Two Problems with Reasoning and Acting in Time*

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

Abstract

Natural language competent embodied cognitive agents should satisfy two requirements. First, they should act in and reason about a changing world, using reasoning in the service of acting and acting in the service of reasoning. Second, they should be able to communicate their beliefs, and report their past, ongoing, and future actions in natural language. This requires a representation of time using a deictic NOW, that models the compositional semantic properties of the English "now". Two problems emerge for an agent that interleaves reasoning and acting in a personal time. The first concerns the representation of plans and reactive rules involving reasoning about "future NOWs". The second emerges when, in the course of reasoning about NOW, the reasoning process itself results in NOW changing. We propose solutions for the two problems and conclude that: (i) for embodied cognitive agents, time is not just the object of reasoning, but is embedded in the reasoning process itself; and (ii) at any time, there is a partonomy of NOWs representing the agent's sense of the current time at different levels of granularity.

*To appear as H. O. Ismail and S. C. Shapiro, Two Problems with Reasoning and Acting in Time. In A. G. Cohn, F. Giunchiglia, & B. Selman, Eds., *Principles of Knowledge Representation and Reasoning: Proceedings of the Seventh International Conference (KR 2000),* Morgan Kaufmann, San Francisco, 2000. All citations should be to the published version.

1 Introduction: *In* and *About* Time

Reasoning about time is something that agents acting in the world ought to be capable of doing. Performing an act *before* another, achieving states that must hold *while* an act is being performed, or reasoning about states that should hold *after* the performance of an act involve, whether implicitly or explicitly, some degree of reasoning about time. Temporal logics are used for reasoning about time and discussing its properties in a precise and explicit manner (van Benthem, 1983, for instance). In these logics, time just happens to be the subject matter of some of its sentences. Except for the presence of terms, predicates, and operators denoting temporal entities and relations; there is nothing about the language that is intrinsically temporal. For example, the logics developed in (van Benthem, 1983) might be applied to one-dimensional space, the rational numbers, or the integers just by changing the denotation of some symbols. Being about time is an extrinsic property of a logic; it only determines the domain of interpretation, maybe the syntax, but not the interpretation and reasoning processes. More specifically, let Δ be a collection of logical formulas (i.e., a knowledge base) and let $\mathcal{A}$ be an acting and reasoning system reasoning with the information in Δ. In particular, think of $\mathcal{A}$ as an embodied cognitive agent acting in the world and of Δ as the contents of its memory. $\mathcal{A}$ is said to be reasoning *about* time if the semantics of some of the sentences in Δ refer to temporal individuals and properties.[1] The assumption here is that this is an accidental situation; the design of the inference rules used by $\mathcal{A}$ is tuned only to the syntax (the domain in which inference takes place) and the semantics of the logical connectives and operators. The semantics of functional terms and predicates, and the reasoning being about time has no effect on how $\mathcal{A}$'s

[1] Of course, this is a very liberal characterization of what reasoning about time is.

inference engine operates.

Not only may reasoning be about time, it could also be in time. What does reasoning *in* time mean? In the technical sense in which we want to interpret "in" and in the context of Δ and $\mathcal{A}$ from above, it means two things.

1. Temporal progression is represented in Δ. That is, at any point, there is a term in Δ which, for the agent $\mathcal{A}$, denotes the current time. Which term denotes the current time depends on *when* one inspects Δ.[2] This gives the agent "a personal sense of time" (Shapiro, 1998, p. 141).

2. Reasoning takes time. By that we do not simply mean the obvious fact that any process happens over an interval of time. What we mean is that it happens over an interval of $\mathcal{A}$'s time. In other words, the term in Δ denoting the current time at the beginning of the reasoning process is different from that denoting the current time at the end.

As we shall argue below, if one were to take the issue of reasoning and acting in time seriously, problems immediately emerge. We are going to present two problems that naturally arise when dealing with a cognitive agent reasoning and acting in certain situations. These are, by no means, unrealistic or exotic situations; they involve simple acting rules and behaviors that agents are expected to be able to exhibit. The main point is that, when it comes to embodied cognitive agents, time is not just a possible object of reasoning; it is deeply embedded into the agent's reasoning processes.

2 The Agent

In this section, we briefly highlight certain design constraints that we impose on our developing theory of agents. Our theory is based on the GLAIR agent architecture (Hexmoor et al., 1993; Hexmoor and Shapiro, 1997). This is a layered architecture, the top layer of which is responsible for high level cognitive tasks such as reasoning and natural language understanding. This level is implemented using the SNePS knowledge representation and reasoning system (Shapiro and Rapaport, 1987; Shapiro and Rapaport, 1992; Shapiro and the SNePS Implementation Group, 1999).

We use "Cassie" as the name of our agent. Previous versions of Cassie have been discussed elsewhere (Hexmoor, 1995; Shapiro, 1989). Those are actually various incarnations of the disembodied linguistically-competent cognitive agent of the SNePS system (Shapiro and Rapaport, 1987; Shapiro and Rapaport, 1995). There are four basic requirements that we believe are reasonable for a theory of embodied cognitive agents.

