is a function symbol, and `Have(gold)` is a proposition-valued functional term denoting the proposition that “I have the gold.”.

A SNePS proposition of the form $P \Rightarrow Q$ is not a sentence denoting True when P is False or Q is True; rather it is a functional term denoting the proposition that, from the agent’s point of view, if I believe P , then I am justified in believing Q (Shapiro, 1993; Chalupsky and Shapiro, 1994). Notice that it is not “if I believe P , then I believe Q ”, because SNePS agents are not logically omniscient. If an agent believes P and $P \Rightarrow Q$, then it will believe Q only when the “rule” $P \Rightarrow Q$ “fires” (to borrow terminology from production systems). Rather than being a function from truth values to truth values, \Rightarrow is a function from propositions to propositions, a proposition-valued function. Moreover, $\sim(P \Rightarrow Q)$ simply denotes the proposition that belief in P is not justification for believing Q . So from $\sim(P \Rightarrow Q)$ it does not follow that P , much to the consternation of some traditionalists. Similarly, $P \Rightarrow (Q \Rightarrow R)$ denotes the

³“Fevahr” is an acronym standing for “Foveal Extra-Vehicular Activity Helper-Retriever”.

⁴Actually, since the Fevahr Cassie uses a natural language interface, the act of finding Bill is represented by the term `act(lex(find), b6)`, where: `find` is a term aligned with the English verb *find*; `lex(find)` is the action expressed in English as “*find*”; and `b6` is a term denoting Bill. However, we will ignore these complications in this paper.

⁵`ifdo` was called `DoWhen` in (Kumar and Shapiro, 1994a).

proposition that if I believe P , then I am justified in believing $Q \Rightarrow R$. Although $P \Rightarrow (Q \Rightarrow R)$ and $\{P, Q\} \& \Rightarrow R$ are logically equivalent⁶, in the sense that they both say that belief in R is justified when both P and Q are believed, they are treated differently by the SNePS inference system. If SNePS backchains into the latter rule, P and Q will be established as parallel subgoals, but if SNePS backchains into the former rule, just P will be a subgoal. Only if P is then inferred will $Q \Rightarrow R$ be believed and Q established as another subgoal. In either case, if SNePS forwardchains into either P or Q , it will then backchain on the other, and R will be asserted if that backchaining is successful. From this point of view, $P \Rightarrow Q$, $\sim(P \Rightarrow Q)$, $P \Rightarrow (Q \Rightarrow R)$, and $\{P, Q\} \& \Rightarrow R$ are all examples of metabeliefs just as are the more obvious cases such as $\text{Believes}(a, p)$, where a denotes some agent, p denotes some proposition, and $\text{Believes}(a, p)$ is a proposition-valued functional term denoting the proposition that a believes p (cf. (Wiebe and Rapaport, 1986; Shapiro and Rapaport, 1991; Rapaport et al., 1997)). The policies $\text{ifdo}(p, a)$, $\text{whendo}(p, a)$, and $\text{wheneverdo}(p, a)$, although not beliefs about beliefs (because policies are not entities to be believed), are metapropositions in the sense that they are SNePS terms whose arguments are propositions and they are used to monitor beliefs.

Categories of Acts

External, Mental, and Control Acts

SNePS actions and, by extension, acts, may be subclassified as either external, mental, or control. External acts either sense or affect the real, virtual, or simulated outside world. An example mentioned above from the Fevahr version of Cassie is $\text{find}(\text{Bill})$. No external acts are predefined in SNePS; they must be supplied by each agent designer.

Mental acts affect the agent's beliefs and adopted policies. There are four: $\text{believe}(p)$ is the act of asserting the proposition p and doing forward inference on it; $\text{disbelieve}(p)$ is the act of unasserting the proposition p , so that it is not believed, but its negation is not necessarily believed; $\text{adopt}(p)$ is the act of adopting the policy p ; $\text{unadopt}(p)$ is the act of unadopting the policy p .

Before believe changes the belief status of a proposition p , it performs a limited form of belief revision (Alchourrón et al., 1985). If $\text{andor}(0, 0)\{\dots, p, \dots\}$ is believed,⁷ it is disbelieved. If $\text{andor}(i, 1)\{p, q, \dots\}$ is believed, for $i = 0$ or $i = 1$, and q is believed, q is disbelieved. Mental acts are metacognitive in that they are agent acts that explicitly affect what the agent believes.

Control acts are the control structures of the SNePS acting system. The SNePS control actions are achieve , do-all , do-one , prdo-one , snif , sniterate , ssequence , withall , and withsome .

If the proposition p is asserted, performing the act $\text{achieve}(p)$ does nothing. Otherwise, performing the act $\text{achieve}(p)$ consists of inferring plans for bringing about the proposition p (by inferring instances of the proposition $\text{GoalPlan}(p, x)$), and then performing do-one on them.

Performing the act $\text{do-all}(\{a_1, \dots, a_n\})$ consists of performing all the acts a_1, \dots, a_n in a nondeterministic order.

Performing the act $[\text{do-one}(\{a_1, \dots, a_n\})]$ consists of nondeterministically choosing one of the acts a_1, \dots, a_n , and performing it.

Performing the act $\text{prdo-one}(\{\text{pract}(x_1, a_1), \dots, \text{pract}(x_n, a_n)\})$ consists of performing one of the acts a_j , with probability $x_j / \sum_i x_i$.

Performing the act $\text{snif}(\{\text{if}(p_1, a_1), \dots, \text{if}(p_n, a_n)[, \text{else}(d)]\})$ consists of using backward inference to determine which of the p_i hold, and, if any do, nondeterministically choosing one of them, say p_j , and performing a_j . If none of the p_i can be inferred, and if $\text{else}(d)$ is included, d is performed. Otherwise, nothing is done.

Performing the act $\text{sniterate}(\{\text{if}(p_1, a_1), \dots, \text{if}(p_n, a_n)[, \text{else}(d)]\})$ consists of using backward inference to determine which of the p_i hold, and, if any do, nondeterministically choosing one of them, say p_j , performing a_j , and then performing the entire sniterate again. If none of the p_i can be inferred, and if $\text{else}(d)$ is included, d is performed. Otherwise, nothing is done.

Performing the act $\text{ssequence}(a_1, a_2)$ consists of performing the act a_1 and then the act a_2 .

⁶The term $\{P_1, \dots, P_n\} \& \Rightarrow \{Q_1, \dots, Q_m\}$ denotes the proposition that if all the P_i are believed then belief is justified in any or all of the Q_j (Shapiro, 1979). If either n or m is 1, the set brackets can be omitted.

