The Cognitive Clock: A Formal Investigation of the Epistemology of Time

Haythem O. Ismail and Stuart C. Shapiro
Department of Computer Science and Engineering
and Center for Cognitive Science
University at Buffalo, The State University of New York
226 Bell Hall
Buffalo, NY 14260-2000
E-mail: {hismail|shapiro}@cse.buffalo.edu
July 26, 2001

Abstract

We present a formal model of time and time perception. The model is used to endow a cognitive agent with a personal sense of time. A personal sense of time is based on a representation of the real-time present that allows the agent to distinguish it from the past and the future, and is continuously changing to reflect the progression of time. This is important for reasoning about, and acting in, dynamic domains and for interacting with other agents (humans, for example). An agent that interleaves reasoning and acting while maintaining a sense of the present faces what has been dubbed *the problem of the fleeting now*. A solution to the problem involves, not only endowing the agent with a sense of temporal progression, but also with a feel for how much time has passed. In this paper, we construct a theory of time that allows just that.

> *Clocks never have a moment's rest; no sooner have they achieved the desired relationship to the current time than time slips out from under their fingers* ... (Smith, 1988, p. 9)

1 Introduction

Cognitive agents reason and act in the real world. In the real world, everything takes time, and for an agent to behave appropriately, it needs to reason about time. Most acting systems in artificial intelligence (AI) incorporate time and reasoning about it either implicitly or explicitly. However, in those systems, time happens to be yet another domain phenomenon that a cognitive agent occasionally needs to reason about. In particular, as represented in those systems, time is totally decoupled from the agent's experience; the agent does not have any personal sense of the real time. Two questions now arise. What is a *personal* sense of time, and why is it important for cognitive agents? Basically, a personal sense of time is based on a representation of the real-time present that (i) allows the agent to distinguish it from the past and the future, and (ii) is continuously changing to reflect the progression of time. This is important for various reasons. First, an agent reasoning about a dynamic domain, one that changes as it is reasoning, needs to distinguish present facts from past ones. In addition, if the agent is itself acting in such an environment, distinguishing the present is crucial for it to identify those facts that are pertinent to its actions. Second, cognitive agents, in addition to acting, may also be *interacting* with other agents (humans, for example). Such interaction, whether in natural language or not, would only be effective and intelligible if the agent has a clear distinction between what is now the case, versus what used to be the case. There are at least two reasons why this distinction is important. First, if the agent expresses its beliefs in natural language, then a distinction between the present, the past, and the future is required for generating sentences with the correct tense. Second, and more important, present facts have a distinguished status for an acting agent. Suppose an agent believes that whenever the fire-alarm sounds it should leave the building. If we tell the agent that the alarm sounded yesterday, it merely needs to remember this fact, and maybe derive some inferences from it. However, if we tell it that the fire-alarm is *now* sounding, it also needs to act on this fact and leave the building.[1] Thus, distinguishing the present is important for correct behavior, not just for intelligible interaction and sound reasoning.

Note that, since facts that are at some time present would eventually be past, the representation of time should be such that the mere flow of time does not require extensive revision of the agent's knowledge. More precisely, expressions in a knowledge base should not be manipulated just because time has passed. This is reasonable both for practical computational efficiency purposes, and since an experience as intimate and frequent as the movement of the present should be as transparent to the knowledge base as possible.

What is needed is a theory of subjective time, one that involves a representation of a moving "now". This should be done in such way that the movement of "now" leaves the knowledge base intact, without the need for any involved *mental surgery* in the agent's knowledge base. As it turns out, an agent that interleaves reasoning and acting while maintaining a sense of the present, represented by the relentlessly moving "now", faces what has been dubbed *the problem of the fleeting now* (Ismail and Shapiro, 2000b). The gist of the problem is that reasoning takes time. It emerges when the agent is reasoning about "now", when the very process

[1]For the problems involved in making the agent act on present facts, see the discussion of the problem of *the unmentionable now* in (Ismail and Shapiro, 2000b).

of reasoning results in "now" moving, and thereby fleeting from the agent's *mental grip*. As shall be argued, a solution to the problem involves, not only endowing the agent with a sense of temporal progression, but also with a *feel* for how much time has passed. In this paper, we construct a theory of time that allows just that.

In Section 2 we review the relevant literature on subjective time. In Sections 3 and 4 we present the problem of the fleeting now, an analysis of where we believe the problem exactly lies, and an informal sketch of our solution to it. Section 5 reviews the two main psychological theories explaining how humans judge durations of time. Section 6 introduces the framework and assumptions underlying our theory of agents. The theory of subjective time and temporal progression is presented in detail in Section 7. Section 8 examines in detail how belief update interacts with the progression of time. In Section 9, we discuss how our agent can have a sense of how much time has passed, which leads to a solution to the problem of the fleeting now in Section 10. An application of the theory to research on the frame problem is outlined in Section 11. In Section 12, a number of theorems are proved to illustrate that the proposed theory allows the agent to have reasonable beliefs about the state of the environment, based on its perception of the environment, its knowledge of its own body, and its subjective sense of time. Finally, in Section 13, we present our conclusions and the implications of our investigation for temporal reasoning.

2 Subjective Time

There are two main views that a theory of time may adopt. These may be called the *subjective* and the *objective* views of time. Historically, the distinction was precisely formalized by (McTaggart, 1908) in his *A series* and *B series*. Basically, the A series is the series of positions in time "running from the far past to the near past to the present, and then from the present to the near future and the far future" (McTaggart, 1908, p. 458). The B series is the "series of positions which runs from earlier to later" (McTaggart, 1908, p. 458). From a cognitive perspective, the A series involves a deictic representation of time, relative to a cognitive being, hence the subjectivity. On that view, the agent is itself located *in* time (metaphorically, the agent is *on* the time line), and maintains a sense of the present time. The B series is a totally objective view of time, independent of the experience of any conscious mind. On that view, the agent exists *out of time* and merely reasons about a static configuration of events. (Lakoff and Johnson, 1999, Ch. 10) present evidence for two versions of the A series: one with a moving time and a stationary *observer*, whose location represents the present, and one with a stationary time and a moving observer. For example, compare *"The deadline is approaching"* to *"We are approaching the deadline"*.

Most AI research on time takes the objective view. On that view, time is yet another domain phenomenon that a reasoning system needs to reason about. Granted, the very introduction of time and change has important impacts on many syntactic, semantic, and ontological commitments. For example, the introduction of time raises questions about the reification of propositions (Shoham, 1987; Bacchus et al., 1991); the nature of temporal individuals (whether times are basic individuals, and if so, whether they are best represented as points, intervals, or both (McDermott, 1982; Allen, 1983; Allen, 1984; Shoham, 1985; Galton, 1990)); and the structure of time (whether it is dense, discrete, branching, etc. (McDermott, 1982; Pinto, 1994; Habel, 1994, for instance)). These issues have effects on the overall structure of a logic (see (Galton, 1995) for a more elaborate discussion). As outlined in Section 1, our research is concerned with endowing an agent with a sense of the present that would put it on the time line, providing it with a subjective view of time. Representing the present requires a logical account of the concept of "now".

One of the most common methods of accounting for the semantics of indexicals (of which "now" is a special case) is that outlined by David Kaplan (Kaplan, 1979). In an intensional Montagovian setting, Kaplan argues for a two-step interpretation function. The *character* of an expression is a function from contexts (of utterances) to *contents*, which are the traditional intensions (i.e., functions from possible worlds to appropriate domain entities, truth values, or set-theoretical structures over the domain). A competent speaker of English would recognize that "now" is used (loosely speaking) to refer to the current time. This knowledge is actually knowledge of the character of "now". The content of "now", on the other hand, is a particular time that does not by itself convey any concept of the present. The semantics of "now" is even more involved. As pointed out by many authors (Prior, 1968; Kamp, 1971; Cresswell, 1990), "now" always refers to the time of utterance even when embedded within a nest of tense operators.[2] This proves to be a very technically involved issue in the semantics of "now". These features of "now" pose unique problems for any logical account of the present. The main concern is what a logic should represent: the character or the content of "now". Of course, logics of objective time do not even need to consider such an issue; only contents are represented, since characters of demonstrative-free expressions are constant functions. On the other hand, systems that need to have a notion of the present have (at least) four choices.

First, either implicitly or explicitly, the present is directly represented in the logic. This is typical of classical Priorian tense logics (Prior, 1967; Rescher and Garson, 1968; Kamp, 1971). Classical tense logic is essentially a temporally-interpreted modal logic with two, rather than one, modal operators, together with their duals. If p is a proposition, "Pp" means that "It has been the case that p" and "Fp" means that "It will be the case that p". By itself, "p" refers to the *current* truth of p. Thus, syntactically, the present is distinguished by having the proposition outside

[2] An exception would be the "now-point" of a narrative (Almeida, 1995; ter Meulen, 1997).

the scope of any tense operators. Semantically, expressions (which may be embedded within tense operators) are interpreted with respect to a particular temporal index representing the present.[3] Although these logics are capable of distinguishing among the past, the present, and the future, they do not represent the passage of time. In a sense, they endorse an implicit objective view of time where the present is a static pivotal point from which one views the temporal configuration of events. An agent reasoning in the world needs to maintain a notion of the present that continuously changes as events occur. Evidently, traditional tense logic, by itself, does not suffice for such an agent.

The second approach, usually adopted in reasoning about actions and plans, is to represent the present using an indexical **now** term. The use of indexical terms, in general, was studied in depth by (Lespérance and Levesque, 1995) with special attention to the case of **now** in (Lespérance and Levesque, 1994).[4] The indexicality of such a term stems from its having a context-dependent interpretation, much in the same spirit of Kaplan's semantics discussed above. However, unlike the English "now", whose content depends on the context of utterance (or assertion), the semantics of the indexical **now** depends on the *evaluation* context. In the context of acting and planning, it is the time of executing a particular instance of a plan that includes occurrences of **now** in its specification. Such an approach facilitates certain kinds of reasoning that seem to require an explicit representation of indexical time (Lespérance and Levesque, 1994). Along the lines of (Lespérance and Levesque, 1994) (and using the same syntax), the following is a possible representation of a plan to get to other side of the street (probably for a rather despondent agent):

$$\mathbf{if}(\mathbf{At}(\mathbf{now},\ \textsc{WalkLightOn}), \textsc{Cross}, \mathbf{noOp})$$

This roughly says that, if, at that the time of performing the action, the walk-light is on, then cross the street; otherwise do nothing. What should be noted is that **now** in the above form does not refer to the time of introducing the form into the knowledge base, or to any other *fixed* time for that matter. It is, in a sense, a place-holder for any time at which the plan is performed.[5]

[3]See (Kamp, 1971; van Benthem, 1983; Cresswell, 1990; van Benthem, 1995) for other approaches that use multiple temporal indices.

[4]Other authors have also used the same or a similar approach (Schmiedel, 1990; Artale and Franconi, 1998; Dix et al., 2001).

[5]On a first pass, it seems to us that the use of the indexical **now** is only *required* in certain descriptions of such future actions. In particular, it seems that *imperative*, rather than *declarative*, representation of plans are the ones that are inherently indexical (see (Huntley, 1984) for a linguistic perspective.) We have managed to rewrite all formulas in (Lespérance and Levesque, 1994) without mentioning **now**. In fact, the discussion of (Perry, 1979), who is a major proponent of indexical beliefs (and is cited by (Lespérance and Levesque, 1994; Lespérance and Levesque, 1995)), is mainly about beliefs required for explaining behavior (see (Millikan, 1990) for a critical examination of Perry's position). More examination of when exactly indexicals are required is needed.

However, as represented, and used, in (Lespérance and Levesque, 1994; Lespérance and Levesque, 1995), such a concept of the present is certainly not the same one represented by the English "now". First, because its interpretation is not dependent on the time of assertion, there are uses of **now** that do not correspond to the *real* present. Second, temporal progression cannot be modeled by such an abstract characterization of the present. For example, it is intuitive to express the English *It is now raining* as, for example, **At**(**now**, RAINING). Since **now** is not interpreted at the time of the assertion, the expression **At**(**now**, RAINING) would essentially mean that it is always raining, since **now** always refers to the mere notion of the current time.

As for the first point, (Lespérance and Levesque, 1995, p.82) explicitly state that their **now** is not intended to represent the English "now". As for the second point, they briefly discuss a solution which we will now consider in some detail. The obvious approach to modeling the passage of time within the theory of (Lespérance and Levesque, 1995) would be to appropriately edit the knowledge base every time "now" changes in order to preserve the truth of its sentences. Thus, **At**(**now**, RAINING) should be replaced by something more appropriate once "now" changes. One problem, of course, is that such updates might be computationally expensive. To get around the problem, (Lespérance and Levesque, 1995, p. 101) suggest that "if all occurrences of 'now' are replaced by a new constant and the fact that this new constant is equal to 'now' is added, then only this single assertion need be updated as time passes." This indeed eliminates the problem of expensive belief update and provides a neat logical and computational account of "now". However, from a cognitive perspective, we find the very idea of erasing sentences from an agent's mind as time passes by far from natural. If such sentences represent beliefs that the agent once held, where do they go, and how come the agent would have no memory of them once time passes? Note that this cannot be explained away as a matter of forgetting, for forgetting is not that selective to always affect beliefs involving "now", nor is it vigorous enough to take effect with every tick of the clock. The only way to explain this mysterious disappearance of beliefs is by arguing that they exist at a lower level of consciousness with respect to other beliefs. If this were the case, why are such beliefs part of the logical theory (which we take to be representing conscious beliefs of the agent)? Thus, we do not see a way of reconciling the notion of "now" as a term in the agent's language of conscious thought with a cognitively plausible account of the passage of time. Note that we are not trying to refute Perry's (or, for that matter, Lespérance and Levesque's) arguments for indexicality (Perry, 1979). The theory presented in this paper incorporates indexical *thought*, but not indexical *thoughts*. That is, the reasoning process is indeed sensitive to indexicality, but the object language does not include any indexical terms.

The third approach to represent "now" is to do it indirectly, by means of a *Now* predicate, where the expression $Now(i)$ means that the current time is represented by the term i. This is exactly the method adopted in *active logic*, originally known as *step logic* (Elgot-Drapkin and Perlis, 1990; Perlis et al., 1991). Active logic is perhaps the only system that takes the issue of reasoning *in* time seriously. Temporal individuals are represented by integers, with the usual numerical order implicitly representing chronological order. Thus, active logic generally adopts a discrete model of time points, though, in (Nirkhe et al., 1997), a point-based representation of intervals is used for reasoning about durative actions and events. Active logic is intended to model an agent's reasoning about its own reasoning processes in time. As such, it is not primarily interested in reasoning *about* time. Thus, one does not find any detailed analyses of different situation types.

In active logic, time moves with every inference step. This movement of time is represented both logically and meta-logically. Logically, this is achieved by a special inference rule that essentially replaces $Now(i)$ by $Now(i+1)$. Meta-logically, assertions are associated with the step, i, of inference at which they were asserted. By having special "inheritance" (inference) rules, assertions made at step i may be rederived at step $i+1$, thereby allowing persistence of information. The system can reason about the present, since, at any step i, $Now(i)$ is asserted. In addition, beliefs about the past could be carried on to the present by yet another special rule of inference allowing the system to derive $Know(\alpha, i)$ from α being derived at step i.[6] Indeed, active logic proves very powerful in a number of applications (Nirkhe et al., 1997; Gurney et al., 1997).

However, the *Now* predicate, though well-suited for the kind of reasoning problems addressed within active logic, does not exactly reflect the behavior of the English "now". This is apparent in the special rules that are tailored just to get the appropriate inferences concerning the current time. Most notably, *Now(i)* is a special case that the general inheritance rule cannot apply to, that is, *Now(i)* cannot be carried on to step $i+1$. In addition, active logic would have to be supplemented with more special rules to correctly deal with occurrences of *Now(i)* within opaque contexts. Another problem with active logic is its use of integers to represent time. Apparently, the use of integers facilitates the expression of some crucial rules of inference (also the *counting* of reasoning steps (Nirkhe et al., 1997)) that depend on having a well-defined notion of the *next moment of time*, represented by the integer successor operator. However, such a representation forces a certain degree of rigidity on the kind of knowledge that may be entered into the system. For example, there is no way to assert at step $i+m$ $(m > 1)$ that a certain event e_2 occurred between events e_1 and e_3 that happened at times i and $i+1$, respectively. In other words, once "now" moves, there is no way to go back and create arbitrary past temporal

[6] No distinction is made between belief and knowledge, so $Know(\alpha, i)$ should actually be read as $Believe(\alpha, i)$. This is intended to mean that, at step i, the system believed α. This is a piece of knowledge that should always be believed by the system, even if, at some later step, α is no longer believed. However, it is not clear how the system may believe that, at some past time, some state held.

locations. This is definitely a big drawback if the system is to be used in general interactions, where assertions need not be only about the present.