R1. **Reasoning in service of acting and acting in service of reasoning.** Cassie uses reasoning in the service of acting in order to decide when, how, and/or whether to act in a certain manner. Similarly, Cassie may act in order to add a missing link to a chain of reasoning. For example, conclusions about the state of the world may be derived based, not only on pure reasoning, but also on looking, searching and performing various sensory acts. For more on this see (Kumar and Shapiro, 1994).

R2. **Memory.** Cassie has a record of what she did and of how the world evolved. A memory of the past is important for reporting to others what happened. This, as shall be seen, constrains the *form* of certain sentences in the logic.

R3. **Natural language competence.** Cassie should be capable of using natural language to interact with other agents (possibly human operators). This means that SNePS representations of the contents of Cassie's memory ought to be linguistically-motivated. By that we mean two things. First, on the technical side, the representations should be designed so that they may be produced by a natural language understanding system, and may be given as input to a natural language generator. Second, at a deeper level, the syntax of the representations and the underlying ontology should reflect their natural language (in our case, English) counterparts. In particular, we admit into the SNePS ontology anything that we can think or talk about (Shapiro and Rapaport, 1987; Shapiro and Rapaport, 1992). For a general review of linguistically-motivated knowledge representation, see (Iwańska and Shapiro, 2000).

R4. **Reasoning in time.** Cassie has a personal sense of time (Shapiro, 1998); at any point, there is a term in the logic that, for Cassie, represents the current time. To represent temporal progression, we use a deictic NOW pointer (Almeida and Shapiro, 1983; Almeida, 1995)– a meta-logical variable that assumes values from amongst the time-denoting terms in the logic.[3] We will use

[2] Note that this means that an agent reasoning in time also reasons about time.

[3] A similar idea has been suggested by (Allen, 1983).

"*NOW" to denote the term which is the value of the meta-logical variable NOW. There are four things to note. First, NOW is not a term in the logic, just a meta-logical variable. Second, "*NOW" is not itself a fixed term in the logic; at any point, it is a shorthand for the term denoting the current time.[4] Third, to maintain a personal sense of time, the value of NOW changes to a new term when, and only when, Cassie acts.[5] Note that this does not preclude Cassie's learning about events that happened between or during times that once were values of NOW. Fourth, for R3, the behavior of *NOW models the compositional semantic properties of the English "now".[6] It is always interpreted as the time of the utterance (or the assertion), it cannot refer to past nor to future times (Prior, 1968; Kamp, 1971; Cresswell, 1990). Note that, given R3, Cassie essentially reasons *in* time, in the sense of Section 1.

Two incarnations of embodied Cassies have been developed based on the above requirements. In the FEVAHR project (Shapiro, 1998) Cassie played the role of a "Foveal Extra-Vehicular Activity Helper-Retriever (FEVAHR)." 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. There have also been several software simulatated versions of the FEVAHR. Cassie, the FEVAHR, operates in a $17' \times 17'$ room containing: Cassie; Stu, a human supervisor; Bill, another human; a green robot; and three indistinguishable red robots. Cassie is always talking to either Stu or Bill—taking statements, questions, and commands from that person (all expressed in a fragment of English), and responding and reporting to that person in English. Cassie can be told, by the person addressing her, to talk to the other person, or to find, look at, go to, or follow any of the people or robots in the room. Cassie can also engage in conversations on a limited number of other topics in a fragment of English, similar to some of the conversations in (Shapiro, 1989).

A more recent incarnation of embodied Cassie is as a robot that clears a field of unexploded ordnance (UXO remediation). This Cassie has only existed as a software simulation. The UXO-Cassie exists in an area consisting of four zones: a safe zone; an operating zone that possibly contains UXOs; a drop-off zone; and a recharging zone. The UXO-Cassie contains a battery that discharges as she operates, and must be recharged in the recharge zone as soon as it reaches a low enough level. She may carry charges to use to blow up UXOs. Her task is to search the operating zone for a UXO, and either blow it up by placing a charge on it, and then going to a safe place to wait for the explosion, or pick up the UXO, take it to the drop-off zone, and leave it there. The UXO-Cassie has to interrupt what she is doing whenever the battery goes low, and any of her actions might fail. (She might drop a UXO she is trying to pick up.) She takes direction from a human operator in a fragment of English, and responds and reports to that operator. There is a large overlap in the grammars of Cassie, the FEVAHR, and the UXO-Cassie.

The requirements listed above, which we believe are quite reasonable, have certain representational and ontological impacts on the formal machinery to be employed. As we have found in our experiments with Cassie, the FEVAHR, and the UXO-Cassie, and as we shall show below, this leads to problems with reasoning with the deictic NOW. Before setting out to discuss these problems, let us first introduce the basic logical infra-structure.

3 The Ontology of a Changing World

3.1 Change

Traditionally, there have been two main models for representing change. First, there is the STRIPS model of change (Fikes and Nilsson, 1971) where, at any time t, the only propositions in the knowledge base are about those states that hold at t. When time moves, propositions are added and/or deleted to reflect the new state of the world. The main obvious problem is that an agent based on such a system does not have any memory of past situations (thus violating R2). Second, there are variants of the situation calculus (McCarthy and Hayes, 1969) where propositions are associated with indicators to when they hold. Indicators may be situations as in the situation calculus (terms denoting instantaneous snapshots of the world) or time-denoting terms (Allen, 1983; Shoham, 1987). In what follows, we shall adopt the second approach for representing change. In particular, our chronological indicators shall be taken to denote times— a decision that is rooted in R3 and R4 from Section 2.