⁷ andor (Shapiro, 1979) is a parameterized connective that takes a set of argument-propositions, and generalizes *and*, *inclusive or*, *exclusive or*, *nand*, *nor*, and *exactly n of*. A formula of the form $\text{andor}(i, j)\{P_1, \dots, P_n\}$ denotes the proposition that at least i and at most j of the P_k s are true.

Performing the act `withall(x, p(x), a(x), [d])` consists of performing backward inference to find entities e such that $p(e)$ is believed, and, if such entities are found, performing a on them all in a nondeterministic order. If no such e is found, and the optional act d is present, d is performed.

Performing the act `withsome(x, p(x), a(x), [d])` is like performing `withall`, but if any entities e such that $p(e)$ are found, a is performed on one of them, nondeterministically chosen.

Primitive, Defined, and Composite Acts

SNePS actions and acts may also be classified as either primitive, defined, or composite. Primitive acts constitute the basic repertoire of a SNePS agent. They are either provided by the SNePS system, itself, or are implemented at the PML. An example predefined action is `believe`; an example primitive action defined at the PML is the Fevahr Cassie's `find`. Because primitive actions are implemented below the KL, SNePS agents have no cognitive insight into how they perform them.

A composite act is one structured by one of the control acts. For example, the Wumpus-World Cassie, whose only primitive turning acts are `go(right)` and `go(left)`, can turn around by performing the composite act, `snsequence(go(right), go(right))`.

A defined act is one that, unlike composite acts, is syntactically atomic, and unlike primitive acts, is not implemented at the PML. If a SNePS agent is to perform a defined act a , it deduces plans p for which it believes the proposition `ActPlan(a,p)`, and performs a `do-one` of them. Such a plan is an act which, itself, can be either primitive, composite, or defined. For example, the Wumpus-World Cassie has a defined act `turn(around)` defined by `ActPlan(turn(around), snsequence(go(right), go(right)))`.

Examples of Metacognition

Self-Awareness

If “a metacognitive reasoner is a system that reasons specifically about itself” (Cox, 2005), it needs a model of itself. There are two aspects to our approach to self-modeling: representation techniques in SNePS at the KL; and PML-KL interaction. At the KL, the agent, itself, is represented by a SNePS term that is in no way different from the SNePS terms representing other agents. Beliefs about itself are represented as propositions containing the self-term as an argument. Beliefs about what the agent is doing and has done are represented using an interval model of time (Allen, 1983), with the times principally connected by `before` and `during` relations (which each represent a disjunction of several of Allen's temporal relations). These beliefs form the agent's episodic memory.

The PML is involved in self-awareness in order to model an agent's sense of embodiedness and of being situated in the world. There are two aspects to this: a source of beliefs; and a locus of indexical and similar information. Beliefs about what the agent is doing and beliefs derived from the agent's sensory organs are put into the KL directly from the implementation of primitive sensory and efferent actions in the PMLa. For example, as mentioned above, `find` is a primitive action for the Fevahr Cassie (Shapiro and Ismail, 2003). The implementation of `find` in the PMLa inserts into the KL the belief that Cassie has found whatever she did. So after being told to `find(Bill)`, Cassie believes, and then says (via her natural language (NL) generation component) “*I found Bill.*” Having the agent's beliefs of what it senses and does inserted directly by the PML models these beliefs as “first-person privileged knowledge.”

Another aspect of being situated in the world is continuous awareness of one's self and of one's general orientation in the world. We accomplish this via a set of PML deictic registers and modality registers, each of which is a variable whose value is one or a set of KL terms.

The deictic registers derive from the theory of the Deictic Center (Duchan et al., 1995), and include: `I`, the register that holds the KL term denoting the agent itself; `YOU`, the register that holds the KL term denoting the individual the agent is talking with; and `NOW`, the register that holds the KL term denoting the current time. It is by using these registers that the NL generator generates “*I found Bill*” in first-person, in past tense, and using the third-person referring expression, “*Bill*”. If she had been talking to Bill at the time, she would have said, “*I found you*” or “*I found you, Bill.*”

Each modality register contains the KL term or terms representing the act(s) the agent is currently performing, or the percept(s) the agent is currently having in each effector or affector modality. The mechanisms for advancing `NOW` use these registers, as they change their contents, to place past acts and percepts in the past (Ismail, 2001). For

example, Fevahr Cassie starts out talking with Stu. If Stu then tells her to “*Talk to Bill.*” and Bill asks her “*Who have you talked to?*”, she says, “*I talked to Stu and I am talking to you.*”

Lack of Knowledge Acting

The optional default act in the SNePS control acts `withall` and `withsome` provide for acting on the basis of lack of knowledge (*cf.* (Moore, 1988)). We are using this facility for a SNePS solution to what we call McCarthy’s second telephone number problem (McCarthy, 1977). In this problem, the agent is asked to call someone. If the agent knows the person’s phone number it should make the call. However, if it doesn’t, it must obtain the number from some directory or other agent, and then dial it. Obtaining the number is to result in the agent’s knowing the person’s phone number.

Our agent’s plan for calling someone is,

```
all(p)(Person(p)
      => ActPlan(call(p),
                withsome(n,
                        Has(p, PhoneNumber, n),
                        dial(n),
                        snsequence(lookup(p),
                                  call(p))))))
```

where `Has(p, PhoneNumber, n)` means that “*p’s phone number is n.*”

In the following demonstration of a SNePS agent solving this problem, material output by the agent is shown in **bold monospace** font. The “`:`” is the system prompt. User input is in `light monospace` font.

```
: Person(?x)?
Person(Stu)
Person(Mike)

: Has(?x, PhoneNumber, ?y)?
Has(Mike, PhoneNumber, N(5, N(5, N(5, N(5, N(6, N(1, 2)))))))

: perform call(Mike)
I am pressing 5.
I am pressing 5.
I am pressing 5.
I am pressing 5.
I am pressing 6.
I am pressing 1.
I am pressing 2.

: perform call(Stu)
I could not find Stu’s phone number in any external information source available to me.
Do you know Stu’s number? yes
What is Stu’s number (e.g. 555-5555)? 555-7890
I am pressing 5.
I am pressing 5.
I am pressing 5.
I am pressing 7.
I am pressing 8.
I am pressing 9.
I am pressing 0.
```

```

: Has(?x, PhoneNumber, ?y)?
Has(Stu, PhoneNumber, N(5, N(5, N(5, N(7, N(8, N(9, 0)))))))
Has(Mike, PhoneNumber, N(5, N(5, N(5, N(5, N(6, N(1, 2)))))))

```

The agent engaged in an information-acquiring act because it lacked the information needed to perform another act. After acquiring the information, it performed the latter act, and then knew the information it acquired.

Metacognition in Consistency Maintenance and Optimization

Any computational system that stores and reasons with information must have some means for adding new information, changing existing information, and removing information. These belief change operations are especially needed to resolve contradictions (also called “consistency maintenance” or “truth maintenance”).