The final approach for representing the present is that of (Almeida, 1987; Almeida, 1995) and (Shapiro, 1998). In this approach, only contents are represented in the logical language; the notion of "now" exists only in the meta-theory. This is achieved by having a meta-logical variable, NOW, that assumes values from amongst the time-denoting terms in the logic. Use of the NOW variable is restricted to the natural-language interface and the reasoning and acting system; it is always replaced by its value (the content) when interacting with the knowledge base. Since the logic is a standard demonstrative-free logic, this approach overcomes all the difficulties outlined above, and is the one we have adopted. On the other hand, it poses certain problems in cases where the general notion of the current time needs to be represented in the logic (Ismail and Shapiro, 2000b). In addition, as pointed out in Section 1, any system that models temporal progression, and in which time passes as reasoning proceeds, faces the problem of the fleeting now.

3 The Vagueness of "now"

> Salesman: *Good morning Madame, may I interest you in a fine . . . ?*
> Woman (annoyed): *Not now.*
> Salesman: *. . . Now?*

What is funny about the above joke is probably the realization that the salesman's reply, though obviously silly, is in some sense valid. The woman's *"not now"* is primarily a show of disinterest but may also be taken to include an implicit invitation for the salesman to approach her later, at some different "now". The salesman's response is silly because our intuition is that it came at the same "now". Nevertheless, there is also a sense in which the response is appropriate, for, strictly speaking, whenever time passes there is a new "now".

What this all means is that we maintain different notions of "now" at the same time. Specifically, there is a set of time intervals, linearly-ordered by the subinterval relation, representing the concept of "now" at different levels of granularity (Ismail and Shapiro, 2000b). What an occurrence of "now" refers to is some element of this set; which one in particular is determined by contextual and pragmatic factors. In the above joke, the salesman conveniently interprets "now" at a level of granularity that is finer than the one picked by the woman. He does that by, probably, choosing to ignore some salient pragmatic factors.

A similar problem raised by (Quine, 1960) may help drive the point home. Quine delves into a long discussion of *radical translation*: "translation of a language of a hitherto untouched people" (Quine, 1960, p. 28). Quine discusses a certain difficulty that a linguist encounters when an accompanying native informant points in the direction of a rabbit and says: "Gavagai". The difficulty is that the mere act of pointing toward a rabbit does not necessarily mean that "Gavagai" is the translation of "Rabbit". For as far as the linguist can tell, the informant may be pointing, not to the rabbit as a whole, but to a part thereof (Quine, 1960, pp. 51–54). The pointing action is ambiguous as to which of a set of successively larger regions of space it is supposed to cover. Similarly, an utterance of "now" *points* to the present time, and is vague since there is always a nest of intervals containing the utterance event.

This vagueness of "now" has seldom been noted in the literature. Like us, (Allen, 1983, pp. 840–841) argues for a hierarchy of intervals to represent the present. However, his argument is primarily motivated by the desire to enhance the computational efficiency of a temporal reasoning system, rather than by genuine reasoning problems. It might seem that even though "now" is vague, this vagueness does not pose serious problems to language understanding. One reply is to witness the silliness of the salesman's response above, which stems from an inappropriate resolution of the vagueness. However, we are not primarily concerned with linguistic problems here. The main problem that we address is one that faces a reasoning and acting agent that maintains a sense of the present time. This problem, dubbed *the problem of the fleeting now* (Ismail and Shapiro, 2000b), is a direct result of ignoring the vagueness of "now".

4 The Fleeting Now

The time is t_1. Consider an agent that needs to know whether some state, s, holds *now.* To achieve its goal, the agent performs some sequence of actions, σ. Performing σ results in the agent observing that s indeed holds *now.* Naturally, σ takes time, and the observation is made at some time, t_2, that is later than t_1. Thus, strictly speaking, the agent observes that "s holds at t_2", whereas its original concern is whether "s holds at t_1". The problem, of course, is that, for the temporally-fanatic agent, the question is about one *now* (henceforth now_1) while the answer is about another (now_2). Just as in the case of the salesman, the agent picks the wrong level of granularity at which to interpret now_1. What is needed is a level relative to which now_1 and now_2 are *indistinguishable* (Hobbs, 1985). That is, now_1 should be interpreted as a coarser interval, t_3, which includes both t_1 and t_2. Thus, now_1 would be the same as now_2.

Figure 1 depicts the situation, where t_σ is the time interval over which σ is performed. As the figure shows, the only thing that we definitely know about t_3 is that it is a super-interval of t_1, i.e., it is a possible interpretation of now_1. For t_3 to also be a super-interval of t_2, it must survive t_σ. Since the nest of "now"s is

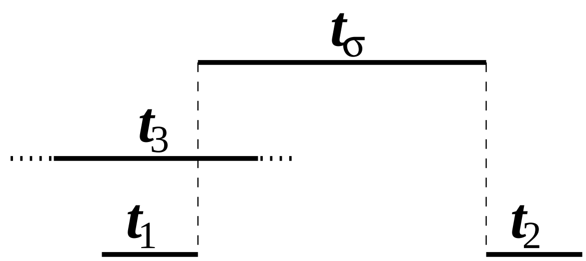

Figure 1: The temporal structure of the fleeting now problem.

potentially infinite, one can always find a *now* that is large enough to persist beyond t_σ. However, there is another restriction on the possible candidates for now_1. This is where knowledge about the state s comes into play. People have intuitions about the typical lengths of intervals over which various states persist. For example, walk-lights are on for about 15 to 30 seconds, meals are typically eaten in 15 to 30 minutes, conferences are held for several days, vacations last for a few weeks, and so on.[7] However now_1 is interpreted, it should be interpreted at a level of granularity that is neither too fine nor too coarse for s. In metrical terms, t_3 should be neither too short, nor too long, for the typical duration of s. If it is too short, the fleeting-now problem emerges. On the other hand, if it is too long, then observing s at t_2 cannot be used (at least defeasibly) to conclude that it holds throughout t_3.

The following spatial analogy may help fix the idea. Think of the nested intervals that are candidates for the interpretation of now_1 as some sort of *cognitive ropes* with various lengths. These ropes are fixed from one end (the left end, assuming time flows from left to right) at t_1. They are also somehow elastic, so that they could be stretched, within a certain limit, beyond their minimum lengths. When asked whether s holds *now*, the agent mentally picks one of these ropes. In particular, it picks one whose length is comparable (maybe within a half order of magnitude (Hobbs, 2000)) to the typical duration of s. The agent then moves to the right, holding the loose end of the rope, until at some point, t_2, it sees that s holds. If the agent is still holding the rope, it may answer the question affirmatively. If, on the other hand, it has run out of rope along the way from t_1 to t_2, it would not be able to conclude whether s holds at now_1. Thus, for the answer to be "yes", the agent needs to be still holding the rope when it sees s. In addition, it should have picked the *right* rope in the first place.

For example, suppose that at 12:15 p.m. sharp an agent is asked: *"Is John having lunch now?"*. To answer the question, the agent walks to John's office, where it sees John busy munching something, with a sandwich in his hand. Assume that this

[7] See (Hobbs, 2000) for an interesting discussion of the mathematical properties of typical measures.

perception event takes place at 12:17. Since now_1 is interpreted at a level of temporal granularity appropriate for "having lunch", its approximate duration is certainly longer than the 2-minute span of t_σ (the time it took the agent to walk to John's office). Thus, the agent may go back to the questioner and answer affirmatively. On the other hand, suppose that John is not in his office, but unbeknown to the questioner, is at home. The agent reasons that if it drives to John's house, it will arrive there at a *now* that is different from now_1, even when interpreted at the coarse lunch-granularity level (i.e., it will run out of rope). In this case, the agent has no way of answering the question, at least as long as it insists on adopting the driving-to-John's plan.

The above intuitions may be easy to state. But how would one formalize them for cognitive robotics applications? How can we give an embodied agent, not only a sense of the passage of time, but also a *feel* of how much time has passed? The rest of the paper attempts to answer these questions.

5 Life without a Clock

How do clocks keep track of time? "A clock is basically an instrument for creating a periodic process and counting the periods" (Dowden, 1998).[8] Thus, the amount of time elapsed is a function of the number of periods counted. In particular, since clock-periods have equal lengths, the amount of time elapsed is simply the number of counted periods multiplied by some constant period length. But agents do not always have access to clocks. For some people, forgetting to wear their watch before leaving to work in the morning may become a really disturbing experience. Throughout the day, they would have a feeling of loss and disorientation because of their inability to precisely keep track of time. Yet, despite these feelings, people do not totally lose their sense of time. They can still behave appropriately, estimate how much time it took to type a report, how long a meeting has lasted, and whether it is time to go home. How is this possible without a clock? Are there other means by which people perceive the passage of time? Certainly; the psychology literature provides various models of time perception (see (Friedman, 1990) for a general discussion). These are categorized into two major categories: biological models that hypothesize the existence of some internal clock or what is referred to as the *pacemaker* (see (Pastor and Artieda, 1996)), and cognitive models that assume its absence (see (Levin and Zakay, 1989)).

Computationally, a pacemaker is simply a process that generates equally-spaced *ticks*. (Ladkin, 1986) discusses a system that uses the workstation clock for representing the progression of time, where *now* is interpreted as the time indicated

[8]Of course, this only applies to the familiar (analog or digital) clocks, not to hour-glasses, for example.

by the clock. Similar methods may also be adopted. It should be noted, however, that, in general, what is needed is some method to *measure* the amount of time elapsed; the exact relation between ticks of the pacemaker and successive values of *now* need not be a one-to-one correspondence. The ticks of the pacemaker are best thought of as providing, not absolute dates, but a feel of the temporal proportions carved out by different events.

However, pacemaker-theories do not explain most of our everyday experiences of time.[9] If there is no internal clock, how else do we get a feel of the duration of what happens? Basically, we do it by knowledge of what *else* happens. For example, (Poynter, 1989) claims that the number of sensory and mental events filling a time interval is one major factor that provides the feel for its duration. He further argues that "whether an event turns out to be a useful marker of time passage depends on the length of time which is to be remembered" (Poynter, 1989, p. 312). For example, if one is interested in a duration on the scale of hours, then knowledge of events happening on the same scale is more helpful than knowledge of those happening on the scales of seconds or years. What this suggests is that it is not just the mere *number* of events filling an interval that matters, but that intuitions about typical durations of events provide rough metrics for measuring time.

The theory of time developed in the following sections is flexible enough to accommodate either of the above-presented views. In Section 9, we shall illustrate how our model may be used in conjunction with the cognitive theory of duration judgment. For technical reasons, however, we shall present a fuller, more elaborate discussion of a pacemaker-based theory. Nevertheless, it should be noted that this is a tactical, rather than a strategic, decision.

6 SNePS, Cassie, and Meinongian Semantics

Before delving into technical details, a brief discussion of the framework within which the analysis is pursued is due. Our theory of agents is based on the GLAIR agent architecture (Hexmoor et al., 1993; Hexmoor, 1995; Hexmoor and Shapiro, 1997). GLAIR is a layered architecture consisting of three levels:

1. **The Knowledge Level (KL):** The level at which conscious reasoning takes place. Our KL is implemented by the SNePS knowledge representation and reasoning and acting system (Shapiro and Rapaport, 1987; Shapiro and Rapaport, 1992; Shapiro and the SNePS Implementation Group, 1999).
2. **The Perceptuo-Motor Level (PML):** The level at which routines for carrying out primitive (or basic) acts are located. This is also the location for other

[9] See (Friedman, 1990, ch. 2) for a lengthy discussion and references.

subconscious activities that allow for the agent's consciousness of its body and surroundings.

3. **The Sensori-Actuator Level (SAL):** The level controlling the operation of sensors and actuators (being either hardware or simulated).

We use "Cassie" as the name of our GLAIR/SNePS-based agent. The KL corresponds to Cassie's "mind", and the SNePS knowledge base is to be interpreted as the contents of Cassie's memory. It is important to stress that SNePS structures do not represent "the world" (although they can). That is, the representations are not based on theories of physics or even (objectivist) metaphysics. For example, ontological questions about the (un)reality of time (McTaggart, 1908) or the status of temporal individuals (Chisholm, 1990; Pianesi and Varzi, 1996, for instance) might be important insofar as they shed more light on the phenomena under investigation. Nevertheless, answers to those questions are not sought and should not be a factor to consider when making decisions regarding SNePS representations. SNePS structures represent the "mind" of a cognitive agent. That is, rather than representing the world, they represent *conceptualizations* of the world. In addition, not any conceptualization will do; after all, the theory of relativity (or any other theory of physics for that matter) is just one possible conceptualization. What we mean is the kind of naive conceptualization that forms the basis of our linguistic expressions; the understanding of the world that underlies the way we *talk* about it. In many cases such an understanding is in obvious contradiction with the laws of physics (Hayes, 1985; Talmy, 1988, for example). Evidently, we talk about *times*; they are part of our understanding of the world. Indeed, just arguing that, say, temporal individuals are not real is by itself an evidence for the reality of their mental representation. Therefore, our representations will include SNePS terms corresponding to temporal individuals.

Given this cognitive interpretation, what do SNePS terms denote? They simply denote mental entities, in particular, Meinongian objects of thought (Rapaport, 1978; Shapiro and Rapaport, 1987; Shapiro and Rapaport, 1991). Mental entities are intensional, that is, their "... identity conditions *do* depend on their manner of representation" (Shapiro and Rapaport, 1987, original emphasis). Accordingly, there is a one-to-one correspondence between the set of SNePS terms and the domain of interpretation: no two distinct terms denote the same mental entity and no two distinct mental entities are denoted by the same term.

There are various ways by which new knowledge may be added to Cassie's mind:

1. **Direct Assertion:** Using either natural language (English) or various interfaces to SNePS (Shapiro and the SNePS Implementation Group, 1999), one may communicate information to Cassie.

2. **Inference:** Using inference, Cassie's memory may be expanded by deriving new beliefs from old ones.

3. **Bodily-Feedback:** The PML may communicate information to the KL as a result of perception or proprioception.

This information or knowledge represents propositions that Cassie believes. Propositions are first-class entities in the SNePS ontology (Shapiro and Rapaport, 1987; Shapiro, 1993). Thus, SNePS-based languages do not include predicate symbols in the traditional sense—symbols denoting tuples of individuals. Rather, there are function symbols denoting functions whose range is the set of propositions. Similarly, logical connectives are interpreted as functions over the domain of propositions rather than that of truth-values (Shapiro, 1993).

Of particular interest to the current study are propositions about *current* states of Cassie and the environment. In what follows, we make the following reasonable assumption: If perception, proprioception, or direct assertion result in Cassie's acquiring a new belief, then forward inference is initiated. Why is this assumption reasonable? An acting agent should always be alert, paying attention to its immediate surroundings and the states thereof in order to direct its actions and decide what to do next. Forward inference provides one way of modeling this state of high-alert; not only does Cassie accept what she is told or directly perceives, but she must also reason about it to see if there is more to the new information than its explicit content. In particular, she should determine whether the new information allows her to infer something about the *current* state. This is particularly important in the case of perception, since perceived states are usually very specific in character that they may, themselves, not be very helpful. In the example of John's lunch from Section 4, we said: "the agent walks to John's office, where it sees John busy munching something, with a sandwich in his hand." What the agent sees is *not* "John having lunch"; rather, it perceives some other state from which, and given everything else that the agent believes, it *infers* that John is having lunch. It would be ridiculous if the agent does not make that inference immediately, particularly given that the state of John's having lunch is something that it is interested in. Ultimately, however, we believe that full-fledged forward inference may not be needed. For example, although it is reasonable for Cassie, having acquired a new belief, to infer everything that she can about current states, it seems that inferring things about, for instance, the far past is not as motivated. In addition, forward inference should also be restricted to those propositions that seem pertinent to what the agent is currently doing or interested in. These are all issues that certainly need to be considered. Nevertheless, we defer their investigation to the future and uphold the assumption that acquiring new information *always* initiates forward inference.