[4]In Kaplan's terms (Kaplan, 1979; Braun, 1995), only *contents*, not *characters*, are represented in the knowledge base.

[5]More generally, this should happen whenever Cassie recognizes a change in the world, including changes in her own state of mind.

[6]Unlike the **now** of (Lespérance and Levesque, 1995).

3.2 States

So far, we have been a little sloppy with our terminology. In particular, we have been using the terms "state" and "proposition" interchangeably. First, let us briefly explain what we mean by "proposition". Propositions are entities that may be the object of Cassie's belief (or, in general, of any propositional attitude). We assume propositions to be first-class entities in our ontology. Cassie's *belief space* is essentially a set of propositions— those that she believes (Shapiro, 1993).

Cassie's beliefs are about *states* holding over time. At any given time, a given state may either hold or not hold. The notion of states referred to here is that found in the linguistics and philosophy of language literature (Vendler, 1957; Mourelatos, 1978; Galton, 1984; Bach, 1986, for instance).[7] A particularly interesting logical property of states is homogeneity; if a state s holds over some time interval t, then it holds over all subintervals of t.

States differ as to their degree of *temporal stability.* Here, we are interested in two types of states: *eternal* and *temporary.* Eternal states are related to the eternal sentences of (Quine, 1960); they either *always* hold or *never* hold. Temporary states, on the other hand, may repetitively start and cease to hold. Examples of eternal states are expressible by sentences such as *God exists*, *Whales are fish*, or *The date of John's graduation is September 5th 2001.* Temporary states are expressed by sentences such as *The litmus paper is red*, *John is in New York*, or *John is running.*

Temporary states starting or ceasing to hold are responsible for changes in the world and, hence, need to be associated with times (see Section 3.1). This association is formally achieved by introducing a function symbol, Holds, that denotes a function from temporary states and times to propositions. Thus, $\mathsf{Holds}(s, t)$ denotes the proposition that state $[\![s]\!]$ holds over time $[\![t]\!]$.[8] Note that this is similar to the situation calculus with reified fluents.

Eternal states do not change with time and, hence, should not be associated with any particular times.[9] If anything, one would need a unary function to map eternal states into propositions. More ontological economy may be achieved though if we make some observations. First, note that, unlike temporary states, eternal states cannot start, cease, or be perceived. They may only be believed to be holding, denied to be holding, asserted to be holding, wished to be holding, etc. That is, an agent can only have propositional attitudes toward eternal states. This means that the set of eternal states is isomorphic to a subset of the set of propositions. Second, all propositions may be thought of to be about eternal states holding. For example, $\mathsf{Hold}(s, t)$ may be thought of as denoting the eternal state of some particular temporary state holding at some particular time. Note that this is eternal since it is either always the case or never the case. Henceforth, we shall make the assumption that eternal states are identical to propositions and will use the two terms interchangeably. In this case, we do not need any functions to map eternal states, thus simplifying our syntax.

3.3 Time

As has been hinted above, we opt for an interval-based ontology of time.[10] We introduce two functional symbols, $<$ and $\sqsubseteq$, to represent the relations of temporal precedence and temporal parthood. More precisely, $t_1 < t_2$ denotes the proposition that $[\![t_1]\!]$ precedes (and is topologically disconnected from) $[\![t_2]\!]$ and $t_1 \sqsubseteq t_2$ denotes the proposition that $[\![t_1]\!]$ is a subinterval of $[\![t_2]\!]$. Because of its homogeneity, a state will be said to hold *NOW if it holds over a super-interval of *NOW.

4 The Problem of the Unmentionable Now

4.1 The Problem

How does introducing the eternal-temporary distinction affect the reasoning process? Consider the following sentence schema:

(1) IF *ant* THEN *cq*

(1) means that if Cassie believes that *ant*[11] holds, then she may also believe that *cq* holds.[12] This works fine if *ant* and *cq* denote eternal states (for example, "IF Mammals(whales) THEN Give-Birth(whales)"). However, if, instead, they denote temporary states, we need

[7] As opposed to the states of (McDermott, 1982) and (Shanahan, 1995) which are more like time points or situations of the situation calculus.

[8] If τ is a term in the logic, we use $[\![\tau]\!]$ to mean the denotation of τ.

[9] Though, with time, Cassie may revise her beliefs about eternal states (Martins and Shapiro, 1988).

[10] See (Allen, 1983) and (Shoham, 1985) for arguments in support of interval semantics.

[11] For convenience, we shall, henceforth, write p in place of $[\![p]\!]$ whenever what we mean is clear from context.

[12] "may" because the rule might not fire, even if Cassie believes that *ant* holds.

to quantify over time; the temporary state-denoting terms by themselves do not say anything about the states holding over time. (2) captures the intended meaning: if Cassie believes that *ant* holds over time t, then she may also believe that *cq* holds over t.

(2) $\forall t$ IF Holds(*ant*, t) THEN Holds(*cq*, t)

(1) and (2) represent sentences for *pure* reasoning. Doing reasoning in the service of acting requires sentences for *practical* reasoning. In particular, let us concentrate on one kind of acting rule: rules about *when* to act.[13] Imagine Cassie operating in a factory. One reasonable belief that she may have is that, when the fire-alarm sounds, she should leave the building. The underlying schema for such a belief is represented in (3) (Kumar and Shapiro, 1994).