SNePS uses a monotonic logic, so contradiction resolution requires that at least one base (input) belief that underlies a contradiction must be retracted to eliminate that contradiction.⁸ The retraction process unasserts the belief, effectively removing it from the current base (also called the current context); both the contradiction and any derived beliefs that were dependent on that base belief are no longer derivable. The retracted belief is retained in the KB as an unasserted belief about which the system can still reason.⁹

Belief change is managed in SNePS by SNeBR, the SNePS Belief Revision system (Martins and Shapiro, 1988; Shapiro, 2000b). As implemented in SNePS 2.6.1, SNeBR detects contradictions, presents the user with possible culprits, provides the interface for base belief retraction, and removes from the active KB derived beliefs whose assertion status depended on any removed base beliefs. Automated retraction of base beliefs (Johnson and Shapiro, 1999; Shapiro and Johnson, 2000) and the belief base re-optimizing operation of reconsideration (Johnson and Shapiro, 2005a; Johnson, 2006) are implemented in other, developmental versions of SNePS.

Key guidelines for consistency maintenance are:¹⁰ when some beliefs are considered more important (or credible) than others, and consistency maintenance forces the removal of some beliefs, then the least important beliefs should be selected for removal; and any beliefs removed during consistency maintenance must be responsible for the inconsistency.

Because each belief is a term in the SNePS logic, metabeliefs can be used to indicate various levels of credibility. For example, we can order the beliefs `Truck(Obj1)` (*Object 1 is a truck*) and `Helicopter(Obj1)` (*Object 1 is helicopter*) to indicate that the first is more credible by asserting `moreCredible(Truck(Obj1), Helicopter(Obj1))`. This belief credibility ordering enables the automation of contradiction resolution. Reasoning with the rule that `all(x)(and(0,1){Truck(x), Helicopter(x)})` (*An object cannot be both a helicopter and a truck*) the system detects a contradiction and, following the guidelines listed above, removes the least credible belief¹¹ and reports,

```
I will remove the following node: Helicopter(Obj1).
```

Ordering beliefs by their source or category (Johnson and Shapiro, 1999; Shapiro and Johnson, 2000) is more reasonable than ordering all beliefs individually. This results in a pre-order over the beliefs; some sources may be equivalent (or incomparable) in strength, and beliefs from the same source (or equivalent sources) would have equivalent strength (unless there is an additional sub-ordering of those beliefs).

Again, because each belief is a term in the language, assigning and ordering sources is easily represented in SNePS:

```
Source(UAV, Truck(Obj1))
(An unmanned aerial vehicle (UAV) is the source of the information that Object 1 is a truck.)
```

```
Source(Satellite, Helicopter(Obj1))
```

⁸In this section, “propositions” are referred to as “beliefs,” whether they are asserted or not.

⁹Note the difference between the *current base*, which is the context comprised of the input propositions that are *currently* believed, and the entire *knowledge base* (KB), which contains *all* base propositions (both asserted and unasserted) as well as derived beliefs. The *active* KB contains only asserted base beliefs and the beliefs that have been derived from them. It is important to retain non-currently believed propositions when reasoning in multiple contexts (Martins and Shapiro, 1983; Shapiro, 2000b), and when allowing for reconsideration, as will be discussed below.

¹⁰These guidelines are paraphrased from (Gärdenfors and Rott, 1995) and (Hansson, 1997).

¹¹In this example, rules are preferred over observations. In the implementation of automated contradiction resolution (Johnson and Shapiro, 1999; Shapiro and Johnson, 2000), which is not included in the SNePS 2.6.1 release, if the system cannot identify a single belief for removal, it advises the user, who must make the final selection.

(Satellite imagery is the source of the information that Object 1 is a helicopter.)

```
all(s1,s2,b1,b2)({Source(s1,b1),Source(s2,b2)})
  &=> {(sourceMoreCredible(s1,s2)
    => moreCredible(b1,b2)),
    (sourceEquallyCredible({s1,s2})
    => equallyCredible({b1,b2}))})
```

(Information from a more credible source is more credible than information from a less credible source, but information from two equally credible sources are equally credible.)

When we then input `sourceMoreCredible(UAV,Satellite)` (*The UAV is more credible than satellite imagery*) the system infers `moreCredible(Truck(Obj1),Helicopter(Obj1))`, which is used during automated contradiction resolution, as before.

Ordering beliefs is also helpful if a system needs to re-optimize the base. This can happen when new input and/or inferences result in changes being made to the base or its ordering. For example, if the new information forces contradiction resolution to remove some belief p from the base, then any belief q that is weaker than p and was removed because it conflicted with p , should be considered for *possible* return to the base. Alternatively, if the new information re-orders the sources and/or beliefs (or further refines the pre-ordering), the base may need to be re-optimized with respect to this new ordering. This optimization process is called reconsideration (Johnson and Shapiro, 2005a; Johnson, 2006) and is implemented using the anytime algorithm of dependency-directed reconsideration (DDR) (Johnson and Shapiro, 2005b; Johnson, 2006).

DDR processes unasserted base beliefs that *might* be able to return to the current base due to a stronger conflicting belief being recently retracted or a re-ordering of the base beliefs. A belief can be re-asserted if its reassertion either (1) does not raise an inconsistency or (2) each inconsistency its return raises can be resolved by retracting one or more strictly weaker beliefs (called belief swapping). Beliefs are processed in a non-increasing order of credibility and are only examined if they are directly linked to some change in the belief base.

This scales well in the case where new information forces a retraction. Each inconsistency resolution typically involves a small subset of base beliefs, and reconsideration would involve a similarly small subset.¹² If a belief is returned with swapping, reconsideration continues but is progressing through ever decreasing credibility weights; if there is no swap—whether or not the belief is reasserted—reconsideration along that branch of the KB is pruned. The computational complexity of the DDR algorithm is minimal when a belief cannot return (no change to the base) and has a higher cost when the base is improved.

Re-ordering base beliefs (both asserted and unasserted) might occur when a source is discovered to be unreliable—forcing information from that source to be considered less credible than once thought. Reconsideration re-optimizes the base to reflect the new ordering by processing retracted base beliefs that are now more credible than conflicting base beliefs that are currently asserted. Unasserted base beliefs that have not changed their order relative to stronger conflicting beliefs need not be processed.

Any implemented system is restricted by the real-world limitations of the size of its memory and the time it takes to reason. There may be times when DDR must be interrupted so that further inferences can be made or actions taken. Because the logic used in SNePS is a paraconsistent logic,¹³ the SNePS system can even reason in an inconsistent space—though, typically, the inconsistency needs to be unrelated to the propositions being used in the reasoning. When time allows, DDR can be re-called to re-optimize the KB and provide a more preferred belief base from which to act and reason.