7 The Cognitive Structure of Time

7.1 The Theory and the Meta-Theory

It is important from the outset to make a clear distinction between two formal theories that we shall be developing below. Let us call these "the theory" and "the meta-theory". The theory is a logical theory representing the contents of Cassie's mind. It is made up of well-formed expressions of a SNePS-based formal language with Meinongian semantics. The meta-theory, as the name indicates, is a theory about the "theory". Whereas the theory represents the contents of Cassie's mind, the meta-theory does not. It represents interesting structures of those contents, deep relations among those structures and how they evolve over time. In a sense it is the model theory of the logic. Two points may illustrate the distinction.

1. Cassie may believe that time is linearly ordered. That is, the "theory" may include a well-formed expression (representing a proposition that Cassie believes) to the effect that any two times are either contemporaneous or one of them precedes the other. Nevertheless, the structure of time in Cassie's mind need not be linear. That is, meta-theoretically, not all pairs of time-denoting terms are related by the precedence relation. For instance, Cassie may have beliefs about two times, t_1 and t_2, without having any (need for a) belief about their order. This point is discussed in detail by (Kamp, 1979).

2. Cassie may hold a belief, represented by a proposition-denoting term in the theory, that time is dense. However, from a meta-theoretical stance, time cannot be dense, for this would imply an infinite number of terms in Cassie's finite mind. That is, although Cassie may believe that time is dense, the structure of time-denoting terms in Cassie's mind is not itself dense.

7.2 The Theory

The formal language that we shall be using to construct the "theory" is a sorted first-order language. All variables are universally-quantified unless otherwise noted.

Sorts. Using standard type-theoretical notation, the set of denotations of symbols of sort Σ is denoted $\mathcal{D}_\Sigma$. The language has four main sets of symbols (sorts).

1. $\mathcal{T}$, where $\mathcal{D}_\mathcal{T}$ is a set of time intervals. We use t, t', and t_n, where $n \in \mathbb{N}$ is a natural number, as meta-variables ranging over members of $\mathcal{T}$.

2. TEMP, where elements of $\mathcal{D}_{\text{TEMP}}$ are temporary states. These may be thought of as relational fluents of the situation calculus (McCarthy and Hayes, 1969)—states of affairs that may, or may not, hold at various times (see (Ismail and Shapiro, 2000b) for more details). Temporary states cover traditional stative situations (for example, the walk-light being green, or Cassie being empty-handed) in addition to dynamic processes (Cassie's crossing the street, her moving to John's office, her picking up a block, etc.). For arguments in favor of this broad interpretation of states, see (Galton, 1984; Galton, 1990). We use s, s', and s_n ($n \in \mathbb{N}$) as meta-variables ranging over members of TEMP.

3. $\mathcal{P}$, where $\mathcal{D}_{\mathcal{P}}$ is a set of propositions. A *finite* subset, β, of $\mathcal{D}_{\mathcal{P}}$ represents Cassie's belief space—the set of propositions that she believes. We use p, p', and p_n ($n \in \mathbb{N}$) as meta-variables ranging over members of $\mathcal{P}$.

4. $\mathcal{Q}$, where $\mathcal{D}_{\mathcal{Q}}$ is a set of amounts. We use q, q', and q_n ($n \in \mathbb{N}$) as meta-variables ranging over elements of $\mathcal{Q}$.

Constants. There is a set of constants corresponding to each of the four sorts.

1. Tc, Tc', Tc_n ($n \in \mathbb{N}$), are constant symbols in $\mathcal{T}$.
2. Sc, Sc', Sc_n ($n \in \mathbb{N}$) are constant symbols in TEMP.
3. Pc, Pc', Pc_n ($n \in \mathbb{N}$) are constant symbols in $\mathcal{P}$.
4. Qc, Qc', Qc_n ($n \in \mathbb{N}$) are constant symbols in $\mathcal{Q}$.

Variables. There is a set of variables corresponding to each of the four sorts.

1. Tv, Tv', Tv_n ($n \in \mathbb{N}$) are variables in $\mathcal{T}$.
2. Sv, Sv', Sv_n ($n \in \mathbb{N}$) are variables in TEMP.
3. Pv, Pv', Pv_n ($n \in \mathbb{N}$) are variables in $\mathcal{P}$.
4. Qv, Qv', Qv_n ($n \in \mathbb{N}$) are variables in $\mathcal{Q}$.

Function Symbols. There is a set of function symbols denoting functions from the Cartesian product of combinations of $\mathcal{T}$, TEMP, $\mathcal{P}$, and $\mathcal{Q}$ to any of these sets. Of particular importance are the following groups of functions.

1. Functions on Propositions:

- $\neg: \mathcal{P} \longrightarrow \mathcal{P}$, where $[\![\neg p]\!]$ is the proposition that it is not the case that $[\![p]\!]$.[10]
- $\wedge: \mathcal{P} \times \mathcal{P} \longrightarrow \mathcal{P}$, where $[\![p \wedge p']\!]$ is the proposition that it is the case that both $[\![p]\!]$ and $[\![p']\!]$.
- $\vee: \mathcal{P} \times \mathcal{P} \longrightarrow \mathcal{P}$, where $[\![p \vee p']\!]$ is the proposition that it is either the case that $[\![p]\!]$, $[\![p']\!]$, or both.
- $\Rightarrow: \mathcal{P} \times \mathcal{P} \longrightarrow \mathcal{P}$, where $[\![p \Rightarrow p']\!]$ is the proposition that if it is the case that $[\![p]\!]$, then it is the case that $[\![p']\!]$.
- $\Leftrightarrow: \mathcal{P} \times \mathcal{P} \longrightarrow \mathcal{P}$, where $[\![p \Leftrightarrow p']\!]$ is the proposition that it is the case that $[\![p]\!]$ if and only if it is the case that $[\![p']\!]$.

2. Temporal Functions:

   - $\prec: \mathcal{T} \times \mathcal{T} \longrightarrow \mathcal{P}$, where $[\![t \prec t']\!]$ is the proposition that $[\![t]\!]$ precedes $[\![t']\!]$.
   - $\sqsubseteq: \mathcal{T} \times \mathcal{T} \longrightarrow \mathcal{P}$, where $[\![t \sqsubseteq t']\!]$ is the proposition that $[\![t]\!]$ is a subinterval of $[\![t']\!]$.
   - $\sqsubset: \mathcal{T} \times \mathcal{T} \longrightarrow \mathcal{P}$, where $[\![t \sqsubset t']\!]$ is the proposition that $[\![t]\!]$ is a proper subinterval of $[\![t']\!]$.

3. Miscellaneous Functions:

   - Holds: $\textsc{Temp} \times \mathcal{T} \longrightarrow \mathcal{P}$, where $[\![\mathsf{Holds}(s,t)]\!]$ is the proposition that the temporary state $[\![s]\!]$ holds throughout the interval $[\![t]\!]$.
   - MHolds: $\textsc{Temp} \times \mathcal{T} \longrightarrow \mathcal{P}$, where $[\![\mathsf{MHolds}(s,t)]\!]$ is the proposition that the temporary state $[\![s]\!]$ *maximally* holds throughout the interval $[\![t]\!]$.
   - SDur: $\textsc{Temp} \times \mathcal{Q} \longrightarrow \mathcal{P}$, where $[\![\mathsf{SDur}(s,q)]\!]$ is the proposition that the typical duration of intervals over which $[\![s]\!]$ *maximally* holds is $[\![q]\!]$.
   - $\mathsf{Dur}: \mathcal{T} \times \mathcal{Q} \longrightarrow \mathcal{P}$, where $[\![\mathsf{Dur}(t,q)]\!]$ is the proposition that the duration of interval $[\![t]\!]$ is $[\![q]\!]$.
   - $<_{\mathcal{Q}}: \mathcal{Q} \times \mathcal{Q} \longrightarrow \mathcal{P}$, where $[\![q <_{\mathcal{Q}} q']\!]$ is the proposition that amount $[\![q]\!]$ is less than amount $[\![q']\!]$.
   - $\mathsf{Equiv}: \Sigma \times \Sigma \longrightarrow \mathcal{P}$, where Σ is any of the four sorts. $[\![\mathsf{Equiv}(\tau,\tau')]\!]$ is the proposition that $[\![\tau]\!]$ and $[\![\tau']\!]$ are coreferential, i.e., the two Meinongian objecta "pick out the same extension in some world" (Maida and Shapiro, 1982, p. 298).

Terms.

1. Constants and variables are terms of the appropriate sort.

[10]For a term, τ, $[\![\tau]\!]$ represents its denotation.

2. If f is a function symbol with domain A and range B, and if x is a tuple in A, then $f(x)$ is a functional term in B.

3. If x is a variable of sort Σ, τ is a term in Σ, and p is a term in $\mathcal{P}$ with one or more occurrences of τ and no occurrences of x, then $\forall x[p\{x/\tau\}]$ and $\exists x[p\{x/\tau\}]$ are terms in $\mathcal{P}$.[11] For all p, $[\![\forall x[p]]\!]$ is the proposition that, for all $\tau' \in \Sigma$, it is the case that $[\![p\{\tau'/x\}]\!]$. Similarly, $[\![\exists x[p]]\!]$ is the proposition that, for some $\tau' \in \Sigma$, it is the case that $[\![p\{\tau'/x\}]\!]$.

Rules of Inference. We assume the existence of standard inference rules (for example, introduction and elimination rules for connectives and quantifiers). However, because of assumptions underlying our theory, a couple of points need to be made explicit. In standard first-order predicate calculus, the interpretation of a rule of inference is that if the premises are true, then so is the conclusion. In our theory, truth is replaced with belief by Cassie. That is, if Cassie believes the premises (i.e., if they are in β), then she *may* also believe the conclusion. As (Shapiro, 1993) pointed out, it is important to note that rules of inference should not be taken as saying that if the premises are in β, then so is the conclusion. In particular, we do not assume β to be closed under deduction (Johnson and Shapiro, 2000b).

Axioms. Basically, axioms are propositions that we assume Cassie to believe, i.e., they form a subset of β. The first two axioms formalize the distributional patterns of states over time. **AS1** asserts that states are *cumulative* and **AS2** states that they are *divisive*.[12] Cumulativity means that, if a state holds over all proper subintervals of some interval, t, then it also holds over t. Divisitivity means that, if a state holds over an interval, then it holds over all of its subintervals. Together, the two axioms state that an interval over which a state holds is *homogeneous*; all of its parts look alike in some respect.

- **AS1.** $[\forall Tv'[Tv' \sqsubseteq Tv \Rightarrow \mathsf{Holds}(Sv, Tv')]] \Rightarrow \mathsf{Holds}(Sv, Tv)$
- **AS2.** $\mathsf{Holds}(Sv, Tv) \Rightarrow \forall Tv'[Tv' \sqsubseteq Tv \Rightarrow \mathsf{Holds}(Sv, Tv')]$

The following axiom makes explicit the semantics of MHolds.

- **AS3.** $\mathsf{MHolds}(Sv, Tv) \Rightarrow [\mathsf{Holds}(Sv, Tv) \wedge \neg\exists Tv'[\mathsf{Holds}(Sv, Tv') \wedge Tv \sqsubset Tv']]$

[11] For any term τ and any two terms τ_1 and τ_2 of the same sort, $\tau\{\tau_1/\tau_2\}$ is the result of replacing all instances of τ_2 in τ by τ_1.

[12] Interestingly, no one terminology is used by all authors addressing these properties. Here we follow the terminology used in (Krifka, 1989).

The following group of axioms characterize various properties of temporal relations. We start with properties of $\prec$.

- **AT1.** $Tv \prec Tv' \Rightarrow \neg[Tv' \prec Tv]$
- **AT2.** $[Tv_1 \prec Tv_2 \wedge Tv_2 \prec Tv_3] \Rightarrow Tv_1 \prec Tv_3$
- **AT3.** $\mathsf{Equiv}(Tv, Tv') \Rightarrow \neg[Tv \prec Tv']$

One may interpret the above axioms as saying that $\prec$ is a strict partial order. This is indeed the intuition. However, note that, technically, this is not correct, for $\prec$ is a function rather than a relation. To be precise, it is the set of pairs corresponding to the maximal subset of the domain of $\prec$ whose image is a subset of the deductive closure of β that may be interpreted as a strict partial order. For convenience, however, we shall allow ourselves to be a little bit sloppy and overlook the distinction.

Similarly, the following axioms characterize $\sqsubseteq$ as a partial order.

- **AT4.** $[Tv \sqsubseteq Tv' \wedge Tv' \sqsubseteq Tv] \Rightarrow \mathsf{Equiv}(Tv, Tv')$
- **AT5.** $[Tv_1 \sqsubseteq Tv_2 \wedge Tv_2 \sqsubseteq Tv_3] \Rightarrow Tv_1 \sqsubseteq Tv_3$
- **AT6.** $\mathsf{Equiv}(Tv, Tv') \Rightarrow Tv \sqsubseteq Tv'$

The interaction between $\prec$ and $\sqsubseteq$ is illustrated by the following axioms. **AT7** and **AT8** state that inclusion preserves order,[13] and **AT9** asserts that all intervals in $\mathcal{D}_{\mathcal{T}}$ are convex.

- **AT7.** $[Tv_1 \prec Tv_2 \wedge Tv_3 \sqsubseteq Tv_1] \Rightarrow Tv_3 \prec Tv_2$
- **AT8.** $[Tv_1 \prec Tv_2 \wedge Tv_3 \sqsubseteq Tv_2] \Rightarrow Tv_1 \prec Tv_3$
- **AT9.** $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_4 \wedge Tv_4 \prec Tv_3] \Rightarrow Tv_4 \sqsubseteq Tv_1$

A strict version of $\sqsubseteq$ is defined by the following axiom.

- **AT10** $Tv_1 \sqsubset Tv_2 \Leftrightarrow [Tv_1 \sqsubseteq Tv_2 \wedge \neg\mathsf{Equiv}(Tv_1, Tv_2)]$

Although we shall not explicitly state theorems to the effect, we can prove that $\sqsubset$ is a strict partial-order. We also assume the existence of axioms defining Equiv as an equivalence relation and $<_{\mathcal{Q}}$ as a strict linear order.

[13] These are the "left monotonicity" and "right monotonicity" axioms of (van Benthem, 1983, p. 67), respectively.

Given the above set of axioms, we may prove the following simple but important results (see Appendix B for the proofs). **TT1** states that temporal part-hood and temporal precedence are contradictory. **TT2** states that if parts of an interval precede each other, then they are *proper* parts of that intervals.

- **TT1.** $Tv_1 \prec Tv_2 \Rightarrow [\neg[Tv_2 \sqsubseteq Tv_1] \wedge \neg[Tv_1 \sqsubset Tv_2]]$
- **TT2.** $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3] \Rightarrow [Tv_2 \sqsubset Tv_1 \wedge Tv_3 \sqsubset Tv_1]$

7.3 Temporal Frames

Having outlined the basic formal machinery, we now move on to investigate the meta-theoretical structures articulating the domain of time intervals. First, let us restrict the discussion to those intervals in Cassie's "consciousness".

Definition 7.1 *For any sort Σ, define $\Psi(\Sigma)$ as the largest subset of $\mathcal{D}_\Sigma$ such that for every $[\![\tau]\!] \in \Psi(\Sigma)$, $[\![\tau]\!] \in \beta$ or there is a term τ' such that $[\![\tau']\!] \in \beta$ and τ is a subterm of τ'.*

Intuitively, $\Psi(\Sigma)$ represents those entities in $\mathcal{D}_\Sigma$ that Cassie conceives of.[14] It should be noted that $\Psi(\mathcal{P})$ is not identical to β, since Cassie may conceive of propositions she does not believe. For convenience, we shall henceforth drop the syntax-semantics distinction. Thus, we shall use "τ" in place of "$[\![\tau]\!]$" and "Σ" in place of "$\mathcal{D}_\Sigma$". In addition, where $p \in \mathcal{P}$, we shall often use "p" where what is intended is "$[\![p]\!] \in \beta$". This should not be confusing given the one-to-one correspondence between terms and their denotations.