(3) When *cond* DO *act*

The intended interpretation for (3) is that when Cassie comes to believe that the condition *cond* holds, she should perform the act *act*. Again, this is fine so long as *cond* denotes an eternal state. If forward inference causes both *cond* and (3) to be asserted in Cassie's belief space, she will perform *act*. What if *cond* denotes a temporary state? Obviously, we need to somehow introduce time since assertions about temporary states holding essentially involve reference to time. Following (2), one may propose the following representation.

(4) $\forall t$ When Holds(*cond*, t) DO *act*

Asserting Holds(*cond*, t_1), for some particular time t_1, (4) would be matched and Cassie would perform *act*. On the face of it, (4) looks very innocent and a straight forward extrapolation of (2). However, a closer look shows that this is, by no means, the case. Using quantification over time works well for inference since the consequent is a proposition that may just happen to be *about* time. Acting, on the other hand, takes place *in* time, resulting in an interesting problem. Table 1 represents a timed sequence of assertions entered into Cassie's belief space. The left column shows the assertion, and the right column shows Cassie's term for the time of the assertion. The problem is that t_0 in (6) may refer to a time preceding t_2 (or even t_1). That is, (6) could be an assertion about the alarm sounding

Assertion		Assertion Time
(5)	$\forall t$ When Holds(Sounds(alarm), t) DO Leave(building)	t_1
(6)	Holds(Sounds(alarm), t_0)	t_2

Table 1: A timed sequence of assertions for the fire-alarm problem.

at some time in the past, something that we should be capable of asserting. Nevertheless, (6) matches (5) and Cassie leaves the building —at t_2— even though there is no danger!

One problem with (5) (and generally (4)) is that nothing relates the time of performing the act to the time at which the state holds. We may attempt to revise (4) by tying the action to that time.

(7) $\forall t$ When Hold(*cond*, t) DO Perform(*act*, t)

Where Perform(*act*, t) is intended to mean that Cassie should perform *act* at time t. However, this alleged semantics of Perform is certainly ill-defined; acts may only be performed *NOW, in the present. Cassie cannot travel in time to perform *act* in the past, at a time over which (she believes that) *cond* held. The basic problem seems to be quantifying over *all* times. What we really want to say is that when the state holds *NOW, perform the act. That is,

(8) When Hold(*cond*, *NOW) DO *act*

However, we cannot mention "*NOW"; it is not itself a term in the logic (see R4 in Section 2). If we replace (5) in Table 1 with the appropriate variant of (8), "*NOW" in the left column would be just a shorthand for the term appearing in the right column, namely t_1. The assertion would, therefore, be very different from what we intended it to be.

Before presenting our solution to the problem, we first need to discuss two approaches that might seem to solve it. We shall show that, although they may appear to eliminate the problem, they actually introduce more drastic ones.

4.2 A NOW Function

The basic problem, as we have shown, is that we cannot mention NOW; there is no unique term in the logic that would, at any point, denote the current time for the agent. The existence of such a term is problematic since its semantics essentially changes with time.

[13] By "rule" we mean a domain rule, expressed in the logical language, which Cassie might come to believe as a result of being told it in natural language. We do *not* mean a rule of inference which would be implemented in the inference engine of the knowledge representation and reasoning system.

One way to indirectly incorporate NOW within the language is to introduce a NOW function symbol. In particular, the expression $\mathsf{NOW}(t)$ would mean that t denotes the current time. Thus, one may express the general acting rule as follows:

(9) $\forall t$ When ($\mathsf{Hold}(cond,\ t) \wedge \mathsf{NOW}(t)$) DO *act*

This might seem to solve the problem, for it necessitates that the time at which the state holds is *a current time* (and at any time, there is a unique one). There are two problems though.

1. $\mathsf{NOW}(t)$ denotes a temporary state. By its very definition, the argument of NOW needs to change to reflect the flow of time. Thus, rather than using $\mathsf{NOW}(t)$, we should use $\mathsf{Holds}(\mathsf{NOW}(t),\ t)$ to express the proposition that $[\![t]\!]$ is the current time. This gets us back where we started, since the expression $\mathsf{Holds}(\mathsf{NOW}(t),\ t)$ would have to replace $\mathsf{NOW}(t)$ in (9). An assertion of $\mathsf{Holds}(\mathsf{NOW}(t_0)$, $t_0)$ with t_0 denoting some past time will cause the agent to perform *act* when it shouldn't.

2. Suppose that, at t_1, Cassie is told that John believes that t_2 is the current time. That is, "Believe(John, $\mathsf{NOW}(t_2)$)" is asserted. At time t_3, the same assertion provides different information about John. In particular, it attributes to John a belief that was never intended to be asserted into Cassie's belief space. The general problem is that, as time goes by, Cassie needs to revise her beliefs. Those may be her own, or other agents', beliefs about what the current time is. In the first case, the revision may be the simple deletion of one belief and introduction of another. In the second case, however, things are much more complicated as demonstrated by the above example. It should be noted that, in any case, the very idea of Cassie *changing her mind* whenever time moves is, at best, awkward, and results in Cassie *forgetting* correct beliefs that she once had (thus violating R2).