Whether for contradiction resolution or base optimization, the system uses metabeliefs to determine *what* it should believe—insuring that it maintains and reasons from the most credible knowledge base possible given the current constraints and needs of the user.

¹²This assumes inconsistency maintenance uses dependency-directed, truth maintenance system (TMS) techniques (de Kleer, 1986; Forbus and de Kleer, 1993) as opposed to some ‘entire base’ technique using SAT testing, where scalability is an issue.

¹³It is *not* the case that everything is derivable from an inconsistency.

Metacognition and Contextual Vocabulary Acquisition

Reading is a cognitive activity that begins by requiring a great deal of conscious processing (in order to assign sounds to letters and meanings to words and sentences). Readers must think about which sounds go with which letters, and then whether the resulting combined sound is a word that they know. For good readers, this process evolves into an automatic process requiring little or no conscious activity (e.g., when looking at a text in their native language, good readers cannot *not* read the passage, since reading has become so automatic). But, for excellent readers, it evolves once again into a conscious activity during which the reader monitors his or her reading (e.g., in order to draw inferences about the passage being read). (See, e.g., (McNamara and Kintsch, 1996; van Oostendorp and Goldman, 1999; Ruddell and Unrau, 2004).) In short, for excellent readers, reading involves metacognition.

But even for readers who are not that excellent, reading contains metacognitive components. Moreover, we believe that practice with at least one of these components (vocabulary acquisition) can improve reading and comprehension skills. In our Contextual Vocabulary Acquisition project,¹⁴ we are developing a computational theory of how readers can “figure out” (i.e., compute) a meaning for an unknown word from the textual context augmented by the reader’s prior knowledge (Ehrlich, 1995; Ehrlich and Rapaport, 1997; Rapaport and Ehrlich, 2000; Rapaport, 2003; Rapaport, 2005; Kibby and Rapaport, 2007). This is a metacognitive task requiring readers to monitor their prior knowledge while reading, i.e., to think about their thoughts in order to improve their vocabulary. We are adapting our theory (implemented in SNePS) to a school curriculum that, we hope, will help improve students’ reading and comprehension abilities by showing them a precise series of steps they can follow (i.e., by providing them with an algorithm) for figuring out meaning from textual context plus prior knowledge.

In this section, we present one example of how our algorithm works. In the *Morte Darthur* (Malory, 1470)—one of the earliest versions of the King Arthur legends—there occurs the following passage (here minimally paraphrased into modern English):

Therewith two of them [i.e., two evil knights] *dressed* their spears and Sir Ulfyus and Sir Brastias [Arthur’s knights] *dressed* their spears; they ran together with great force. (Malory, 1470, p. 15)

Even a very excellent reader who can manage to penetrate the archaic grammar (itself probably a metacognitive task) is likely to be puzzled by this use of ‘dress’, given that for modern readers, ‘dress’ means—without question—“to put clothes on”. Did these knights really put clothes on their spears?

Thus, puzzled readers compare what they are reading *with what they antecedently believed* about dressing and spears in order to discover an inconsistency. This is therefore a metacognitive activity; metacognition tells the reader that this makes no sense. So perhaps we should take the passage metaphorically; perhaps the “clothes” they put on their spears were some sort of sheaths. But that, too, makes no sense: One would expect them to *remove* such sheaths before battle, not put them on. Such (conscious) thinking about a cognitive process is also a metacognitive task.

A bit later, Merlin advises King Arthur on the disposition of troops:

When it is daylight, dress your battle right before the Northern host and the strong passage-way, so that they may see all your host. (Malory, 1470, p. 19)

Here, the context strongly suggests that ‘battle’ means “troops”. Does this mean that the great King Arthur puts clothes on his troops before battle? Again, even a metaphorical reading according to which he helps his troops don their *armor* makes little sense; surely, King Arthur has better things to do with his time.

We have conflicts among what the text says, the reader’s knowledge of the meaning of ‘dress’, and the reader’s prior beliefs that spears don’t wear clothing and that kings don’t put clothing on their troops.

Our implemented algorithm handles this situation as follows (as adapted from (Rapaport and Ehrlich, 2000)): First, the system’s prior-knowledge base for this example includes the following two beliefs: $\text{all}(x, y) (\text{dresses}(x, y) \Rightarrow \{\text{wears}(y, \text{clothesOf}(y)), \text{Isa}(\text{clothesOf}(y), \text{clothing})\})$ (If x dresses y , then y wears y ’s clothes); $\text{all}(x, y) (\{\text{Instance}(x, \text{spear}), \text{Isa}(y, \text{clothing})\} \&\Rightarrow \sim\text{wears}(x, y))$ (Spears don’t wear clothing.)

Next, the system reads a sequence of passages containing ‘dress’, adapted from (Malory, 1470), interspersed with questions and requests for definitions. First the system reads,

Passage D1 King Arthur dressed himself.

¹⁴<http://www.cse.buffalo.edu/~rapaport/CVA/>

Then, asked to define “dress”, the system produces,

Definition D1 A person can dress itself; result: it wears clothing.

Then it reads,

Passage D2 King Claudas dressed his spear

At this point, the system infers that King Claudas’s spear wears clothing. However, challenged with

Question D2 What wears clothing?

a contradiction is detected, and SNeBR is invoked, automatically replacing the prior belief that if x dresses y , then y wears y ’s clothes. (*Rather than the prior belief that spears don’t wear clothing*, because of the occurrence of a verb in the antecedent, since people tend to revise beliefs about verbs rather than beliefs about nouns (Gentner, 1981)). There follow several passages in the text in which dressing spears precedes fighting. Rather than *rejecting* the prior definition, it is *modified*. The system decides that to dress is *either* to put clothes on *or* to prepare for battle:

Definition D3 A person can dress a spear or a person; result: the person wears clothing or the person is enabled to fight.

Such disjunctive definitions are consistent with dictionary-like definitions of polysemous words.¹⁵ We plan to investigate a method to *induce* a more general definition: In the present case, further experience with such phrases as ‘salad dressing’, ‘turkey dressing’, and so on, should lead the reader to decide that ‘dress’ more generally means something like “prepare” (for the day, by putting on clothes; for battle, by preparing one’s spear; for eating, by preparing one’s salad; for cooking, by preparing one’s turkey; and so on).

Metacognition arises in this situation in a variety of ways. First, there is metacognitive activity when the reader breaks out of the predominantly subconscious process of reading when faced with an inconsistency; the reader believes that what she or he is reading is inconsistent with other things that she or he believes. This belief about a belief is metacognitive. Second, when the reader decides to re-interpret what she or he reads in order to eliminate the inconsistency, there is metacognitive activity, because this is a decision about what to believe. It is a cognitive process (deciding) about another cognitive process (what to believe). Finally, the search for a more general meaning that subsumes the apparently inconsistent ones is metacognitive, because it involves reasoning *about* not merely *with* one’s (other) beliefs.