Definition 7.2 *An interval $t \in \Psi(\mathcal{T})$ is an* **atomic interval** *if there is no $t' \in \Psi(\mathcal{T})$ such that $t' \sqsubset t$.*[15]

Note that an interval being atomic is not an intrinsic property of the interval itself, it is totally dependent on Cassie's *state of mind*. Cassie's state of mind may be represented by the set,

$$\Psi(\mathcal{T}) \cup \Psi(\textsc{Temp}) \cup \Psi(\mathcal{P}) \cup \Psi(\mathcal{Q})$$

This set changes with time, since Cassie may acquire more knowledge as time passes by. Consider the following sentences.

[14] See (Shapiro, 1991, p. 143) for a different presentation of the same notion.

[15] Atomic intervals are, therefore, similar to the *moments* of (Allen and Hayes, 1985).

(1) The vase fell on the floor.

(2) John tried to catch it.

Sentence (1) reports a punctual event. After hearing (1), Cassie would have certain beliefs about the time, t, of that event. None of these beliefs, however, are about any subinterval of t. In that *state of mind*, t is an atomic interval. Later, after hearing (2), this may change. Cassie would have a belief about a subinterval, t', of t; the interval over which John tried catching the falling vase. Now Cassie is in a state of mind in which t is not atomic. Indeed, to be more precise, one should relativize all definitions to Cassie's state of mind. We choose, however, to be more liberal while stressing that all of our definitions should be interpreted in the context of a particular state of mind.

The relations $\prec$ and $\sqsubset$ articulate the set $\Psi(\mathcal{T})$ giving rise to what we call temporal frames.

Definition 7.3 *A set $\Phi \subseteq \Psi(\mathcal{T})$ is a* **temporal frame** *if there exists $t \in \Psi(\mathcal{T})$ such that, for every $t' \in \Phi$, $t \sqsubset t'$ (is in β) or $t' = t$.*[16]

The above definition may be interpreted as saying that a temporal frame is a set of intervals that share a common subinterval. It should be noted that, in the above definition, "$t \sqsubset t'$ (is in β) or $t' = t$" is not equivalent to "$t \sqsubseteq t'$ (is in β)", the latter is equivalent to "$t \sqsubset t'$ (is in β) or $\mathsf{Equiv}(t', t)$ (is in β)".

The above definition is by itself not very interesting since it covers a lot of trivial cases. For instance, any singleton subset of $\Psi(\mathcal{T})$ is a temporal frame. A more conservative notion is required.

Definition 7.4 *A temporal frame Φ is* **maximal** *if there is no temporal frame Φ' such that $\Phi \subset \Phi'$.*

It could be shown that, if Φ is a maximal temporal frame (henceforth, **MTF**), then $\langle \Phi, \lambda x \lambda y(\beta \vdash x \sqsubset y \vee x = y)\rangle$ is a poset with a smallest element (which makes it a *meet semilattice* (Link, 1998)).[17]

Observation 7.1 *If Φ is an MTF, then the poset $\langle \Phi, \lambda x \lambda y(\beta \vdash x \sqsubset y \vee x = y)\rangle$ has a smallest element.*

[16]Note that '=' is a meta-theoretical predicate representing term (or denotation) identity. This is not to be confused with the object language 'Equiv' which corresponds to coreference of denotations.

[17]We are overloading '$\vee$', using it as meta-theoretical disjunction.

Proof. To show that $\langle \Phi, \lambda x \lambda y(\beta \vdash x \sqsubset y \vee x = y)\rangle$ has a smallest element, we need to prove that there is some $t \in \Psi(\mathcal{T})$ such that:

(i) for every $t' \in \Phi$, $\langle t, t'\rangle \in \lambda x \lambda y(\beta \vdash x \sqsubset y \vee x = y)$ and

(ii) $t \in \Phi$.

By Definition 7.3, there exists $t \in \Psi(\mathcal{T})$ such that for every $t' \in \Phi$ $t \sqsubset t' \in \beta$ or $t' = t$, which proves (i). Now, suppose that such an interval, t, is not in Φ. Then the set $\Phi \cup \{t\}$ is a super-set of Φ. Since, by designation, t is a sub-interval of every element of Φ, then, by Definition 7.3, $\Phi \cup \{t\}$ is itself a temporal frame. But, by Definition 7.4, this implies that Φ is not maximal, which leads to a contradiction. Therefore, $t \in \Phi$, which proves (ii). $\square$

The smallest elements of MTFs have an interesting property: they are atomic.[18]

Proposition 7.1 *If Φ is an MTF, then a smallest element of Φ is an atomic interval.*

Proof. Let t be a smallest element of Φ. Assume that t is not atomic. By Definition 7.2, there is some $t' \in \Psi(\mathcal{T})$ such that $t' \sqsubset t$. Since, for all $t'' \in \Phi$, $t'' \neq t$ implies $t \sqsubset t''$, then, by the transitivity of $\sqsubset$, t' is a proper sub-interval of all elements of Φ. Thus, by Definition 7.3, $\Phi \cup \{t'\}$ is a temporal frame. If $t' \notin \Phi$, then $\Phi \subset \Phi \cup \{t'\}$ and, by Definition 7.4, Φ is not an MTF, leading to a contradiction. On the other hand, if $t' \in \Phi$, then t cannot be a smallest element of Φ, which also leads to a contradiction. Therefore, t must be atomic. $\square$

Not only is a smallest element of an MTF atomic, but it is the *only* atomic interval therein.[19]

Theorem 7.1 *For every MTF, Φ, there is one and only one atomic interval in Φ.*

Proof. Let Φ be an MTF. By Observation 7.1, Φ contains a smallest interval, t, which, by Proposition 7.1, is atomic. Now, we need to show that t, is the only atomic interval in Φ. Since t is a smallest element of $\langle \Phi, \lambda x \lambda y(\beta \vdash x \sqsubset y \vee x = y)\rangle$, then, for every $t' \in \Phi$, $t' \neq t$ implies that $t \sqsubset t' \in \beta$. Thus, by Definition 7.2, for every $t' \in \Phi$, $t' \neq t$ implies that t' is not atomic. Therefore, t is the only atomic interval in Φ. Since Φ is arbitrary, then the result applies to all MTFs. $\square$

---

[18]In what follows, we shall be talking about "a smallest element of Φ" (where Φ is an MTF) referring to a smallest element of the poset $\langle \Phi, \lambda x \lambda y(\beta \vdash x \sqsubset y \vee x = y)\rangle$.

[19]The notion of MTFs is related to that of *filters* (van Benthem, 1983, Ch. I.4). Informally, van Benthem's definition of filters maps as follows onto our system. A subset F of $\Psi(\mathcal{T})$ is a filter if (i) $t \in F$ implies that all super-intervals of t are in F and (ii) t and t' are in F implies that their maximal common sub-interval (if one exists) is also in F. Accordingly, an MTF is a filter, but not every filter is an MTF. In particular, a filter may include more than one atomic interval.

The following corollary directly follows.

Corollary 7.1 *Every MTF has a unique smallest element.*

Proof. Follows directly from Observation 7.1, Proposition 7.1, and Theorem 7.1. □

Given the above results, one can outline an algorithm for computing the MTFs of a set $\Psi(\mathcal{T})$.

1. For every atomic interval $t_i \in \Psi(\mathcal{T})$, let $\Phi_i \longleftarrow \{t_i\}$.
2. For every $t \in (\Psi(\mathcal{T}) - \bigcup_i \Phi_i)$
    3. For every Φ_i
        4. If $t_i \sqsubset t$, let $\Phi_i \longleftarrow \Phi_i \cup \{t\}$.

The resulting Φ_i's are the MTFs of $\Psi(\mathcal{T})$. Note that, for each Φ_i, t_i is its unique smallest element. The intuitive construction of MTFs represented by the above algorithm assumes the following result which, together with Theorem 7.1, point to a one-to-one correspondence between MTFs and atomic intervals.

Theorem 7.2 *For every atomic interval, t, there is one and only one MTF to which t belongs.*

Proof. Let t be an atomic interval. Consider the set consisting of *all* intervals, t', satisfying the property $t \sqsubset t' \in \beta$ or $t' = t$. By Definition 7.3, this set is a temporal frame, and is maximal since it includes all such t's. Now, we need to show that t belongs to a unique MTF. Suppose that Φ_1 and Φ_2 are distinct MTFs to which t belongs. Then there is an interval, t', that belongs to Φ_1 and not Φ_2. Since t is atomic, then, by Theorem 7.1, it is the smallest element of both Φ_1 and Φ_2 and, therefore, $t \sqsubset t'$. By Definition 7.3, the set $\Phi_2 \cup t'$ is a temporal frame which, since $t' \notin \Phi_2$, is a super-set of Φ_2. Thus, by Definition 7.4, Φ_2 is not maximal, leading to a contradiction. Therefore, t belongs to only one MTF. Since t is arbitrary, then the result applies to all atomic intervals. □

Given Theorems 7.1 and 7.2, we shall use the notation $\Phi(t)$ (where t is atomic) to refer to the unique MTF whose smallest element is t. Going back to the above algorithm for the construction of MTFs, we can show that the union of the MTFs of a set $\Psi(\mathcal{T})$ is identical to $\Psi(\mathcal{T})$. More specifically, we can make the following observation.

Observation 7.2 *The collection of MTFs of $\Psi(\mathcal{T})$ constitutes a minimal cover of $\Psi(\mathcal{T})$.*

Proof. First, we show that the set of MTFs is a cover of $\Psi(\mathcal{T})$ and, then, we prove its minimality. Let C be the collection of MTFs of $\Psi(\mathcal{T})$. To show that C covers $\Psi(\mathcal{T})$, we only need to prove that, for every $t \in \Psi(\mathcal{T})$,

$$t \in \bigcup_{\Phi_i \in C} \Phi_i.$$

Consider an arbitrary $t \in \Psi(\mathcal{T})$. If t is atomic, then, by Theorem 7.2, $t \in \Phi(t)$. On the other hand, if t is not atomic, then there must be some atomic interval, t', such that $t' \sqsubset t$. But, by Definition 7.4, $t \in \Phi(t')$. Therefore, for every $t \in \Psi(\mathcal{T})$, there is some MTF, Φ, such that $t \in \Phi$. Thus, for every $t \in \Psi(\mathcal{T})$,

$$t \in \bigcup_{\Phi_i \in C} \Phi_i,$$

which means that the collection of MTFs cover $\Psi(\mathcal{T})$. To prove that the covering is minimal, we need to show that for every $\Phi \in C$, there is some $t \in \Psi(\mathcal{T})$ such that

$$t \notin (\bigcup_{\Phi_i \in C} \Phi_i) \setminus \Phi.$$

Let Φ be an arbitrary member of C. By Theorem 7.1, there is a unique atomic interval, t, such that $t \in \Phi$. By Theorerm 7.2, Φ is the only MTF to which t belongs. Therefore,

$$t \notin (\bigcup_{\Phi_i \in C} \Phi_i) \setminus \Phi.$$

Since Φ is an arbitrary member of C, then C minimally covers $\Psi(\mathcal{T})$. $\square$

It should be noted that the collection of MTFs does not partition $\Psi(\mathcal{T})$. The reason is that a non-atomic interval, t, may span more than one MTF, i.e., MTFs are not disjoint.

MTFs are, in general, similar to the situations of the situation calculus (McCarthy and Hayes, 1969); they represent snapshots, not of the universe, but of Cassie's conceptualization of it. A general, not essentially maximal, temporal frame corresponds to the *states of affairs* of situation semantics (Barwise and Perry, 1983). Note, though, that whereas the latter is objective, the former is mental. As fluents may hold in situations, states hold in MTFs.

Definition 7.5 *For every $s \in \Psi(\text{TEMP})$ and temporal frame Φ, s* **holds in** *Φ, if there is some $t \in \Phi$ such that* $\mathsf{Holds}(s, t)$.

Note, in particular, that if s holds in an MTF, $\Phi(t)$, then (by **AS2**) Cassie may conclude that $\mathsf{Holds}(s, t)$. In addition, if there is some t' such that $\mathsf{Holds}(s, t')$ and

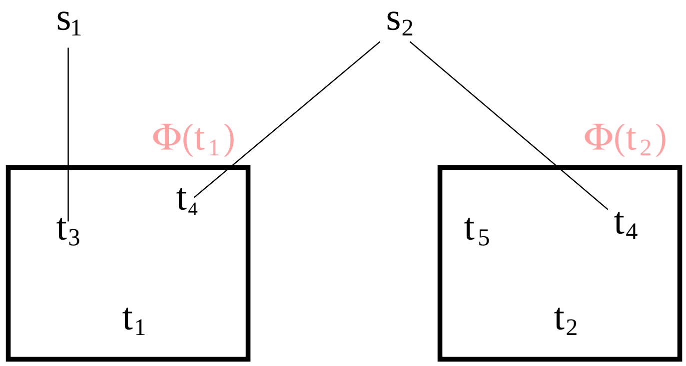

Figure 2: Two MTFs, $\Phi(t_1)$ and $\Phi(t_2)$. The state s_1 holds in $\Phi(t_1)$ and the state s_2 holds in both $\Phi(t_1)$ and $\Phi(t_2)$.

there is some atomic interval, t, such that $t \sqsubset t'$, then s holds in $\Phi(t)$ since otherwise $\Phi(t)$ would not be maximal. Thus, an MTF with a smallest element, t, corresponds to the set of all states that, Cassie believes or may conclude they, hold over t. Figure 2 shows two MTFs with the states holding in them. The MTFs are represented by rectangles, and their contents by meta-variables. The smallest element of an MTF is shown near the center of the bottom side of the rectangle. States are represented by meta-variables, with lines connecting states and time intervals standing for the Holds relation.

A precise characterization of the difference between two MTFs may be given an epistemic interpretation.

Definition 7.6 *The* **epistemic distance** *between two MTFs, Φ_1 and Φ_2, denoted $d_e(\Phi_1, \Phi_2)$, is the cardinality of their symmetric difference. That is, $d_e(\Phi_1, \Phi_2) = |\Phi_1 \triangle \Phi_2|$.*

For example, the epistemic distance between the two MTFs shown in Figure 2 is 4. Depending on their epistemic distance, two MTFs are more, or less, similar to each other as far as the states of affairs they correspond to are concerned. In fact, the set of MTFs together with d_e form a metric space.

Observation 7.3 *The function d_e defines a metric over the set of MTFs.*

Proof. For d_e to be a metric, it must satisfy the following (where $\Phi(t_1)$, $\Phi(t_2)$, and $\Phi(t_3)$ are MTFs):

(a) $d_e(\Phi(t_1), \Phi(t_2)) \geq 0$.

(b) $d_e(\Phi(t_1), \Phi(t_2)) = 0$ if and only if $\Phi(t_1) = \Phi(t_2)$.

(c) $d_e(\Phi(t_1), \Phi(t_2)) = d_e(\Phi(t_2), \Phi(t_1))$.

(d) $d_e(\Phi(t_1), \Phi(t_3)) \leq d_e(\Phi(t_1), \Phi(t_2)) + d_e(\Phi(t_2), \Phi(t_3))$.

The proof is straight-forward and follows from the definiton of d_e. In particular, the cardinality of the symmetric difference is a metric over any class of finite sets. (a), (b), and (c) are obvious, and (d) follows from the fact that $A \triangle B \subseteq (A \triangle C) \cup (C \triangle B)$, for any sets A, B, and C. $\square$

The above observation constrains epistemic distance to be non-negative. However, being a metric over *MTFs* imposes yet another constraint.

Theorem 7.3 *For any two distinct MTFs, $\Phi(t_1)$ and $\Phi(t_2)$, $d_e(\Phi(t_1), \Phi(t_2)) \geq 2$.*

Proof. Since $\Phi(t_1)$ and $\Phi(t_2)$ are distinct MTFs, then, by Thereom 7.2, t_1 and t_2 are distinct atomic intervals. By Theorem 7.1, an MTF may have one and only one atomic interval. It follows that, $\{t_1, t_2\} \subseteq \Phi(t_1) \triangle \Phi(t_2)$. Therefore, $d_e(\Phi(t_1), \Phi(t_2)) \geq 2$. $\square$

As shown above, the relation $\sqsubset$ provides the internal structure of MTFs. The relation $\prec$ provides the external structure.