4.3 The Assertion Time

Inspecting the second row of Table 1, one may think that part of the problem is the inequality of the time appearing in the right column (t_2) to that in the left column (t_0). Indeed, if somehow we can ensure that these two times are identical, the problem may be solved. (Kamp, 1971) proposes an ingenious mechanism for correctly modeling the compositional properties of the English "now" (namely, that it always refers to the time of the utterance even when embedded within the scope of tense operators). Basically, Kamp defines the semantic interpretation function relative to two temporal indices rather than only one as in traditional model theory. The two times may be thought of as the Reichenbachian event and speech times (Reichenbach, 1947). We shall not review Kamp's proposal here; rather, based on it, we shall introduce an approach that might seem to solve the problem.

The basic idea is to move the assertion time appearing in the right column of Table 1 to the left column. That is, to formally *stamp* each assertion with the time at which it was made. Formally, we introduce a symbol Asserted that denotes a function from propositions and times to propositions. Thus, "$\mathsf{Asserted}(p,\ t_a)$" denotes the proposition that Cassie came to believe p at t_a. We then replace (4) by (10).

(10) $\forall t$ When $\mathsf{Asserted}(\mathsf{Holds}(cond,\ t),\ t)$ DO *act*

That is, Cassie would only perform *act* when she comes to believe that *cond* holds, at a time at which it actually holds. This will indeed not match any assertions about past times and apparently solves the problem. However, there are at least two major problems with this proposal.

1. Introducing the assertion time results in problems with simple implications as that in (1). In particular, due to its semantics, the assertion time of the antecedent need not be that of the consequent; one may come to believe in *ant* at t_1 and infer *cq* later at t_2. The problem is that the time at which the inference is made cannot be known in advance. Essentially, this is the same problem that we started with; we only know that the inference will be made at some unmentionable future *NOW.

2. So far, we have only discussed the problem in the context of forward chaining. The same problem also emerges in some cases of backward reasoning in the service of acting. For example, Cassie might have a plan for crossing the street. Part of the plan may include a conditional act: *"If the walk-light is on, then cross the street"*. Note that this is a conditional act, one that involves two things: (i) trying to deduce whether the walk-light is on, and (ii) crossing the street or doing nothing, depending on the result of the deduction process. Evidently, to formalize the act, we have the same difficulty that we have with (4). Using the assertion time proposal, one might represent the act as shown in (11), where the act following

Forward(s_1)

1. Perform usual forward chaining on s_1.
2. If s_1 = Holds(s_2, *NOW) then Forward(s_2).

Backward(s_1)

1. If s_1 is eternal then perform usual backward chaining on s_1.
2. Else Backward(Holds(s_1, *NOW)).

Figure 1: Modified forward and backward chaining procedures.

"THEN" is to be performed if the state following "ActIf" holds.

(11) $\forall t$ ActIf Asserted(Holds(On(walk-light), t), t) THEN Cross(street)

However, attempting to deduce Asserted(Holds(On(walk-light), t), t) will succeed even if t matches some past time, t_0, at which it was asserted that the walk-light is on. Hence, introducing the assertion time only solves the problem with forward but not backward reasoning.

4.4 A Solution

What is the problem? At its core, the problem is that we need to make some assertions about future acts that refer to unmentionable future *NOWs. Those *NOWs would only be known at the time of acting. Even their being *future* is not something absolute that we know about them; they are only future with respect to the assertion time. We somehow need to introduce *NOW only when it is known— at the time of acting. Our proposal is to eliminate reference to time in rules like (4) (or acts like (11) for that matter) and let the inference and acting system introduce *NOW when it is using these rules. Thus, instead of (4), we shall use (3) for both cases where *cond* is eternal or temporary.

(3) When *cond* DO *act*

Figure 1 outlines modified forward and backward chaining procedures. The input to these procedures is a state (eternal or temporary) s_1. Note that *NOW is inserted, by the procedures themselves at the time of reasoning. This guarantees picking up the appropriate *NOW. Going back to the fire-alarm example, consider the timed sequence of assertions in Table 2

Assertion		**Assertion Time**
(12)	When Sounds(alarm) DO Leave(building)	t_1
(13)	Holds(Sounds(alarm), t_0)	t_2
(14)	Holds(Sounds(alarm), t_3)	t_3

Table 2: Fire-alarm scenario for the modified chaining procedures.

Forward(Holds(Sounds(alarm), t_0))
1. Holds(Sounds(alarm), t_0) doesn't match Sounds(alarm)
2. $t_0 \neq t_2$

Figure 2: Forward inference on (13) at t_2 does not lead to acting.

(which is a variant of Table 1). As illustrated in Figure 2, asserting (13) at time t_2 does not cause Cassie to leave the building. First, note that (13) does not match (12) and hence the act of leaving the building will not be activated by step 1 of the **Forward** procedure. Second, since t_0 is not identical to *NOW (t_2), the recursive call to **Forward** in step 2 will not be performed. Thus, Cassie will, correctly, not leave the building just because she is informed that the fire-alarm sounded in the past. On the other hand, as illustrated in Figure 3, at t_3 the fire-alarm actually sounds. Still, (14) does not match (12). However, since t_3 is itself *NOW, step 2 results in **Forward** being called with "Sounds(alarm)" (which matches s_2). By step 1, of the recursive call, this will match (12) resulting in Cassie, correctly, leaving the building. Similarly, we may replace (11) by (15):

(15) ActIf On(walk-light)THEN Cross(street).