Metacognition in a Math-Capable Computational Agent

Mathematical problem-solving often requires some degree of metacognitive reasoning. The importance of metacognition in mathematics education is well known (cf. (Cohors-Fresenborg and Kaune, 2001; Gartmann and Freiberg, 1995; Schoenfeld, 1992)) and should not be ignored when developing math-capable AI agents.

The applicability of multiple solution-methods to any given problem is a driving force for metacognition in mathematics. The differences between solution methods may involve different mathematical contexts (e.g., a student may know how to solve a problem both algebraically and geometrically) or subtle variations (e.g., a child may count from either the smaller or larger addend when count-adding two numbers). Furthermore, a student may use (or improvise) a “non-standard” solution method in certain problem contexts (e.g., checking to see whether 4152342 is even to find the remainder of $4152342 / 2$, rather than performing the long division as with the other numbers). These abilities suggest two metacognitive requirements for a math-capable AI agent: the ability to represent multiple procedures for the same task; the ability to select between procedures using contextual information.

In this section, we will describe how these abilities can be supported in SNePS agents. The focus will be on mathematical procedures at the level of counting and arithmetic.

Multiple Plans for a Single Act

Even in the early mathematical domain of arithmetic, a plurality of procedures arises for each of the arithmetic operations. There are at least three ways two natural numbers, x and y , can be added: count-addition; arithmetic addition; and calculator addition.

¹⁵Their resemblance to “weakening” in belief revision is worthy of further investigation; see, e.g., (Benferhat et al., 2004).

In *count-addition*, $x + y$ is computed by counting from x for y numbers, and seeing what number has been reached. The result of the counting procedure is the result of the addition procedure.

In *arithmetic addition*, single-digit sums for all pairs of single-digit numbers are given. The numbers x and y are treated as strings of decimal digits, and the notion of a “carry” is used when a single-digit sum exceeds 9. The sum $x + y$ is computed by performing columnwise sums for each increasing place-value in x and y , and concatenating these results in the appropriate way.

In *calculator addition*, $x + y$ is computed by performing a sequence of button presses on a calculator (e.g., entering x , pressing +, entering y , pressing =).

Each of these procedures specifies a potentially infinite set of ActPlans, one for each x and y pair.¹⁶ These are specified in SNePS as follows:

```
all(x,y)({Number(x),Number(y)}
  &=> ActPlan(Add(x,y),
    {CountAdd(x,y),
     ArithAdd(x,y),
     CalcAdd(x,y)})).
```

The propositions representing the result for each specific plan should be different, so that Cassie will remember which way of addition was used. The three ways of adding two numbers require three propositions: CountSum(x, y, z), ArithSum(x, y, z), and CalcSum(x, y, z). Whenever Cassie performs a series of arithmetic operations, she will have a trail of result propositions that indicate which plan was chosen at each step. This memory of how a plan was performed is essential for metacognitive reasoning.

In addition to a specific result propositions, there should also be a general result proposition so the sum may be used in other operations without regard to how it was computed:

```
all(x,y,z)({CountSum(x,y,z),
  ArithSum(x,y,z),
  CalcSum(x,y,z)}
  v=> Sum(x,y,z)).
```

When Cassie is asked to perform an operation, say perform CountAdd(2,3), she will believe CountSum(2,3,5) and, through forward inference, Sum(2,3,5). The question “What is the sum of 2 and 3?” can be asked in two ways: Sum(2,3,?x)?, which is plan-independent; and CountSum(2,3,?x)?, which depends on Cassie having used the CountAdd method to obtain her sum.

Without any additional machinery, Cassie will nondeterministically select an addition plan whenever she is asked to perform Add(x, y). Metacognition requires that Cassie makes a more informed decision. A plan should be chosen when it is an appropriate way of doing addition *in the current context*. To determine the contextual relevance of a plan, it will be useful to categorize the different kinds of arithmetic procedures.

Categorization of Mathematical Procedures

Clearly, there are many more procedures for addition than those given in the previous section. The three listed above are representative of three different procedural categories. Count-addition is a procedure in which an operation is performed by appealing to a more basic operation. In fact, we can build up all of the fundamental arithmetic operations in this way: addition is achieved by counting, subtraction is inverted addition, multiplication is repeated addition, and division is inverted multiplication. These are categorized as *semantic* procedures because they involve a semantic interpretation of the given procedure in terms of another (simpler) procedure. Arithmetic addition is a procedure characterized by syntactic operations in a positional number system and uses columnwise operations. This is done *without* a semantic interpretation for each of the single-digit sums (i.e., it doesn’t tell us what addition *is* in terms of another concept). Such procedures are categorized as *syntactic* procedures. Calculator addition requires non-mathematical operations (i.e., button presses) and an external (trusted) source. Such procedures are categorized as *extended* procedures.

¹⁶Because SNePS uses the unique variable binding rule (Shapiro, 1986), which stipulates that no one term can substitute for different variables, separate ActPlans must be defined for adding a number to itself (“doubling a number”).

The borders between these procedural categories are not sharp ones. Arithmetic addition, for example, usually requires a student to mark single-digit sums on a piece of paper; using an extended medium is a characteristic of an extended procedure. However, even with an imperfect categorization, it is important to assign some category to each of the agent’s procedures. A procedure’s category will be indicative of the situations in which a procedure will be most useful.

Semantic routines will be most useful for an agent in a context of explanation. This might include anything from a tutoring system for a human user to an agent taking a Turing Test (in which mathematical understanding is tested alongside natural-language understanding). Syntactic routines will be most useful when an agent must “show its work” or leave some indication of its activities. The utility of extended routines is very much dependent on the extended medium being used. In general, extended routines may be most useful in various contexts of discovery, or when a syntactic or semantic routine is simply unavailable. An agent that must obtain mathematical results quickly (perhaps reacting in a time sensitive setting) and use those results without explanation could apply an extended routine that uses a calculator. Extended routines may also be applied when real-world objects (external to the agent) must be manipulated to find a solution. The problem context will ultimately determine which category is most appropriate.

Enforcing a problem context in SNePS can be done in two ways: (1) by extending the antecedent constraints for `ActPlans` and (2) by redefining the SNeRE `do-one` primitive act to guide Cassie’s decision-making.