Definition 7.7 *An MTF, $\Phi(t_1)$,* **precedes** *another MTF, $\Phi(t_2)$, (or $\Phi(t_2)$* **follows** *$\Phi(t_1)$) if and only if $t_1 \prec t_2$.*

Of course, since the precedence relation over MTFs is based on $\prec$, it is a strict partial order. Although there may be a situation where Cassie does not have any beliefs about the relative order of various MTFs, some structure may still be retrieved. In particular, MTFs form clusters corresponding to their intersections.

Definition 7.8 *For every $t \in \Psi(\mathcal{T})$, the* **span** *of t is the set $Span(t) = \{\Phi \colon \Phi$ is an MTF and $t \in \Phi\}$. For any $s \in \Psi(\textsc{Temp})$ and any set, A, of MTFs, s* **spans** *A if there is some $t \in \Psi(\mathcal{T})$ such that $\mathsf{Holds}(s,t)$ and $A \subseteq Span(t)$.*

The MTFs in the span of an interval, t, correspond to different *pieces* of t. Since intervals are convex, those MTFs form clusters that, although not internally ordered, are certainly *closer* to each other than to MTFs not containing t. Such clusters of MTFs form *episodes* in Cassie's memory: a collection of related and temporally contiguous events (see (Rumelhart et al., 1972; Tulving, 1972)). Cassie may not know the exact order of intervals within an episode, but she may know that one episode is earlier or later than another if they correspond to the spans of some time intervals, t_1 and t_2, where $t_1 \prec t_2$ (or vice versa). Note that this knowledge is essentially based on **AT7** and **AT8**.

Algorithm `move_NOW`

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\beta \longleftarrow \beta \cup \{*\mathsf{NOW} \prec t\}$.
3. $\mathsf{NOW} \longleftarrow t$.

Figure 3: The algorithm `move_NOW`.

7.4 The Passage of Time

A particularly interesting subset of MTFs forms a linearly-ordered chain corresponding to the experienced progression of time. Cassie's sense of temporal progression is modeled by a deictic metavariable, NOW, that assumes values from amongst the members of $\mathcal{T}$ (Almeida and Shapiro, 1983; Almeida, 1995; Shapiro, 1998; Ismail and Shapiro, 2000b). At any time, NOW *points* to a particular member of $\Psi(\mathcal{T})$. This represents Cassie's sense of the current time at the finest level of granularity (see Section 3).

The movement of time is represented by NOW's *moving* (changing its value) to a different term in $\Psi(\mathcal{T})$. Depending on what NOW exactly represents, there may, or may not, be restrictions on its movement. For example, if NOW represents a narrative *now-point* (Almeida and Shapiro, 1983; Almeida, 1995; ter Meulen, 1997), then there may be no restrictions at all on the values it assumes; NOW may freely hop around in $\Psi(\mathcal{T})$. This is because narration may go back and forth in time and may be about temporally unrelated episodes. On the other hand, if NOW represents the *real* present for an acting agent, which is how we are using it, then there certainly are restrictions on its movement. First, whenever it moves, NOW moves to a new term. That is, a change in the value of NOW is always associated with a change in $\Psi(\mathcal{T})$, since at least the new present enters into Cassie's consciousness. Second, values of NOW form a chain of times linearly-ordered by $\prec$.[20] For any $\Psi(\mathcal{T})$, NOW is always pointing to the greatest element of the chain (the *newest* present). The movement of time is thus modeled by the algorithm `move_NOW` shown in Figure 3. *NOW denotes the term pointed to by NOW (i.e., "*" is a dereferencing operator). Note that although a change in NOW (step 3) is always associated with a change in Cassie's state of mind (step 2), the converse is not necessarily true. It all depends on what causes NOW to move. In the current status of our theory, NOW moves whenever Cassie becomes aware of a change in the environment. The "environment" here does not include Cassie's own state of mind. Thus, Cassie's noticing that the walk-light turns from red to green, her starting to move, or her sensing that her battery is low (for a battery-operated Cassie) results in NOW moving. However, mere inferences

[20] Thus, the presented model is silent about the issue of *forgetting*.

Algorithm `initialize_NOW`

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\mathsf{NOW} \longleftarrow t$.

Figure 4: The algorithm `initialize_NOW`.

that do not involve any interaction with the environment (for example, inferring she can cross the street having perceived that the walk-light is green), and that change Cassie's state of mind, do not move NOW. Thus, the actual movement of NOW (i.e., the implementation of `move_NOW`) is taken care of at the PML. Generally, however, NOW may move with every inference step providing Cassie with a fine-grained sense of temporal progression.[21]

Algorithm `move_NOW` takes care of *changing* the value of NOW; we still need to account for *initializing* it. This is illustrated in Figure 4. Note that algorithm `initialize_NOW` is identical to algorithm `move_NOW` without the second step which requires a previous value of NOW. We make the following reasonable assumptions about the temporal career of the variable NOW.

1. Algorithms `move_NOW` and `initialize_NOW` are the only places where NOW is set. The first time NOW is set is by step 2 of `initialize_NOW`. Subsequent changes to NOW are the result of step 3 of `move_NOW`.
2. At any time there is at most one execution of algorithm `move_NOW` going on.

Together, the above two assumptions mean that the value of NOW changes sequentially with time. For ease of notation, we shall use numerical subscripts to refer to the successive values of NOW. Thus, $*\mathsf{NOW}_i$ is the value of NOW at a time earlier than the time at which $*\mathsf{NOW}_j$ is the value of NOW if and only if i is less than j. Thus, $*\mathsf{NOW}_1$ is the first value of NOW, $*\mathsf{NOW}_2$ is the second, and so on. If i is the largest subscript, then both $*\mathsf{NOW}$ and $*\mathsf{NOW}_i$ refer to the latest value of NOW. Algorithm `move_NOW` guarantees that this real-time ordering of the values of NOW corresponds to a $\prec$-chain of those values.

Theorem 7.4 *For all $i \in \mathbb{N}$, $\beta \vdash *\mathsf{NOW}_i \prec *\mathsf{NOW}_{i+1}$.*

Proof. Let $i \in \mathbb{N}$. According to the above-stated assumptions, the value of NOW can change from $*\mathsf{NOW}_i$ to $*\mathsf{NOW}_{i+1}$ only by executing algorithm `move_NOW`. At the time of executing the algorithm, $*\mathsf{NOW}_i$ is the latest value of NOW, and is, therefore,

[21] This is the idea behind *Active Logic* (Elgot-Drapkin and Perlis, 1990; Perlis et al., 1991). See Section 2 for an overview of that system.

identical to $*\mathsf{NOW}$. Step 1 introduces a new interval t. By step 2, $*\mathsf{NOW} \prec t \in \beta$. By step 3, t becomes the $i+1^{st}$ value of NOW, and is, thus, identical to $*\mathsf{NOW}_{i+1}$. Therefore, $\beta \vdash *\mathsf{NOW}_i \prec *\mathsf{NOW}_{i+1}$. Since i is arbitrary, then the result applies to all $i \in \mathbb{N}$. $\square$

7.5 Types of Intervals

The introduction of NOW induces a partitioning of the set of time intervals. In particular, we need to distinguish between those intervals that are introduced by step (1) of `move_NOW` (or `initialize-NOW`), and those introduced by assertions about states holding.

Definition 7.9 *An interval,* $t \in \Psi(\mathcal{T})$, *is a* **situation interval** *if there is some* $s \in \Psi(\textsc{Temp})$ *such that* $\mathsf{MHolds}(s,t)$ *and, in that case,* t *is* **a situation interval associated with** s. *An interval that is not a situation interval is a* **reference interval**.

A situation interval is a maximal stretch of time over which Cassie believes that some state holds. It is similar to what are referred to as *temporal traces* in the literature (Krifka, 1989; Link, 1998, for instance). For example, if the state associated with situation interval t is that of the walk-light being green, then, for Cassie to linguistically express t, she would use (variations of) the noun phrase "the time of the walk-light being green".[22] Situation intervals are, thus, ontologically-dependent on states (Chisholm, 1990; Pianesi and Varzi, 1996).[23]

The role of situation intervals in the theory is very crucial; they are used in such a way to model the persistence of states as NOW moves. To assert that a state, s, holds in the present, two propositions are involved:

1. $\mathsf{MHolds}(s,t)$, where t is a newly-introduced interval, and
2. $*\mathsf{NOW} \sqsubset t$.

As NOW moves, s's persistence is modeled by including each new value of NOW as a sub-interval of t. Of course, this requires that t be associated *only* with the state s. This is captured by the following axiom.

[22] Cassie's use of "the time" or "a time" in expressing t depends on whether there is more than one event of s holding.

[23] For the sake of completeness and precision, a note should be made here. In order not to complicate the discussion and introduce notions that are beyond the scope of this paper, we are only considering *states* and not *events*. In a more complete presentation of our theory, events play an important role, and situation intervals are associated, not only with states holding, but also with events occurring.

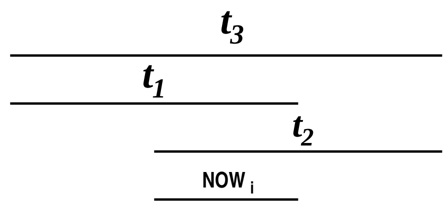

Figure 5: An impossible situation: s maximally holds over the overlapping intervals t_1 and t_2.

Axiom 7.1 *For every every* $s \in \Psi(\textsc{Temp})$ *and* $i \in \mathbb{N}$, *if* $\beta \vdash \mathsf{Holds}(s, {*}\mathsf{NOW}_i)$, *then there exists a unique situation interval,* $t \in \Psi(\mathcal{T})$, *such that*

1. $\beta \vdash \mathsf{MHolds}(s, t)$ *and* $\beta \vdash {*}\mathsf{NOW}_i \sqsubset t$, *and*
2. *for every* $s' \in \Psi(\textsc{Temp})$, *if* $\beta \vdash \mathsf{MHolds}(s', t)$, *then* $s' = s$.

The uniqueness of the situation interval required by the axiom may actually be proved. However, to do that, we will need to introduce more notions that would complicate the exposition and that are only needed for the proof. Therefore, we only sketch an informal proof here. Suppose that there are two distinct intervals, t_1 and t_2, such that they both satisfy (1) in Axiom 7.1. Thus, t_1 and t_2 overlap, their common sub-interval being ${*}\mathsf{NOW}_i$. Without loss of generality, assume that t_1 starts before t_2. The situation is shown in Figure 5. Now, consider the interval t_3, the sum of t_1 and t_2. By the cumulativity of states, s holds over t_3. But since both t_1 and t_2 are proper sub-intervals of t_3, then axiom **AS3** is violated. That is, s does not hold maximally over t_1 or t_2.

Reference intervals, on the other hand, are not associated with particular states; they designate temporal perspectives from which Cassie *views* a situation. In particular, reference intervals are used to represent different granularities of the present. Among the collection of reference intervals, an important sub-collection is the collection of NOW-**intervals**, that made up of those intervals that were once the value of NOW. NOW-intervals are intervals representing the present at the finest level of granularity.

Axiom 7.2 *For all* NOW*-intervals,* t, *the following holds:*

1. t *is a reference interval.*
2. *There is no reference interval,* $t' \in \Psi(\mathcal{T})$, *such that* $t' \sqsubset t$.

Requiring NOW-intervals to be reference intervals excludes them form being maximal intervals over which a state is asserted to hold. That is, a proposition of the form MHolds(s,*NOW$_i$), though syntactically and semantically valid, is pragmatically not possible. This reflects our discussion above that a super-interval of *NOW is introduced whenever a state is asserted to be holding in the present.

The reference interval pointed to by NOW (i.e., *NOW) is expressible by the English "now". A reference interval may also be given a *value*, "3:45 p.m." for instance. Otherwise, Cassie cannot linguistically-express reference intervals, they only determine the tense and aspect of sentences produced by Cassie (Almeida, 1995). Reference intervals are similar (but not identical) to the reference times of (Reichenbach, 1947), and are not to be confused with the reference intervals of (Allen, 1983).

7.6 The Chain of Experience

The following axiom states a principle that follows from our informal assumption that NOW moves whenever there is a change.

Axiom 7.3 (The First Principle of Change) *NOW *is always atomic.*

Informally, suppose that *NOW is not atomic, then there is some t such that $t \sqsubset$*NOW. By Axiom 7.2, t cannot be a reference interval. Therefore, t is a situation interval. By Definition 7.9, there is a state s that maximally holds over t. But if s maximally holds *within* *NOW, then Cassie became aware of changes (s starting and/or ceasing to hold) without NOW moving. Since NOW moves whenever Cassie is aware of a change, then *NOW must be atomic. Note that, *theoretically*, a general NOW-interval need not be atomic. For example, Cassie may be told that some state held within *NOW$_i$, where *NOW$_i \prec$ *NOW. Thus, NOW-intervals are atomic so long as they are present (i.e., pointed to by NOW), as they become past, they may no longer be atomic. Nevertheless, note that except for *NOW, which is expressible by "now", one cannot refer to NOW-intervals in natural language. That is, once they become past, we cannot tell Cassie anything about them. Thus, we shall assume as a working hypothesis that *all* NOW-intervals are atomic.[24] This way, we can talk about NOW-**MTF**s. An MTF, Φ(*NOW$_i$), is said to be the i^{th} NOW-MTF. Φ(*NOW) will be referred to as the **current** MTF. Corresponding to the chain of NOW-intervals, NOW-MTFs form a chain ordered by the "precedes" relation, with the current MTF being the last element in the chain.

The chain of NOW-MTFs is not only distinguished for being linearly-ordered; there is a genuine difference between NOW-MTFs and other MTFs in Cassie's mind.

[24] Even if a NOW-interval is given a value, such as "3:45", we assume that these values are, linguistically, moments of time within which it is not reasonable to assert that something happened.

NOW-MTFs comprise direct, first-person, experiences by Cassie. At least some of the states that hold in a NOW-MTF have been directly perceived by Cassie. Cassie's knowledge of states holding in other MTFs is either the result of inference or direct assertion, but never bodily feedback which may only take place in the present, within some NOW-MTF. As shall be shown later, it is because of this distinction that Cassie may have a *feel* of the duration of a NOW-MTF (or the smallest element thereof), but only *knowledge* of the duration of a non-NOW-MTF.

NOW moves when, and only when, Cassie becomes aware of a change. There are two comments to make about this assertion. First, not any change moves NOW. For example, if Cassie infers that the walk-light changed from red to green yesterday, NOW shouldn't move. NOW moves when, and only when, some state holding in the current MTF ceases to hold, or some state not holding in the current MTF starts to hold. Note that, in such cases, NOW *must* move to reflect the state's holding, or not holding, being past. Second, the "when" part is sanctioned by the First Principle of Change (Axiom 7.3). The "only when" part is validated by another principle regarding the epistemic distances among NOW-MTFs. First, a definition.

Definition 7.10 *Two* NOW-*MTFs,* $\Phi(*\mathsf{NOW}_i)$ *and* $\Phi(*\mathsf{NOW}_j)$*, are* **epistemically equivalent** *if* $d_e(\Phi(*\mathsf{NOW}_i), \Phi(*\mathsf{NOW}_j)) \leq 2$.

What does the epistemic equivalence of $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_j)$ imply? It implies that, if distinct, $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_j)$ differ only in their smallest elements, their location in time if you will.

Proposition 7.2 *If* $\Phi(*\mathsf{NOW}_i)$ *and* $\Phi(*\mathsf{NOW}_j)$ *are distinct, epistemically-equivalent* NOW-*MTFs, then* $\Phi(*\mathsf{NOW}_i) \triangle \Phi(*\mathsf{NOW}_j) = \{*\mathsf{NOW}_i, *\mathsf{NOW}_j\}$.