If Cassie is told to perform this conditional act at t_1, the procedure **Backward** would be called with "On(walk-light)" as an argument. Since this is a temporary state, backward chaining will be performed on Holds(On(walk-light), *NOW), thus querying the

Forward(Holds(Sounds(alarm), t_3))
1. Holds(Sounds(alarm), t_3) doesn't match Sounds(alarm)
2. $t_3 = t_3$
   Forward(Sounds(alarm))
   1. Sounds(alarm) matches Sounds(alarm) so Leave(building)

Figure 3: Forward inference on (14) at t_3 does lead to acting.

knowledge base about whether the walk-light is on at t_1, the time we are interested in.

5 The Problem of the Fleeting Now

Imagine the following situation. At t_1, we tell Cassie to perform the act represented in (15). The modified backward chaining procedure initiates a deduction process for Holds(On(walk-light), t_1). Using acting in service of reasoning, Cassie decides to look toward the walk-light in order to check if it is on. In order to look, Cassie moves her head (or cameras, if you will). Since time moves whenever Cassie acts, NOW moves to a new time, t_2. Cassie notices that the walk-light is indeed on. This sensory information is represented in the form of an assertion "Holds(On(walk-light), t_3)", where *NOW $\sqsubseteq t_3$. By the homogeneity of states, this means that "Holds(On(walk-light), *NOW)". However, *NOW is t_2, not t_1, the time that we were originally interested in. Thus, the deduction fails and Cassie does not cross the street even though the walk-light is actually on!

It should be noted that this problem is not a result of the modified inference procedures. The general problem is that the very process of reasoning (which in this case involves acting) may result in changing the state in which we are interested. We are interested in the state of the world at a specific time. Sensory acts are essentially durative and whatever observations we make would be, strictly speaking, about a different time.[14] It is in the "strictly speaking" part of this last sentence that we believe the solution to the problem lies. The following sentences could be normally uttered by a speaker of English.

(16) I am *now* sitting in my office.

(17) I *now* exercise everyday.

(18) I am *now* working on my PhD.

The word "now" in the above sentences means basically the same thing: the current time. However, there are subtle differences among the three occurrences of the word. In particular, the "now" in each case has a different *size*. The "now" in (18) is larger than that in (17) which is larger than that in (16). The same observation has been made by (Allen and Kautz, 1988, p. 253). Evidently, we conceive of "now's" at different levels of granularity. The problem outlined above really lies in our treatment of *NOW at a level of granularity that is too fine for the task Cassie is executing.

[14]Interestingly, this is the gist of the uncertainty principle in quantum physics.

We are interested in a level relative to which t_1 and t_2 would be *indistinguishable* (à la (Hobbs, 1985)).

Granularity in general, and temporal granularity in particular, has been discussed by many authors (Hobbs, 1985; Habel, 1994; Euzanet, 1995; Pianesi and Varzi, 1996; Mani, 1998). However, these approaches, though quite detailed in some cases, only provide insights into the issue; they do not represent directly implementable computational solutions. What we are going to do here is *sketch* an approach, one that we intend to further pursue and refine in future work. The main idea is to give up thinking of values of NOW as single terms. Instead, each *NOW may have a rich structure of subintervals which are themselves *NOWs at finer levels of granularity. Our approach is to think of the meta-logical variable NOW not as taking the values of time-denoting terms, but rather of totally-ordered sets of time-denoting terms. More precisely, $\langle$NOW, $\sqsubseteq\rangle$ is a totally ordered poset.

We can think of this poset as a stack, such that the greatest and least elements are the bottom and top of the stack elements, respectively. The symbol "*NOW" is now to be interpreted as referring to the top of the stack of NOWs. Moving from one level of granularity to another corresponds to pushing or popping the stack. In particular, to move to a finer granularity, a new term is pushed onto the stack, and thus becomes *NOW. On the other hand, to move to a coarser granularity, the stack is popped. At any level, the movement of time is represented by replacing the top of the stack with a new term. Symbolically, we represent these three operations, illustrated in Fig. 4, as: $\downarrow$NOW, $\uparrow$NOW, and $\updownarrow$NOW respectively (the last one is motivated by realizing that replacement is a pop followed by a push).

Using this mechanism, the problem outlined above may be solved as follows.

1. Cassie wonders whether "Holds(On(walk-light), t_1)".
2. Cassie decides to look towards the walk-light.
3. $\downarrow$NOW (*NOW $= t_2$).
4. Cassie looks toward the walk-light.
5. $\updownarrow$NOW (*NOW $= t_3$).
6. Cassie senses that the walk-light is on. That is, an assertion "Holds(On(walk-light), t_4)" is made, with $t_3 \sqsubseteq t_4$.
7. $\uparrow$NOW (*NOW $= t_1$).