Antecedent Constraints: An act such as `Add(x, y)` can only be meaningfully performed by Cassie if `x` and `y` are restricted to the domain of numbers. This is achieved via the `Number(x)` proposition. Given a contextual proposition `InContext(x)` indicating that the agent is in current context `x`, we might have:

```
all(x, y) ( {Number(x), Number(y)}
  &=> { (InContext(Explanation)
    => ActPlan(Add(x, y),
      CountAdd(x, y)) ),
    (InContext(ShowingWork)
    => ActPlan(Add(x, y),
      ArithAdd(x, y)) ),
    (InContext(Discovery)
    => ActPlan(Add(x, y),
      CalcAdd(x, y)) ) } ).
```

Unless each antecedent is satisfied, including the contextual proposition, Cassie will not consider acting upon an `ActPlan`. This is an *a priori* and “conscious” consideration (i.e., it will explicitly appear in Cassie’s chain of reasoning before any act is attempted). Further refinement of antecedent constraints will lead to more specific contextual behavior. For example, if an agent is told to add 2 and 3, `CountAdd` (i.e., counting from 2 for 3 numbers) will not be an expensive operation. On the other hand, if an agent is adding 123 and 728123, `CountAdd` is very costly (i.e., counting from 123 for 728000 numbers!). We might add antecedent constraints forcing an agent to perform `CountAdd` when the numbers are single digits (via a `Digit(x)` proposition), and `ArithAdd` for multi-digit numbers.

Redefining the do-one Act: Another technique for enabling context-driven action is to modify the `do-one` primitive act. When Cassie is able to derive multiple plans for an act, the SNeRE executive cycle schedules a `do-one` on the set of those plans. In its unmodified form, `do-one` selects one of the plans at random. It can be modified to make a more deliberative decision by incorporating path-based reasoning (Shapiro, 1991) to determine the current context (using an `InContext(x)` proposition as in the previous section) and the procedural category of each plan being considered. The latter piece of metaknowledge can be stored in a proposition `ActCategory(x, y)` as follows:

```
ActCategory(CountAdd(x, y), Semantic).
ActCategory(ArithAdd(x, y), Syntactic).
ActCategory(CalcAdd(x, y), Extended).
```

With an appropriately re-defined `do-one`, making an assertion about context guarantees that the appropriate procedure is chosen (e.g., selecting `CountAdd(x, y)` when `InCategory(Semantic)` is asserted).

In contrast to the method of adding antecedent constraints, this method is online and “subconscious” for Cassie (i.e., it happens after Cassie has decided to perform the general act $\text{Add}(x, y)$ and uses path-based inference to decide upon a plan in the PML layer).

Sensitivity to context (as provided by either of the above techniques) is a useful application of metacognition. It will also lead to a more human-like behavior of the arithmetic task (i.e., we tend to count-add on our fingers only for small numbers and sum up larger numbers in columns and rows).

Self Regulation

The capacity for control and self regulation is an important feature of mathematical metacognition: “How well do you keep track of what you’re doing when (for example) you’re solving problems, and how well (if at all) do you use the input from those observations to guide your problem solving action” (Schoenfeld, 1987, p. 190).

Given a standard algebraic problem such as $x - 2 = 0$, consider the following two solution procedures (solving for x):

A	B
$x - 2 = 0$	$x - 2 = 0$
$+2 \quad +2$	$-2 \quad -2$
$x = 2$	$x - 4 = -2$
	$+4 \quad +4$
	$x = 2$

Clearly procedure A is a more direct solution. Procedure B seems to be a roundabout way towards the solution. The “directness” or “roundaboutness” of a procedure is a metacognitive judgment. This judgment is based on the fact that procedure B uses more operations (of comparable complexity) than procedure A.

Goldfain (Goldfain, 2005) describes a technique with which SNePS agents can enumerate their own actions *while they are acting*. By counting up the arithmetic routines in a mathematical task (e.g., an algebraic task), an agent can get a rough sense of a plan’s efficiency. An agent must be provided with enough logic to establish which operations are comparable in complexity (this is even more metacognitive information).

A math-capable agent should have some sense of whether each step in its reasoning is bringing it any closer to a solution. For the example above, adding y to both sides to produce $x - 2 + y = y$ leaves us with something more complex than the original problem. An agent who performs such a step should stop in its tracks; detecting that it is moving “farther” away from an answer. We would never teach a student a rule to *always* make every line of work ‘simpler’ than the previous one. However, there are several mathematical procedures where this would at least be a good rule-of-thumb. A SNePS implementation of this ability is beyond the scope of this paper, but it is worth further investigation.

Conclusions

The SNePS knowledge representation, reasoning, and acting system has several features that facilitate metacognition in SNePS-based agents. The most prominent is the fact that propositions are represented in SNePS as terms rather than as logical sentences. The effect is that propositions can occur as arguments of propositions, acts, and policies without limit, and without leaving first-order logic. The SNePS acting subsystem is integrated with the SNePS reasoning subsystem in such a way that there are acts (believe and disbelieve) that affect what an agent believes, there are acts (snif, sniterate, withall, witsome) that specify both knowledge-contingent acts and lack-of-knowledge default acts, and there are policies (ifdo, wheno, wheneverdo) that serve as “daemons”, triggering acts when certain propositions are believed or wondered about.

The GLAIR agent architecture supports metacognition by specifying the PML as the source of self-awareness. Affective and effective actions implemented in the PML are the source of first-person privileged knowledge about what the agent is sensing and doing. Deictic and modality registers located in the PML provide the agent a sense of situatedness in its world.

Several SNePS-based agents have taken advantage of these facilities to engage in self-awareness and metacognition: Fevahr Cassie has a model of herself and episodic memory of what she has done; a SNePS agent uses lack-of-knowledge acting to acquire information it needs to make a telephone call; an enhanced SNePS belief revision system can use belief credibility and source credibility information, represented as metaknowledge, to do automatic belief revision on knowledge from multiple sources; a SNePS agent modifies its prior beliefs of word meaning to accommodate unusual usages in a text; a version of Cassie uses knowledge of the context to choose an appropriate method of arithmetic. The ongoing implementation and development of SNePS may allow for further applications of metacognition.

Acknowledgments

This work was supported, in part, by the US Army Communications and Electronics Command (CECOM), Ft. Monmouth, NJ through contracts with CACI Technologies and Boozee-Allen & Hamilton. The authors appreciate the comments of other current members of the SNePS Research Group, and the contributions of the many past members.