Proof. Since $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_j)$ are distinct, then, by Theorem 7.2,

$$\{*\mathsf{NOW}_i, *\mathsf{NOW}_j\} \subseteq \Phi(*\mathsf{NOW}_i) \triangle \Phi(*\mathsf{NOW}_j).$$

But since $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_j)$ are epistemically-equivalent, then, by Definition 7.6, $|\Phi(*\mathsf{NOW}_i) \triangle \Phi(*\mathsf{NOW}_j)| \leq 2$. Thus, by Theorem 7.3,

$$|\Phi(*\mathsf{NOW}_i) \triangle \Phi(*\mathsf{NOW}_j)| = 2.$$

Therefore, $\Phi(*\mathsf{NOW}_i) \triangle \Phi(*\mathsf{NOW}_j) = \{*\mathsf{NOW}_i, *\mathsf{NOW}_j\}$. $\square$

Other than their smallest elements, two epistemically-equivalent NOW-MTFs share all their intervals. Most importantly, they share all their situation intervals, and thus the states that hold in $\Phi(*\mathsf{NOW}_i)$ are exactly those that hold in $\Phi(*\mathsf{NOW}_j)$. There are two points to note.

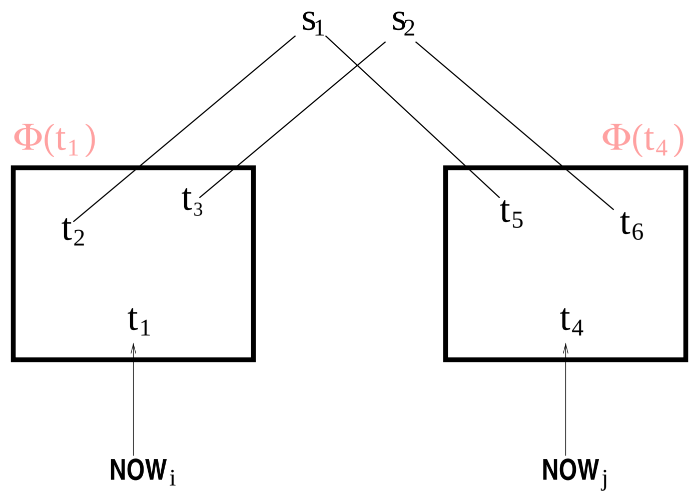

Figure 6: Two NOW-MTFs sharing the same set of states but not epistemically equivalent.

1. Epistemic equivalence (formally, $\lambda x \lambda y(d_e(x,y) \leq 2)$) is an equivalence relation. Symmetry stems from the commutativity of d_e (property (c) of d_e in the proof of Observation 7.3). Transitivity is based on noting that $\Phi(*\mathsf{NOW}_i) \setminus \{*\mathsf{NOW}_i\} = \Phi(*\mathsf{NOW}_j) \setminus \{*\mathsf{NOW}_j\}$, which follows from Proposition 7.2. Reflexivity follows from symmetry and transitivity, or could be independently established using property (b) of d_e.

2. As pointed out above, if $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_j)$ are epistemically equivalent, then the same collection of states hold in both. The converse is not true, however. That is, two NOW-MTFs may correspond to the same collection of states, yet fail to be epistemically equivalent. The reason is that the above definition is based on MTFs, mere collections of intervals, not states corresponding to MTFs. Thus, what matters is not whether the two MTFs share the same states, but rather, whether they share the same *events* of the same states holding. Figure 6 depicts two NOW-MTFs with the same states holding in them, yet with an epistemic distance of 6.

To ensure that NOW moves only when there is some change, we adopt the following principle.

Axiom 7.4 (The Second Principle of Change) *For every $i \in \mathbb{N}$, $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_{i+1})$ are not epistemically equivalent.*

Thus, the structure of Cassie's memory is *concise* (Williams, 1986): Cassie's mind is not populated with chains of NOW-MTFs that only differ in their smallest elements. If a situation interval is in $\Phi(\text{*NOW}_i) \setminus \Phi(\text{*NOW}_{i+1})$, then a state holding in $\Phi(\text{*NOW}_i)$ has ceased to hold; if it is in $\Phi(\text{*NOW}_{i+1}) \setminus \Phi(\text{*NOW}_i)$, then a state not holding in $\Phi(\text{*NOW}_i)$ has started to hold.

Theoretically, two consecutive NOW-MTFs may be disjoint. Nevertheless, except for the perceptually-crudest of agents, this seems very unlikely. First, this is certainly not the case for humans; our deep sense of the continuity (or density) of time would probably vanish if it wasn't for the strong overlap between consecutive NOW-MTFs. If many changes appear to us to happen simultaneously, we would probably not be able to make any sense of the world. Second, as pointed out by many authors (McDermott, 1982; Shoham and McDermott, 1988; Morgenstern, 1996; Shanahan, 1997, to mention a few), most states are persistent, and change is generally an exception. Thus, typically, a state would span more than one MTF. However, this does not rule out cases in which two consecutive NOW-MTFs may have no states in common. For example, consider an incarnation of Cassie where her only task is to monitor the state of some gauges. The readings of gauges are prone to very rapid change, and Cassie only *samples* those readings, say, every five seconds. Naturally, the reading of any gauge would change from one sample to the next. In such a situation, every sample corresponds to a different NOW-MTF, and one may expect consecutive samples to be totally different from one another.

The above principle notwithstanding, it should be noted that the epistemic distance between two NOW-MTFs, $\Phi(\text{*NOW}_i)$ and $\Phi(\text{*NOW}_j)$ does not necessarily increase with their *temporal distance*, $|i - j|$. Figure 7 shows an example where $\Phi(t_4)$ separates the epistemically equivalent NOW-MTFs $\Phi(t_1)$ and $\Phi(t_6)$. Intuitively, NOW moves from t_1 to t_4 as a result of s_3 starting to hold. It then moves to t_6 when s_3 ceases to hold. Thus, except for their location in time, the states of affairs corresponding to $\Phi(t_1)$ and $\Phi(t_6)$ are identical. Figure 8 shows a situation where the two epistemically equivalent, $\Phi(t_1)$ and $\Phi(t_6)$, are separated by $\Phi(t_4)$ which is missing a situation interval, t_2 that belongs to both. Note, however, that such a situation is impossible, for the convexity of intervals necessitates that t_2 be in $\Phi(t_4)$ (see **AT9**). Indeed, one may prove the following result.

Proposition 7.3 *If* $\Phi(\text{*NOW}_i)$ *and* $\Phi(\text{*NOW}_{i+n})$ $(n > 1)$ *are epistemically-equivalent* NOW-*MTFs, then for every* m*,* $0 < m < n$*,* $\Phi(\text{*NOW}_i) \setminus \{\text{*NOW}_i\} \subset \Phi(\text{*NOW}_{i+m})$.

Proof. Since $\Phi(\text{*NOW}_i)$ and $\Phi(\text{*NOW}_{i+n})$ are epistemically-equivalent, then, by Proposition 7.2, for every t, $t \in \Phi(\text{*NOW}_i) \setminus \{\text{*NOW}_i\}$ implies $t \in \Phi(\text{*NOW}_{i+n})$. Now suppose that $t \in \Phi(\text{*NOW}_i) \setminus \{\text{*NOW}_i\}$ and that it is not the case that $t \in \Phi(\text{*NOW}_{i+m})$, for some m $(0 < m < n)$. Therefore, $\text{*NOW}_i \sqsubset t$ and $\text{*NOW}_{i+n} \sqsubset t$, but $\text{*NOW}_{i+m} \not\sqsubset t$. But, by Theorem 7.4 and the transitivity of $\prec$ (**AT2**), $\text{*NOW}_i \prec \text{*NOW}_{i+m}$ and $\text{*NOW}_{i+m} \prec \text{*NOW}_{i+n}$. By the convexity of t (**AT9**) and **TT2**, $\text{*NOW}_{i+m} \sqsubset t$, which leads to a contradiction. Therefore, $t \in \Phi(\text{*NOW}_{i+m})$,

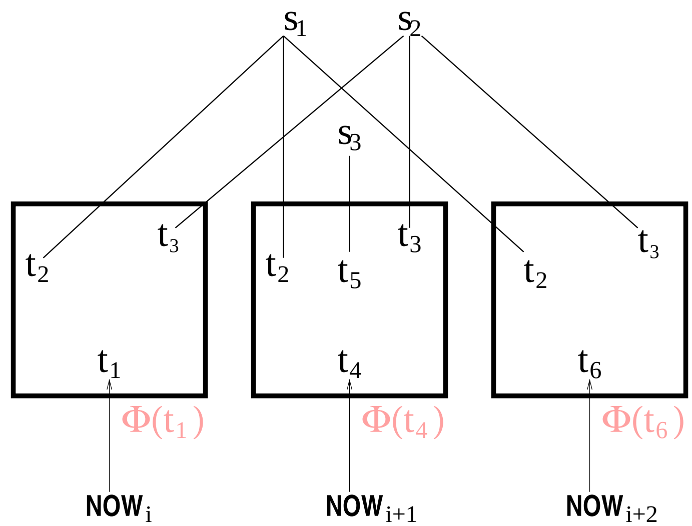

Figure 7: MTF $\Phi(t_4)$ separates the epistemically equivalent $\Phi(t_1)$ and $\Phi(t_6)$.

for all m $(0 < m < n)$. Since t is arbitrary, then $\Phi(*\mathsf{NOW}_i) \setminus \{*\mathsf{NOW}_i\} \subset \Phi(*\mathsf{NOW}_{i+m})$, for all m $(0 < m < n)$. $\square$

8 The Dynamics of Time

In the previous section, we investigated the meta-theoretical structure of time and outlined general principles that govern the movement of NOW. In this section, we look in more detail at how Cassie's beliefs evolve over time and how this interacts with the dynamics of NOW. Before doing that, however, we need to digress.

8.1 Consistency Over Time

As an agent acting and reasoning in a dynamic world, Cassie needs to be capable of handling failure. Failure manifests itself in two ways. First, Cassie may fail to perform some action or achieve some goal. This issue is investigated elsewhere (Ismail and Shapiro, 2000a) and will not be discussed in any detail here. Second, which is what concerns us, Cassie may fail to *reason* correctly about the domain. This happens when Cassie's beliefs space, β, is inconsistent, i.e., when it contains contradictory beliefs. In the theory presented here, there will be times when Cassie would make default assumptions about the persistence of states (see, in particular, Sections 10 and 11). It is possible that these assumptions be simply false, and should Cassie be aware of that, β would be inconsistent. Belief revision within

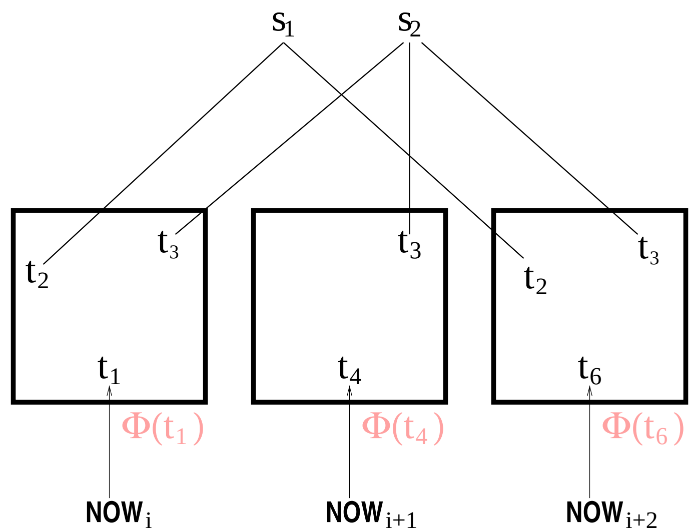

Figure 8: An impossible situation; t_2 must be in $\Phi(t_4)$ by the convexity of intervals.

the SNePS tradition has gone a long way (Martins and Shapiro, 1988; Johnson and Shapiro, 2000a; Johnson and Shapiro, 2000b; Johnson and Shapiro, 2000c) and inconsistency may, in fact, be handled appropriately. However, belief revision in a theory like the one we are presenting here is more complicated. In particular, we impose certain meta-theoretical constraints on the representation of time and its progression (for example, see Axioms 7.1 and 7.2). Whatever belief revision might do to resolve a contradiction, it should do it while observing these constraints. Currently, the belief revision system is not thus integrated with our theory of time. Therefore, should β become inconsistent, we would not be able to verify that the principles and axioms constraining our theory are not violated. It should be noted, however, that this would only be the case if the inconsistency involves beliefs about time, and, more specifically, if it involves beliefs about states holding in the present. Therefore, in what follows, we will make the assumption that beliefs about current states are never contradicted. More precisely, we will assume that, at no time, do pairs of propositions of the form $\mathsf{Holds}(s,^*\mathsf{NOW})$ and $\neg\mathsf{Holds}(s,^*\mathsf{NOW})$, explicitly or implicitly, co-exist in β. This would allow us to present results that we would not otherwise be able to formally prove, unless a complete theory of temporal belief revision is presented. Note that these results are not *wrong*, it is just that some of them presume consistency.

Another kind of inconsistency, one that our theory tolerates and, indeed, endorses, exists at the non-logical, meta-theoretical level. This kind of inconsistency is symptomatic of the need to move NOW. It involves violations of the principles and constraints of the theory that are not the result of inconsistencies in β, but

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s, t)\}$.
3. `move_NOW`.
4. $\beta \longleftarrow \beta \cup \{\mathsf{*NOW} \sqsubset t\}$.

Figure 9: What happens when Cassie senses that the state s holds.

are a natural side-effect of the progression of time. For example, Theorem 12.3 in Section 12.3 states that whenever Cassie is perceiving some state holding, then she believes that it does. Naturally, there will be short periods of time when this theorem does not apply. In particular, those are the times it takes to update Cassie's belief space (and other components of the system) in order to reflect the newly-perceived information. We will, therefore, assume that those are 0-periods—that it takes no time, at the theoretical level, to perform any required updates and *restabilize* the system.

8.2 A Note on the Frame Problem

Algorithm `move_NOW` outlined in Section 7.4 merely represents how temporal progression is modeled. However, it does not express everything that happens when NOW moves. Whenever NOW moves, a new MTF is constructed, namely $\Phi(\mathsf{*NOW})$. What do we know about this MTF, and how can it be constructed? Suppose the new MTF is the i^{th} NOW-MTF. One thing we learned in Section 7.6 is that $\Phi(\mathsf{*NOW}_i)$ is not epistemically equivalent to $\Phi(\mathsf{*NOW}_{i-1})$. If NOW moves due to some state's ceasing, then the situation interval, in $\Phi(\mathsf{*NOW}_{i-1})$, associated with that state should not be in $\Phi(\mathsf{*NOW}_i)$. On the other hand, if NOW moves due to some state's starting to hold, then a new situation interval associated with that state should be in $\Phi(\mathsf{*NOW}_i)$. The algorithm in Figure 9 outlines what happens when Cassie senses that some state, s, holds. Cessation will be discussed below. But, in this section, we are concerned with a slightly different issue.

It is often the case that many states holding in $\Phi(\mathsf{*NOW}_{i-1})$ continue to hold in $\Phi(\mathsf{*NOW}_i)$, and there should be some way of incorporating them into the new MTF. Of course what is lurking between the lines here is the frame problem (McCarthy and Hayes, 1969). We shall not delve into a long discussion of, nor propose any ingenious solutions to, the notorious problem; the literature is full of such discussions and proposals (Shoham and McDermott, 1988; Kautz, 1986; Shoham, 1986; Pylyshyn, 1986; Brown, 1987; Reiter, 1991; Morgenstern, 1996; Shanahan, 1997, to mention a few). Rather, we are going to make some informal remarks. Within the theoretical framework that has been developed in the previous sections, the frame problem manifests itself in determining which states that held in $\Phi(*\mathsf{NOW}_{i-1})$ continue to hold in the new $\Phi(*\mathsf{NOW}_i)$. First of all, we do not believe that there is a single, albeit elusive, solution to the problem; depending on the type of state, there may be different ways of determining whether it holds in $\Phi(*\mathsf{NOW}_i)$.

Armed with a theory of belief revision (Martins and Shapiro, 1988; Johnson and Shapiro, 2000c), we may adopt an off-the-shelf monotonic approach to solving the frame problem (Reiter, 1991; Thielscher, 1999, for example). These approaches provide logical solutions to the problem. That is, they involve Cassie's *reasoning* about whether a certain state continues to hold. In many cases this is reasonable. In particular, the frame problem is usually discussed in one of two contexts:

1. the context of planning, where an agent needs to predict the state of the environment following any of its planned actions, and

2. the context of reasoning about (what have confusingly come to be known as) narratives. Basically, a narrative is a world description in the form of a sequence of events and some states that hold at various points. Given a narrative, a reasoning system is asked to make predictions about which states hold (or not) at various points in the history of the world (the notorious Yale Shooting scenario is a typical example (Hanks and McDermott, 1987)).