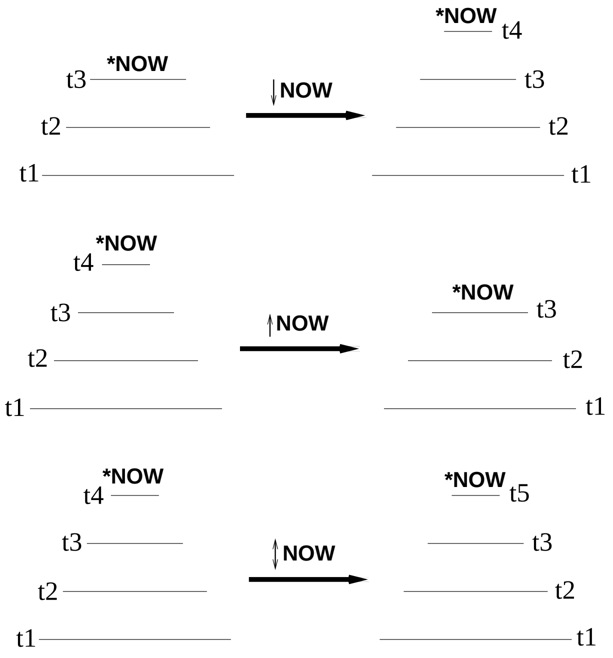

Figure 4: Operations on the stack of NOWs.

8. Assert "$t_1 \sqsubseteq t_4$".

9. Cassie realizes that "Holds(On(walk-light), t_1)".

The above sketches a solution to the problem, it obviously does not present a complete theory. To arrive at such a theory, various questions need to be solved. First, when to push and pop the stack is still not exactly clear. Here, we decide to push into a finer granularity when acting is performed in service of reasoning (step 3). In general, one might propose to perform a push any time the achievement of a goal requires achieving sub-goals. Popping is the complementary operation, it could be performed after each sub-goal has been achieved (step 7). The most problematic step in the above solution is step 8. The motivation behind it is simple and, we believe, reasonable. When Cassie notices that the walk-light is on at t_3, it is reasonable for her to assume that it was on over a period starting before and extending over the *NOW *within* which she is checking the walk-light, namely t_1. Of course, this presupposes certain intuitions about the relative lengths of the period of the walk-light being on (t_4) and of that over which Cassie acts (t_1). In a sense, this is a variant of the frame problem (McCarthy and Hayes, 1969); given that a state s holds over interval t_1, does it also hold over a super-interval, t_2, of t_1?[15]

[15]In the traditional frame problem, t_2 is a successor, not a super-interval, of t_1.

Our main objective here is not to provide a complete solution to the problem. Rather, we want to point the problem out, and propose some ideas about how it may be solved. Future research will consider how the proposal outlined above may be extended and refined into a concrete theory of temporal granularity that could be applied to reasoning, acting, and natural language interaction.

6 Conclusions

A reasoning, acting, natural language competent system imposes unique constraints on the knowledge representation formalism and reasoning procedures it employs. Our commitment to using a common representational formalism with such a multi-faceted system uncovers problems that are generally not encountered with other less-constrained theories. For example, a memoryless agent may use the STRIPS model of change, in which case representing temporal progression, and having to deal with the problems it raises, would not be required. A logical language that is not linguistically-motivated need not represent a notion of the present that reflects the unique semantic behavior of the natural language "now"— an issue that underlies the two problems discussed.

The problem of "the unmentionable now" results from the inability to refer to future values of the variable NOW. Since *NOW can only refer to the time of the assertion (mirroring the behavior of the English "now"), one cannot use it in the object language to refer to the future. Such reference to future now's is important for specifying conditional acts and acting rules. Our solution is to eliminate any reference to those times in the object language, but to modify the forward and backward chaining procedures so that they insert the appropriate values of NOW at the time of performing a conditional act or using an acting rule. The problem of "the fleeting now" emerges when, in the course of reasoning about (the value of) NOW, the reasoning process itself results in NOW changing. The solution that we sketched in Section 5 is based on realizing that, at any point, the value of NOW is not a single term, but rather a stack of terms. Each term in the stack corresponds to the agent's notion of the current time at a certain level of granularity, with granularity growing coarser towards the bottom of the stack. Temporal progression and granularity shifts are modeled by various stack operations.

An agent that reasons about its actions, while acting, and has a personal sense of time modeled by the relentlessly moving NOW is different in a non-trivial sense from other agents. The problem of "the unmentionable now" shows that, for such an agent, time is not just an external domain phenomenon that it, accidentally, needs to reason about. Rather, time is an internal factor that the very deep inferential processes need to take into account. The problem of "the fleeting now" moves the issue of temporal granularity from the rather abstract realm of reasoning and language understanding down to the real world of embodied acting.

7 Acknowledgments

The authors thank the members of the SNePS Research Group of the University at Buffalo for their support and comments on the work reported in this paper. This work was supported in part by ONR under contract N00014-98-C-0062.

References

Allen, J. (1983). Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(11).

Allen, J. F. and Kautz, H. A. (1988). A model of naive temporal reasoning. In Hobbs, J. and Moore, R., editors, *Formal Theories of the Commonsense World*, pages 261–268. Ablex Publishing Corporation, Norwood, NJ.

Almeida, M. (1995). Time in narratives. In Duchan, J., Bruder, G., and Hewitt, L., editors, *Deixis in Narrative: A Cognitive Science Approach*, pages 159–189. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.