References

- Alchourrón, C. E., Gärdenfors, P., and Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2):510–530.
- Ali, S. S. and Shapiro, S. C. (1993). Natural language processing using a propositional semantic network with structured variables. *Minds and Machines*, 3(4):421–451.
- Allen, J. F. (1983). Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11):832–843.
- Benferhat, S., Kaci, S., Berre, D. L., and Williams, M.-A. (2004). Weakening conflicting information for iterated revision and knowledge integration. *Artificial Intelligence*, 153(1–2):339–371.
- Brachman, R. J. and Levesque, H. J., editors (1985). *Readings in Knowledge Representation*. Morgan Kaufmann, San Mateo, CA.
- Chalupsky, H. and Shapiro, S. C. (1994). SL: A subjective, intensional logic of belief. In *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*, pages 165–170, Hillsdale, NJ. Lawrence Erlbaum.
- Cohors-Fresenborg, E. and Kaune, C. (2001). Mechanisms of the taking effect of metacognition in understanding processes in mathematics teaching. In *Developments in Mathematics Education in German-speaking Countries, Selected Papers from the Annual Conference on Didactics of Mathematics*, pages 29–38. SUB Göttingen, Ludwigsborg.
- Cox, M. T. (2005). Metacognition in computation: A selected research review. *Artificial Intelligence*, 169(2):104–141.
- de Kleer, J. (1986). An assumption-based truth maintenance system. *Artificial Intelligence*, 28(2):127–162.
- Duchan, J. F., Bruder, G. A., and Hewitt, L. E., editors (1995). *Deixis in Narrative: A Cognitive Science Perspective*. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.
- Ehrlich, K. (1995). *Automatic Vocabulary Expansion through Narrative Context*. PhD dissertation, Technical Report 95-09, Department of Computer Science, State University of New York at Buffalo, Buffalo, NY.
- Ehrlich, K. and Rapaport, W. J. (1997). A computational theory of vocabulary expansion. In *Proceedings of the 19th Annual Conference of the Cognitive Science Society*, pages 205–210, Mahwah, NJ. Lawrence Erlbaum Associates.
- Forbus, K. D. and de Kleer, J. (1993). *Building Problem Solvers*. MIT Press, Cambridge, MA.

- Gärdenfors, P. and Rott, H. (1995). Belief revision. In Gabbay, D. M., Hogger, C. J., and Robinson, J. A., editors, *Handbook of Logic in Artificial Intelligence and Logic Programming*, volume 4, pages 35–132. Clarendon Press, Oxford.
- Gartmann, S. and Freiberg, M. (1995). Metacognition and mathematical problem solving: Helping students to ask the right questions. *The Mathematics Educator*, 6(1):9–13.
- Gentner, D. (1981). Some interesting differences between nouns and verbs. *Cognition and Brain Theory*, 4:161–178.
- Goldfain, A. (2005). Counting and early mathematical understanding. <http://www.cse.buffalo.edu/~ag33/CAEMU.pdf>.
- Hansson, S. O. (1997). Semi-revision. *Journal of Applied Non-Classical Logic*, 7:151–175.
- 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., editors, *Proceedings of The Sixth Florida AI Research Symposium (FLAIRS 93)*, pages 325–329. The Florida AI Research Society.
- 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., editors, *Expertise in Context*, pages 383–404. AAAI Press/MIT Press, Cambridge, MA.
- Ismail, H. O. (2001). *Reasoning and Acting in Time*. PhD dissertation, Technical Report 2001-11, University at Buffalo, The State University of New York, Buffalo, NY.
- Iwańska, Ł. M. and Shapiro, S. C. (2000). *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language*. AAAI Press/The MIT Press, Menlo Park, CA.
- Johnson, F. L. (2006). *Dependency-Directed Reconsideration: An Anytime Algorithm for Hindsight Knowledge-Base Optimization*. PhD dissertation, Department of Computer Science and Engineering, University at Buffalo, The State University of New York, Buffalo, NY.
- Johnson, F. L. and Shapiro, S. C. (1999). Says who?—incorporating source credibility issues into belief revision. Technical Report 1999-08, Department of Computer Science and Engineering, University at Buffalo, The State University of New York, Buffalo, NY.
- Johnson, F. L. and Shapiro, S. C. (2005a). Dependency-directed reconsideration: Belief base optimization for truth maintenance systems. In *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI-05)*, pages 313–320. AAAI Press, Menlo Park, CA.
- Johnson, F. L. and Shapiro, S. C. (2005b). Improving recovery for belief bases. In Morgenstern, L. and Pagnucco, M., editors, *IJCAI-05 Workshop on Nonmonotonic Reasoning, Action, and Change (NRAC'05): Working Notes*, pages 65–70. IJCAI, Edinburgh, Scotland.
- Kibby, M. W. and Rapaport, W. J. (2007). Contextual vocabulary acquisition as computational philosophy and as philosophical computation. *Journal of Experimental and Theoretical Artificial Intelligence*. forthcoming.
- Kumar, D. (1996). The SNePS BDI architecture. *Decision Support Systems*, 16(1):3–19.
- Kumar, D. and Shapiro, S. C. (1994a). Acting in service of inference (and *vice versa*). In Dankel II, D. D., editor, *Proceedings of The Seventh Florida AI Research Symposium (FLAIRS 94)*, pages 207–211. The Florida AI Research Society.
- Kumar, D. and Shapiro, S. C. (1994b). The OK BDI architecture. *International Journal on Artificial Intelligence Tools*, 3(3):349–366.
- Lehmann, F., editor (1992). *Semantic Networks in Artificial Intelligence*. Pergamon Press, Oxford.
- Maida, A. S. and Shapiro, S. C. (1982). Intensional concepts in propositional semantic networks. *Cognitive Science*, 6(4):291–330. Reprinted in (Brachman and Levesque, 1985, pp. 170–189).