In such settings, abstract reasoning is the only way to account for the persistence of states, and almost all the work that has been done on the frame problem is concerned with the kinds of axioms and reasoning systems required for a robust and efficient account of persistence that allows only *natural* predictions. The situation that we have here is neither one of planning nor of reasoning about narratives. We have Cassie out there, in the world, reasoning, acting, and perceiving the environment while maintaining a sense of time. We are not primarily interested in reasoning about past or future states. (This is not to say that these are trivial or unimportant issues; they certainly are not.) Rather, our main concern is to naturally account for Cassie's awareness of *present* states. Granted, reasoning is still needed for this task—for projecting states from $\Phi(*\mathsf{NOW}_{i-1})$ onto the new MTF $\Phi(*\mathsf{NOW}_i)$. Nevertheless, there is at least a subset of the states holding in $\Phi(*\mathsf{NOW}_{i-1})$ that should be incorporated in the new MTF without any reasoning on the part of Cassie. Let us motivate this with an example. The following (successor state) axiom appears in (Reiter, 1998, p. 553):

$$going(l, l', do(a, s)) \equiv (\exists t)a = startGo(l, l', t) \vee going(l, l', s) \wedge \neg(\exists t)a = endGo(l, l', t).$$

The axiom states conditions under which it could be concluded that the state of going from one location to another holds in a particular situation. Now, suppose that this is the *current* situation (or, in our, terminology, the current NOW-MTF). The crucial question here is *who is going?* There is some agent that does not appear in the axiom but is certainly implied. The axiom itself represents some useful piece of knowledge that an agent, Cassie for example, may use to reason about states of "going". But using this axiom to conclude that some agent is *still* "going" makes sense only when the implicit agent is *not* Cassie, the reasoning agent. If Cassie is the one who is "going", then she can conclude that she is still "going", not because of her general knowledge of causes and effects, nor because it is reasonable to *assume* so, but because *she* is actually "going". Agents do not need to reason about the *current* states of their own bodies, or those of the perceived environment, they have first-person direct access to those states through perception and proprioception.

What we are suggesting here is that, for states of Cassie's own body, or of the perceived environment, whether they hold in the new MTF should be taken care of, not at the KL, but at the PML, as part of temporal progression routines. How may that be done? This is the topic of the next section.

8.3 Modalities

To account for the persistence of bodily or perceived states, we employ a set of meta-logical variables, $\mathcal{M}$, corresponding to various agent modalities. (As shall be seen below, we interpret "modality" in a very broad sense.) The set $\mathcal{M}$ is partitioned into two sets: $\mathcal{M}_{\text{prop}}$ for proprioception, and $\mathcal{M}_{\text{per}}$ for perception. Each variable in $\mathcal{M}_{\text{prop}}$ contains a proposition representing the state currently occupying the corresponding modality. On the other hand, $\mathcal{M}_{\text{per}}$ represents what each modality conveys, about the external world, to Cassie. For example, a particular incarnation of Cassie may have the following set of proprioception modality variables:

$$\mathcal{M}_{\text{prop}} = \{\mathsf{VISION}_{\text{prop}}, \mathsf{AUDITION}_{\text{prop}}, \mathsf{LOCOMOTION}_{\text{prop}}, \mathsf{HANDLING}_{\text{prop}}\}$$

Such a Cassie has visual and auditory capabilities in addition to, maybe, wheels for movement and an arm for manipulating objects. Different processes performed by Cassie, or states of her body, occupy different modalities. For example, $\mathsf{VISION}_{\text{prop}}$ may contain the proposition "I am looking toward the green light" or "My eyes are shut" (where the first person refers to Cassie), $\mathsf{LOCOMOTION}_{\text{prop}}$ may contain the proposition "I am moving from l to l'" or "I am standing still", $\mathsf{AUDITION}_{\text{prop}}$ may contain "I am talking to Stu", and so on. Thus, $\mathcal{M}_{\text{prop}}$ represents what each modality is being *used for*.

The same Cassie may have a similar set of perception modality variables:[25]

[25] Typically, there are fewer perception modalities than there are proprioception modalities. For example, the locomotion system does not convey anything about the external environment, only its own state, and likewise for the handling system (unless we consider tactile perception).

$$\mathcal{M}_{\text{per}} = \{\mathsf{VISION}_{\text{per}}, \mathsf{AUDITION}_{\text{per}}\}$$

For example, $\mathsf{VISION}_{\text{per}}$ may contain the proposition "The green light is on" or "The block is on the table" and $\mathsf{AUDITION}_{\text{per}}$ may contain the proposition "The alarm is sounding" or "The radio is on".

In general, modality variables point to elements of β. If the modality corresponding to the variable $\mu \in \mathcal{M}_{\text{prop}}$ is occupied by some state, $[\![s]\!]$, then $*\mu = \mathsf{MHolds}(s,t)$, where t is the situation interval associated with s, such that $*\mathsf{NOW} \sqsubset t$. A similar idea underlies perception modality variables with the following provision. $*\mu = \mathsf{MHolds}(s,t)$ if $[\![s]\!]$ is a state perceived via the modality corresponding to μ. One main difference between perception and proprioception modality variables is that the former, but not the latter, take as values *sets* of propositions. Thus, for every $\mu \in \mathcal{M}_{\text{per}}$, there is a set of states and associated situation intervals such that $*\mu = \{\mathsf{MHolds}(s_i, t_i) : *\mathsf{NOW} \sqsubset t_i\}$. For example, Cassie not only sees that the block is on the table, but also that it is, say, red, and that she is close enough to pick it up. Thus, in general, a single perception modality may simultaneously convey more than one piece of information about the external world. On the other hand, a proprioception modality may be occupied by only *one* bodily state or process. This idea is stated precisely by the following axiom.

Axiom 8.1 *For every $i \in \mathbb{N}$, and $\mu \in \mathcal{M}_{\text{prop}}$, at $[\![*\mathsf{NOW}_i]\!]$, there is one, and only one, $s \in \Psi(\textsc{Temp})$, such that $[\![s]\!]$ occupies the modality corresponding to μ.*

In the above statement, $[\![*\mathsf{NOW}_i]\!]$ refers to the interval of *real* time during which the value of NOW is $*\mathsf{NOW}_i$. The important thing to note about this axiom is that it does *not* state that, at any time, a proprioception modality *variable* is occupied by a unique term in our logic; this is something that our theory has to ensure. The axiom merely states the *semantic* counterpart of this assertion. Note also that, not only do we require any proprioceptual modality to be occupied by a single state, but that there is *always* some state occupying any given proprioceptual modality. This might seem strange given that there may be times when Cassie is not *using* some, or all, of her modalities. However, as mentioned above, we interpret the notion of modality in a very broad sense; passive states such as "I am standing still" or "My eyes are shut" *occupy* modalities in our theory (in this case, $\mathsf{LOCOMOTION}_{\text{prop}}$ and $\mathsf{VISION}_{\text{prop}}$, respectively). Thus, even if Cassie is not using its sensors and effectors (in the traditional sense of "using"), states corresponding to these resources being idle (or available) are legitimate modality occupiers.

In what follows, we make a couple of assumptions about modalities and states that occupy, or are perceived by, them. First, we assume that, for every bodily state, there is a unique set of proprioception modalities that it occupies if and when it holds. Formally, there is a function, Mod_{prop}, from $\textsc{Temp}$ to the power set of

$\mathcal{M}_{\text{prop}}$ mapping each state to the proprioception modality variables corresponding to modalities it uses when it holds. Note that, for non-bodily states, the value of Mod_{prop} is simply the empty set. Thus, bodily states may be formally identified as those members of TEMP for which the value of Mod_{prop} is a non-empty set.

An important property of Mod_{prop} is stated by the following axiom.

Axiom 8.2 *For every* $s \in \text{TEMP}$, $\mu \in \mathcal{M}_{\text{prop}}$, *and* $i \in \mathbb{N}$, *at* $[\![*\mathsf{NOW}_i]\!]$, *if* $[\![s]\!]$ *occupies the modality corresponding to* μ, *then* $\mu \in \text{Mod}_{\text{prop}}(s)$ *and, for every* $\mu' \in \text{Mod}_{\text{prop}}(s)$ $[\![s]\!]$ *occupies the modality corresponding to* μ'.

For example, if $\text{Mod}_{\text{prop}}(s) = \{\mathsf{VISION}_{\text{prop}}, \mathsf{LOCOMOTION}_{\text{prop}}\}$, then, whenever s holds, it occupies *both*, and *only*, $\mathsf{VISION}_{\text{prop}}$ and $\mathsf{LOCOMOTION}_{\text{prop}}$. The following axiom is also needed to reflect, at the KL, what Axioms 8.1 and 8.2 require at the PML.

Axiom 8.3 *For every* $s_1, s_2 \in \text{TEMP}$, *if* $\text{Mod}_{\text{prop}}(s_1) \cap \text{Mod}_{\text{prop}}(s_2) \neq \{\}$, *then, for all* $t \in \Psi(\mathcal{T})$, *if* $\beta \vdash \mathsf{Holds}(s_1, t)$, *then* $\beta \vdash \neg\mathsf{Holds}(s_2, t)$.

Basically, Axiom 8.3 is a constraint on the domain theory, a principle that the knowledge engineer should adopt in axiomatizing the domain.

Second, we assume that perception modalities are mutually exclusive. That is, the same state cannot be perceived via two distinct perception modalities; each modality presents a unique perspective of the external environment. Note, however, that if two distinct states, s and s', are perceived via two distinct perception modalities, μ_{per} and μ'_{per}, it could still be the case that the fact that s holds entails that s' holds (or vice versa), or that both states entail that some third state, s'', holds. For example, looking at the alarm clock, Cassie may *see* that it is 7 a.m. At the same time, she may also *hear* the alarm, which, given other background knowledge, would entail what she visually perceives. The main point is that, if a state is perceivable, then it is perceivable via one, and only one, modality.[26] Formally, we introduce a *partial* function, Mod_{per}, from TEMP to $\mathcal{M}_{\text{per}}$, such that $[\![s]\!]$ may only be perceived via the modality corresponding to $\text{Mod}_{\text{per}}(s)$.

Axiom 8.4 *For every* $s \in \text{TEMP}$, $\mu \in \mathcal{M}_{\text{per}}$, *and* $i \in \mathbb{N}$, *at* $[\![*\mathsf{NOW}_i]\!]$, *if* $[\![s]\!]$ *is perceived via the modality corresponding to* μ, *then* $\mu = \text{Mod}_{\text{per}}(s)$.

In addition, we assume that bodily states cannot be perceived. That is, the set of bodily states and the set of perceivable states are disjoint.

[26]Note that, ultimately, this might just be a simplifying assumption of our theory.

Algorithm `setup_new_MTF`

1. `move_NOW`
2. For all $\mu \in \mathcal{M}_{\text{prop}}$
   3. If there are s and t such that $*\mu = \mathsf{MHolds}(s, t)$ then $\beta \longleftarrow \beta \cup \{\text{*}\mathsf{NOW} \sqsubset t\}$.
4. For all $\mu \in \mathcal{M}_{\text{per}}$
   5. For all s and t such that $\mathsf{MHolds}(s, t) \in *\mu$
      6. $\beta \longleftarrow \beta \cup \{\text{*}\mathsf{NOW} \sqsubset t\}$.

Figure 10: Algorithm `setup_new_MTF`.

Axiom 8.5 *For every $s \in$ TEMP, if* $\text{Mod}_{\text{prop}}(s) \neq \{\}$*, then* $\text{Mod}_{\text{per}}(s)$ *is undefined.*

Modality variables are set at the PML when bodily or perceived states start or cease to hold. Thus, in the algorithm of Figure 9, if s is one of these states, step (2) is followed by pointing all the appropriate modality variables (those corresponding to modalities occupied by, or perceiving, s) to the new proposition $\mathsf{MHolds}(s, t)$. Constructing the new MTF may thus be characterized by the algorithm in Figure 10. The crucial thing here is that currently holding bodily or perceived states are smoothly incorporated in the new MTF without the need for any reasoning, reflecting Cassie's continuous sense of the states of her body and the perceived environment. Of course, the algorithm in Figure 10 presupposes the correct setting of modality variables. As pointed out above, this should be taken care of by a revised version of algorithm `state_start` (and a similar account for cessation). To proceed and introduce such a revision, however, we need to carefully examine the different issues involved in state transitions.

8.4 Transitions

Before delving into what exactly is involved in a transition from one value of NOW to another (or from one NOW-MTF to the next), we need to be precise and explicit about the technical use of some English expressions that shall recur throughout the rest of the paper. These are expressions that allude to Cassie's knowledge of state transitions over time.

Definition 8.1 *For every $s \in$TEMP and $i \in \mathbb{N}$ ($i > 1$), we say that* **Cassie determines that** s **holds at** $\text{*}\mathsf{NOW}_i$ *if, at* $[\![\text{*}\mathsf{NOW}_i]\!]$*,* $\beta \vdash \mathsf{Holds}(s,\text{*}\mathsf{NOW}_i)$ *and, at*

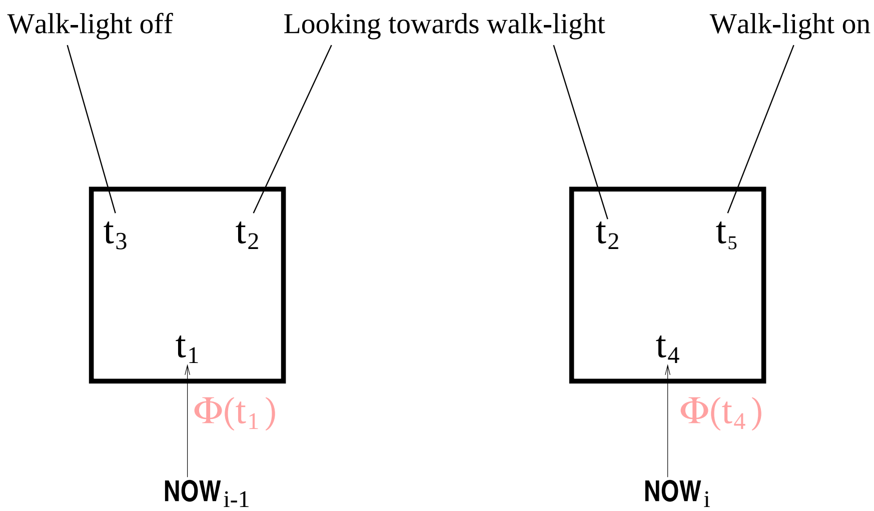

Figure 11: Cassie determines that the walk-light is on at $*\mathsf{NOW}_i$.

$[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$.

Typically, Cassie determines that s holds at $*\mathsf{NOW}_i$ if the transition from $*\mathsf{NOW}_{i-1}$ to $*\mathsf{NOW}_i$ involves Cassie's coming to believe that s holds. For example, suppose that s is the state of the walk-light being on. Further, suppose that the time is $[\![*\mathsf{NOW}_{i-1}]\!]$ and that Cassie sees (and, therefore, believes) that the walk-light is *not* on. In this case, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$.[27] Now, if the walk-light turns on, then NOW moves to $*\mathsf{NOW}_i$ and $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$. This situation can be described by saying that Cassie determines that the walk-light is on at $*\mathsf{NOW}_i$ (see Figure 11).

The reader should be careful about the intuition behind the above definition. In particular, note that the qualifications "at $[\![*\mathsf{NOW}_i]\!]$" and "at $[\![*\mathsf{NOW}_{i-1}]\!]$" are very crucial and cannot be dropped. Consider the situation depicted in Figure 12. At $[\![*\mathsf{NOW}_{i-1}]\!]$ Cassie is *not* looking toward the walk-light and, thus, does not know whether it is on (i.e., at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$). Now, Cassie turns to the walk-light, and two things happen: NOW moves to $*\mathsf{NOW}_i$ (due to the change in Cassie's orientation), and Cassie sees that the walk-light is on. In this case, as in the above one, at $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$, and we can describe this situation by saying that Cassie determines that the walk-light is on at $*\mathsf{NOW}_i$. Unlike the above situation, however, Cassie may later come to believe that the walk-light was actually on at $[\![*\mathsf{NOW}_{i-1}]\!]$, when she was not looking. This does not change anything; it is still the case that Cassie determines that the walk-light is on at $[\![*\mathsf{NOW}_i]\!]$, according

[27] In fact, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \vdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$. We are assuming consistency of β throughout the discussion.