Almeida, M. and Shapiro, S. (1983). Reasoning about the temporal structure of narrative texts. In *Proceedings of the Fifth Annual Meeting of the Cognitive Science Society*, pages 5–9, Rochester, NY.

Bach, E. (1986). The algebra of events. *Linguistics and Philosophy*, 9:5–16.

Braun, D. (1995). What is character? *Journal of Philosophical Logic*, 24:227–240.

Cresswell, M. J. (1990). *Entities and Indices*. Kluwer Academic Publishers, Dordrecht.

Euzanet, J. (1995). An algebraic approach to granularity in qualitative time and space representation. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 894–900, San Mateo, CA. Morgan Kaufmann.

Fikes, R. E. and Nilsson, N. J. (1971). STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2(3/4):189–208.

Galton, A. (1984). *The Logic of Aspect*. Clarendon Press, Oxford.

Habel, C. (1994). Discreteness, finiteness, and the structure of topological spaces. In Eschenbach, C., Habel, C., and Smith, B., editors, *Topological Foundations of Cognitive Science*, pages 81–90. Papers from the Workshop at the FISI-CS, Buffalo, NY.

Hexmoor, H. (1995). *Representing and Learning Routine Activities*. PhD thesis, Department of Computer Science, State University of New York at Buffalo, Buffalo, NY.

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, Menlo Park, CA / Cambridge, MA.

Hobbs, J. (1985). Granularity. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pages 432–435, Los Altos, CA. Morgan Kaufmann.

Iwańska, Ł. M. and Shapiro, S. C., editors (2000). *Natural Language Processing and Knowledge Representation: Language for Knowledge and Knowledge for Language*. AAAI Press/The MIT Press, Menlo Park, CA.

Kamp, H. (1971). Formal properties of 'now'. *Theoria*, 37:227–273.

Kaplan, D. (1979). On the logic of demonstratives. *Journal of Philosophical Logic*, 8:81–98.

Kumar, D. and Shapiro, S. C. (1994). Acting in service of inference (and vice versa). In Dankel, II, D. D., editor, *Proceedings of the Seventh Florida Artificial Intelligence Research Symposium*, pages 207–211, St. Petersburg, FL. The Florida AI Research Society.

Lespérance, Y. and Levesque, H. (1995). Indexical knowledge and robot action– a logical account. *Artificial Intelligence*, 73(1–2):69–115.

Mani, I. (1998). A theory of granularity and its application to problems of polysemy and underspecification of meaning. In Cohn, A., Schubert, L., and Shapiro, S., editors, *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR '98)*, pages 245–255, San Francisco, CA.

Martins, J. and Shapiro, S. C. (1988). A model for belief revision. *Artificial Intelligence*, 35(1):25–79.

McCarthy, J. and Hayes, P. (1969). Some philosophical problems from the standpoint of artificial intelligence. In Meltzer, D. and Michie, D., editors, *Machine Intelligence*, volume 4, pages 463–502. Edinburgh University Press, Edinburgh, Scotland.

McDermott, D. (1982). A temporal logic for reasoning about processes and plans. *Cognitive Science*, 6:101–155.

Mourelatos, A. P. D. (1978). Events, processes, and states. *Linguistics and Philosophy*, 2:415–434.

Pianesi, F. and Varzi, A. (1996). Refining temporal reference in event structures. *Notre Dame Journal of Formal Logic*, 37(1):71–83.

Prior, A. N. (1968). "Now". *Noûs*, 2(2):101–119.

Quine, W. V. O. (1960). *Word and Object*. The MIT Press, Cambridge, MA.

Reichenbach, H. (1947). *Elements of Symbolic Logic*. Macmillan Co., New York, NY.

Shanahan, M. (1995). A circumscriptive calculus of events. *Artificial Intelligence*, 77:249–284.

Shapiro, S. (1989). The CASSIE projects: An approach to natural language competence. In *Proceedings of the 4th Portugese Conference on Artificial Intelligence*, pages 362–380, Lisbon, Portugal. Springer-Verlag.

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

Shapiro, S. (1998). Embodied Cassie. In *Cognitive Robotics: Papers from the 1998 AAAI Fall Symposium*, pages 136–143, Menlo Park, CA. AAAI Press. Technical report FS-98-02.

Shapiro, S. and Rapaport, W. (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. and Rapaport, W. (1992). The SNePS family. *Computers and mathematics with applications*, 23(2–5):243–275. Reprinted in F. Lehman, ed. *Semantic Networks in Artificial Intelligence*, pages 243–275. Pergamon Press, Oxford, 1992.

Shapiro, S. and Rapaport, W. (1995). An introduction to a computational reader of narratives. In *Deixis in Narrative: A Cognitive Science Perspective*, pages 79–105. Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.

Shapiro, S. 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.

Shoham, Y. (1985). Ten requirements for a theory of change. *New Generation Computing*, 3:467–477.

Shoham, Y. (1987). Temporal logics in AI: Semantical and ontological considerations. *Artificial Intelligence*, 33:89–104.

van Benthem, J. (1983). *The Logic of Time*. D. Reidel Publisihng Company, Dordrecht, Holland.

Vendler, Z. (1957). Verbs and times. *The Philosophical Review*, 66(2):143–160.