- Malory, T. (1470). *Le Morte Darthur*. Collier Books, New York. Edited by R. M. Lumiansky, 1982. Page references are to the 1982 edition.
- Martins, J. P. and Shapiro, S. C. (1983). Reasoning in multiple belief spaces. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 370–373. Morgan Kaufmann, San Mateo, CA.
- Martins, J. P. and Shapiro, S. C. (1988). A model for belief revision. *Artificial Intelligence*, 35:25–79.
- McCarthy, J. (1977). Epistemological problems of artificial intelligence. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence (IJCAI-77)*, volume 2, pages 1038–1044. International Joint Conferences on Artificial Intelligence.
- McNamara, D. S. and Kintsch, W. (1996). Learning from texts: Effects of prior knowledge and text coherence. *Discourse Processes*, 22:247–288.
- Moore, R. C. (1988). Autoepistemic logic. In Smets, P., Mamdani, E. H., Dubois, D., and Prade, H., editors, *Non-Standard Logics for Automated Reasoning*, pages 105–127. Academic Press, London.
- Moses, L. J. and Baird, J. A. (1999). Metacognition. In Wilson, R. A. and Keil, F., editors, *The MIT Encyclopedia of the Cognitive Sciences*, pages 533–535. MIT Press, Cambridge, MA.
- Orilia, F. and Rapaport, W. J., editors (1998). *Thought, Language, and Ontology: Essays in Memory of Hector-Neri Castañeda*. Kluwer Academic Publishers, Dordrecht.
- Pape, D., Anstey, J., Dolinsky, M., and Dambik, E. J. (2003). Ygdrasil—a framework for composing shared virtual worlds. *Future Generation Computer Systems*, 19(6):1041–1049.
- Rapaport, W. J. (2003). What is the ‘context’ for contextual vocabulary acquisition? In Slezak, P. P., editor, *Proceedings of the 4th International Conference on Cognitive Science/7th Australasian Society for Cognitive Science Conference (ICCS/ASCS-2003)*, volume 2, pages 547–552, Sydney, Australia.
- Rapaport, W. J. (2005). In defense of contextual vocabulary acquisition: How to do things with words in context. In Dey, A., Kokinov, B., Leake, D., and Turner, R., editors, *Modeling and Using Context: 5th International and Interdisciplinary Conference (Context 2005)*, Lecture Notes in Artificial Intelligence 3554, pages 396–409, Berlin. Springer-Verlag.
- Rapaport, W. J. and Ehrlich, K. (2000). A computational theory of vocabulary acquisition. In (Iwańska and Shapiro, 2000), pages 347–375.
- 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 and Rapaport, 1998, pp. 235–294).
- Ruddell, R. B. and Unrau, N. J. (2004). *Theoretical Models and Processes of Reading*. International Reading Association, Newark, DE, 5th edition.
- Russell, S. J. (1999). Metareasoning. In Wilson, R. A. and Keil, F., editors, *The MIT Encyclopedia of the Cognitive Sciences*, pages 539–541. MIT Press, Cambridge, MA.
- Santore, J. F. and Shapiro, S. C. (2003). Crystal Cassie: Use of a 3-D gaming environment for a cognitive agent. In Sun, R., editor, *Papers of the IJCAI 2003 Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions*, pages 84–91. IJCAI.
- Schoenfeld, A. H. (1987). What’s all the fuss about metacognition? In Schoenfeld, A., editor, *Cognitive Science and Mathematics Education*, pages 189–215. Lawrence Erlbaum Associates, Hillsdale, NJ.
- Schoenfeld, A. H. (1992). Learning to think mathematically: Problem solving, metacognition, and sense-making in mathematics. In Grouws, D., editor, *Handbook for Research on Mathematics Teaching and Learning*, pages 334–370. MacMillan, New York.

- Shapiro, S. C. (1971). A net structure for semantic information storage, deduction and retrieval. In *Proceedings of the Second International Joint Conference on Artificial Intelligence*, pages 512–523. Morgan Kaufmann, San Mateo, CA.
- Shapiro, S. C. (1979). The SNePS semantic network processing system. In Findler, N. V., editor, *Associative Networks: The Representation and Use of Knowledge by Computers*, pages 179–203. Academic Press, New York.
- Shapiro, S. C. (1986). Symmetric relations, intensional individuals, and variable binding. *Proceedings of the IEEE*, 74(10):1354–1363.
- Shapiro, S. C. (1989). The CASSIE projects: An approach to natural language competence. In Martins, J. P. and Morgado, E. M., editors, *EPIA 89: 4th Portuguese Conference on Artificial Intelligence Proceedings, Lecture Notes in Artificial Intelligence 390*, pages 362–380. Springer-Verlag, Berlin.
- Shapiro, S. C. (1991). Cables, paths and “subconscious” reasoning in propositional semantic networks. In Sowa, J., editor, *Principles of Semantic Networks: Explorations in the Representation of Knowledge*, pages 137–156. Morgan Kaufmann, Los Altos, CA.
- 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*, pages 136–143. AAAI Press, Menlo Park, California.
- Shapiro, S. C. (2000a). An introduction to SNePS 3. In Ganter, B. and Mineau, G. W., editors, *Conceptual Structures: Logical, Linguistic, and Computational Issues*, volume 1867 of *Lecture Notes in Artificial Intelligence*, pages 510–524. Springer-Verlag, Berlin.
- Shapiro, S. C. (2000b). SNePS: A logic for natural language understanding and commonsense reasoning. In (Iwańska and Shapiro, 2000), pages 175–195.
- Shapiro, S. C. (2004). A logic of arbitrary and indefinite objects. In Dubois, D., Welty, C., and Williams, M.-A., editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Ninth International Conference (KR2004)*, pages 565–575, Menlo Park, CA. AAAI Press.
- Shapiro, S. C. and Ismail, H. O. (2003). Anchoring in a grounded layered architecture with integrated reasoning. *Robotics and Autonomous Systems*, 43(2–3):97–108.
- Shapiro, S. C. and Johnson, F. L. (2000). Automatic belief revision in SNePS. In Baral, C. and Truszczyński, M., editors, *Proceedings of the 8th International Workshop on Non-Monotonic Reasoning NMR2000*. Available as [arXiv:cs.AI/0003011](https://arxiv.org/abs/cs.AI/0003011), March 6, 2000.
- Shapiro, S. C. and Kandefer, M. (2005). A SNePS approach to the wumpus world agent or Cassie meets the wumpus. In Morgenstern, L. and Pagnucco, M., editors, *IJCAI-05 Workshop on Nonmonotonic Reasoning, Action, and Change (NRAC’05): Working Notes*, pages 96–103. IJCAI, Edinburgh, Scotland.
- Shapiro, S. C., Martins, J. P., and McKay, D. P. (1982). Bi-directional inference. In *Proceedings of the Fourth Annual Meeting of the Cognitive Science Society*, pages 90–93, Ann Arbor, MI.
- Shapiro, S. C. and Rapaport, W. J. (1987). SNePS considered as a fully intensional propositional semantic network. In Cercone, N. and McCalla, G., editors, *The Knowledge Frontier*, pages 263–315. Springer-Verlag, New York.
- Shapiro, S. C. and Rapaport, W. J. (1991). Models and minds: Knowledge representation for natural-language competence. In Cummins, R. and Pollock, J., editors, *Philosophy and AI: Essays at the Interface*, pages 215–259. MIT Press, Cambridge, MA.
- 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 The SNePS Implementation Group (2004). *SNePS 2.6.1 User's Manual*. Department of Computer Science and Engineering, University at Buffalo, The State University of New York, Buffalo, NY.
- Sperber, D. (1999). Metarepresentation. In Wilson, R. A. and Keil, F., editors, *The MIT Encyclopedia of the Cognitive Sciences*, pages 541–543. MIT Press, Cambridge, MA.
- van Oostendorp, H. and Goldman, S. R. (1999). *The Construction of Mental Representations During Reading*. Lawrence Erlbaum Associates, Mahwah, NJ.
- Wiebe, J. M. and Rapaport, W. J. (1986). Representing *de re* and *de dicto* belief reports in discourse and narrative. *Proceedings of the IEEE*, 74(10):1405–1413.