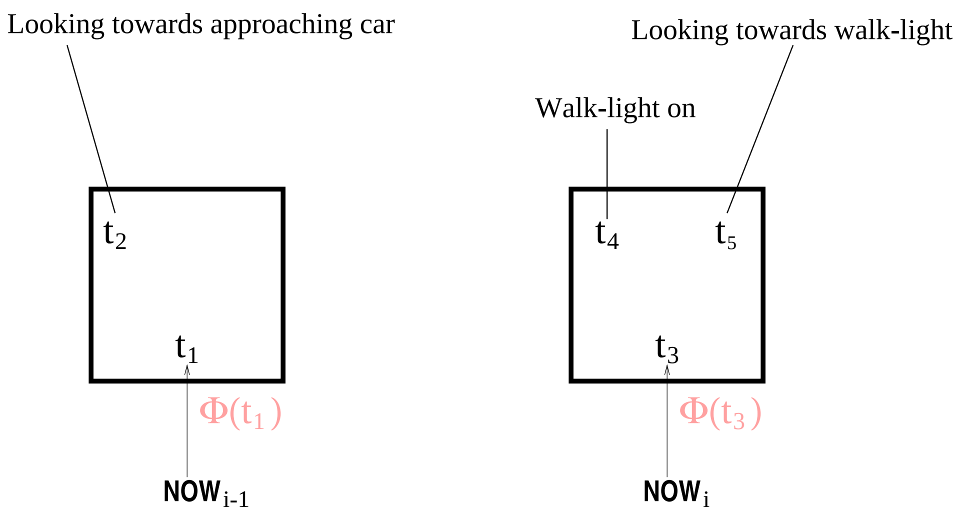

Figure 12: Cassie determines that the walk-light is on at $*\mathsf{NOW}_i$. She does not know whether the walk-light is on at $[\![*\mathsf{NOW}_{i-1}]\!]$

to Definition 8.1. The important thing is not that $\beta \not\vdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$, but that this is the case at $[\![*\mathsf{NOW}_{i-1}]\!]$, the time at which $*\mathsf{NOW}_{i-1}$ is itself $*\mathsf{NOW}$.

To distinguish the above two scenarios, we need a couple of definitions that are more specific than Definition 8.1

Definition 8.2 *For every $s \in$*TEMP *and $i \in \mathbb{N}$ ($i > 1$), we say that* **Cassie determines that s starts to hold at $*\mathsf{NOW}_i$** *if, at $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$ and, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \vdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$.*

The above definition applies to the situation in Figure 11, but not that in Figure 12. The following definition singles out the situation in Figure 12.

Definition 8.3 *For every $s \in$*TEMP *and $i \in \mathbb{N}$ ($i > 1$), we say that* **Cassie determines that s persists at $*\mathsf{NOW}_i$** *if, at $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$ and, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \not\vdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$ and $\beta \not\vdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$.*

The use of "persists" in the above definition will be clear below (Section 10.4). The intuition is that, when one observes a state holding (not starting to hold), one assumes that it *was* holding and will continue to hold for a while. That is, one assumes that the state persists (cf. (McDermott, 1982, pp. 122–123)). Note, however, that the scenario of Figure 12 presents a peculiar instance of Cassie's determining the persistence of a state, namely an instance in which she *starts to perceive* the state.

Definition 8.4 *For every* $s \in$ TEMP *and* $i \in \mathbb{N}$ *(*$i > 1$*), we say that* **Cassie starts to perceive** s **at** $*\mathsf{NOW}_i$ *if*

1. *at* $[\![*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ *is perceived via the modality corresponding to* $\mathrm{Mod}_{\mathrm{per}}(s)$,
2. *at* $[\![*\mathsf{NOW}_{i-1}]\!]$, $[\![s]\!]$ *is not perceived via the modality corresponding to* $\mathrm{Mod}_{\mathrm{per}}(s)$, *and*
3. *at* $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$.

The above definition is required since the case of starting to perceive a state has a particular significance for the state transition algorithms to be presented in Section 8.6. In particular, it involves updating elements of $\mathcal{M}_{\mathrm{per}}$. The first two conditions above capture the transition from not perceiving to perceiving the state s. The third condition merely rules out the possibility of Cassie's determining that s starts through perception, rather than merely starting to perceive it. One important difference between the two cases is that Cassie may indeed believe that s holds over $*\mathsf{NOW}_{i-1}$, for example by being told so. In such a case the transition described by Definition 8.4 is only one of the state of Cassie's perceptual modalities.

Definitions 8.2 and 8.3 describe the different situations in which Cassie determines that a state, s, holds. These may be viewed as involving a transition that Cassie undergoes from one belief state to another. In particular, with respect to s, Cassie undergoes a transition from some sort of a negative belief state to a positive one. More specifically, note that what is common among Definitions 8.1, 8.2, and 8.3, is that Cassie explicitly believes that s holds at the *later* value of NOW. Intuitively, there are similar situations where the opposite is the case.

Definition 8.5 *For every state* $s \in$ TEMP *and for every* $i \in \mathbb{N}$ *(*$i > 1$*), we say that* **Cassie determines that** s **ceases to hold at** $*\mathsf{NOW}_i$ *if, at* $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_i)$ *and, at* $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$.

For example, in the situation represented by Figure 11, Cassie determines that the state of the walk-light's being off ceases at $*\mathsf{NOW}_i$. Similarly, we can give a definition similar to Definition 8.3. This would capture a situation in which Cassie suddenly moves from having an explicit belief that s holds to having no way of telling whether it does. For reasons to be discussed in Section 11, we do not believe that this is a very realistic situation. But even if it is, such a situation does not have any major role to play in our theory and, thus, there is no need to introduce an expression describing it. But a situation similar to that embodied in Definition 8.4 is indeed significant.

Definition 8.6 *For every* $s \in$ TEMP *and* $i \in \mathbb{N}$ *(*$i > 1$*), we say that* **Cassie ceases to perceive** s **at** $*\mathsf{NOW}_i$ *if*

1. *at* $[\![*\mathsf{NOW}_{i-1}]\!]$, $[\![s]\!]$ *is perceived via the modality corresponding to* $\mathrm{Mod}_{\mathrm{per}}(s)$,

2. *at* $[\![*\mathsf{NOW}_{i}]\!]$, $[\![s]\!]$ *is not perceived via the modality corresponding to* $\mathrm{Mod}_{\mathrm{per}}(s)$, *and*

3. *at* $[\![*\mathsf{NOW}_{i}]\!]$, $\beta \not\vdash \neg\mathsf{Holds}(s, *\mathsf{NOW}_i)$.

In the situation depicted in Figure 12, Cassie ceases to perceive that a car is approaching at $*\mathsf{NOW}_i$. Note that this does not mean that Cassie ceases to *believe* that the car is approaching; she only ceases to perceive it, but there might be other reasons for her to believe that it is still there. Indeed, the third condition in Definition 8.6 allows this possibility.

As pointed out in Section 6, Cassie may determine that a state holds through various means: perception, proprioception, direct assertion, or inference. We should now define how the theory reflects each of these modes of determination. First, let us define the following convenient notion of *directly* determining that a state holds.

Definition 8.7 *For every* $s \in$TEMP *and* $i \in \mathbb{N}$ $(i > 1)$, *we say that* **Cassie directly determines that** s **holds at** $*\mathsf{NOW}_i$ *if*

1. *Cassie determines that* s *holds at* $*\mathsf{NOW}_i$, *and*

2. *there is some* $P \subset \mathcal{P}$ *such that*

   (a) *at* $[\![*\mathsf{NOW}_{i-1}]\!]$, *for every* $p \in P$, $\beta \not\vdash p$,

   (b) *at* $[\![*\mathsf{NOW}_{i}]\!]$, $P \subseteq \beta$, *and*

   (c) $P \vdash \mathsf{Holds}(s, *\mathsf{NOW}_i)$.

Thus, Cassie directly determines that s holds if she acquires some *new* pieces of information which, by themselves, allow Cassie to infer that s holds. That is, Cassie does not need to use any background knowledge to determines that s holds; the new information is sufficient. Note that this means that, at a time when β is empty, Cassie can determine that s holds only directly.

Given this definition, we may now be more precise about what it means to determine that a state holds through perception, proprioception, or direct assertion.

Definition 8.8 *For every* $s \in$TEMP *and* $i \in \mathbb{N}$ $(i > 1)$, *we say that* **Cassie determines that** s **holds at** $*\mathsf{NOW}_i$ **through perception** *if*

1. *Cassie directly determines that* s *holds at* $*\mathsf{NOW}_i$ *and*

2. *at* $[\![*\mathsf{NOW}_{i}]\!]$, $[\![s]\!]$ *is perceived via the modality corresponding to* $\mathrm{Mod}_{\mathrm{per}}(s)$.

Definition 8.9 *For every $s \in$*Temp *and $i \in \mathbb{N}$ ($i > 1$), we say that* **Cassie determines that s holds at $*\mathsf{NOW}_i$ through proprioception** *if*

1. *Cassie directly determines that s holds at $*\mathsf{NOW}_i$ and*
2. *at $[\![*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ occupies the modalities corresponding to elements of* $\mathrm{Mod}_{\mathrm{prop}}(s)$.

Definition 8.10 *For every $s \in$*Temp *and $i \in \mathbb{N}$ ($i > 1$), we say that* **Cassie determines that s holds at $*\mathsf{NOW}_i$ through direct assertion** *if*

1. *Cassie directly determines that s holds at $*\mathsf{NOW}_i$,*
2. *at $[\![*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ is not perceived via the modality corresponding to* $\mathrm{Mod}_{\mathrm{per}}(s)$, *and*
3. *at $[\![*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ does not occupy any of the modalities corresponding to elements of* $\mathrm{Mod}_{\mathrm{prop}}(s)$.

Note that, given the above definitions, direct assertion is simply the direct determination of a state holding by means other than perception and proprioception. A similar relation holds between determination through inference and direct determination.

Definition 8.11 *For every $s \in$*Temp *and $i \in \mathbb{N}$ ($i > 1$), we say that* **Cassie determines that s holds at $*\mathsf{NOW}_i$ through inference** *if*

1. *Cassie determines that s holds at $*\mathsf{NOW}_i$ and*
2. *Cassie does not directly determine that s holds at $*\mathsf{NOW}_i$.*

The crucial difference between this and Definition 8.7 is that the new information is not sufficient for Cassie to determine that s holds; she must also make use of *old* background information. Similar to Definitions 8.8 through 8.11, we can provide definitions for determining persistence, onsets, and cessations.

8.5 The Many Scenarios of Change

In this section, we present some insights into the nature of state change. This is done through a careful examination of the different situations in which Cassie determines that a change has taken place. Those insights will be presented as a set of principles amending the two principles of change of Section 7.6 (Axioms 7.3 and 7.4). It should be noted, however, that these principles are not part of the formal theory (and, hence, will not be presented as axioms); they are general guidelines that will be formally crystallized in a number of algorithms to be presented in the next section. The first of these principles follows directly from Axiom 8.1.

The Third Principle of Change. *For every bodily state $[\![s]\!]$, $[\![s]\!]$ ceases to hold when and only when some other bodily state, $[\![s']\!]$, starts to hold.*

The above principle underlines our previous discussion of how it is that there is always some bodily state occupying a given proprioceptual modality. For example, suppose that Cassie is moving, and that this occupies the $\mathsf{LOCOMOTION}_{\mathrm{prop}}$ modality. If Cassie stops moving, the $\mathsf{LOCOMOTION}_{\mathrm{prop}}$ modality gets occupied by the state of her standing still. Similarly, if $\mathsf{VISION}_{\mathrm{prop}}$ is occupied by the state of Cassie's looking toward the walk-light, a cessation of the aforementioned state results in $\mathsf{VISION}_{\mathrm{prop}}$ being occupied by some other state, say Cassie's looking toward an approaching car. The important thing here is that, for bodily states, we have a single unified account for onsets and cessations, since they always occur together. That is, we only need to consider what happens when a bodily state starts to hold; a bodily state's ceasing is only one of the things that happen as another bodily state starts (or the other way around). This shall be reflected in the construction of state transition algorithms.

The Fourth Principle of Change. *For every $s \in \Psi(\textsc{Temp})$ and for every $i \in \mathbb{N}$ ($i > 1$), if Cassie determines that s ceases to hold at $*\mathsf{NOW}_i$, then this happens through inference or direct assertion.*

The basic insight behind the above principle is that one only perceives or feels (through proprioception) the presence of a state, not its absence. Our bodies always provide *positive* information about states, never negative ones. Thus, one does not directly perceive that the walk-light is not on; rather, one perceives that it is off, and *infers* that it is not on. Alternately, one my come to believe that the walk-light is not on if they are explicitly told so. According to Definition 8.5, determining cessation of a state hinges on coming to believe that it does not hold, which, as we noted, cannot happen through perception or proprioception. What this means is that we only need to account for cessation in cases of inference and direct assertion. Again, this will be embodied in state transition algorithms.

The Fifth Principle of Change. *For every $i \in \mathbb{N}$ ($i > 1$), Cassie may determine that more than one state holds at $*\mathsf{NOW}_i$.*

The above principle embodies our discussion in Section 7.6 that, though generally stable, the world might present perceptually-crude agents with multiple simultaneous changes. But even for agents with fine-grained perception of time (humans, for example) the above principle still holds. In particular, it is often characteristic of the patterns of change exhibited by bodily states. As per the third principle of change, the cessation of a bodily state s is (simultaneous to) the onset of some other bodily state s'. If s occupies more than one modality, then its cessation would typically correspond to the onset of multiple states, each occupying one of those modalities. That is, as a bodily state ceases, multiple states rush in, occupying the modalities just made available. For example, let s be the state of Cassie's searching for a block. In a possible implementation, we have:

$$*\mathsf{VISION}_{\text{prop}} = *\mathsf{LOCOMOTION}_{\text{prop}} = \text{“I am searching for a block”}$$

As s ceases, we get the following situation.

$$*\mathsf{VISION}_{\text{prop}} = \text{“I am looking at a block”}$$
$$*\mathsf{LOCOMOTION}_{\text{prop}} = \text{“I am standing still”}$$

These two states (looking at a block and standing still) would virtually start simultaneously as the searching activity comes to an end.

The Sixth Principle of Change. *For every bodily state s, and for every $i \in \mathbb{N}$ ($i > 1$), if Cassie determines that s holds at $*\mathsf{NOW}_i$ through proprioception, then Cassie determines that s starts to hold at $*\mathsf{NOW}_i$.*

The above principle simply states that the only possible situation in which Cassie determines that a state holds through proprioception is when she determines that it *starts* to hold. In particular, Cassie cannot proprioceptually determine that a state persists (as per Definition 8.3). This is reasonable since Cassie does not, all of a sudden, sense a bodily state holding; she has to, first, sense it starting to hold, given the continuous proprioceptual feedback. Thus, determining that a state persists is exclusive to non-bodily states—those that can be *discovered* in the midst of holding. In fact, the above principle embodies a major property that distinguishes bodily states from non-bodily states.[28]

But note the wording of the above principle. We were careful enough to require that the principle applies only if a bodily state is determined to be holding through *proprioception*. Is there another way by which an agent may determine that a state of its own body holds? Theoretically, the agent may be *told* anything, including information about its bodily states. For humans, such information would only be accepted if it conforms with proprioceptual information (at least in normal situations). For robots, however, we believe that this is an implementation decision. One may design a robot that accepts information about its own body from outsiders. Alternatively, the robot may simply reject any such information or choose some

[28]Of course, humans sometimes violate this principle—for example, when they lose consciousness and get back to their senses to find certain bodily states persisting. But perhaps something similar to the above principle is why such unfortunate incidents are very perplexing.

intermediate route, following the lead of humans. We shall discuss what such design decisions may entail in Section 12.2. For now, however, we will adopt the more lenient stance; we shall develop the *theory* so that Cassie is allowed to believe assertions about her bodily states holding, even if she does not *feel* those states. Note that, given Axiom 8.3, this would typically involve belief revision that favors information coming from one source over another (Johnson and Shapiro, 2000a). A particular implementation may choose to disable this mechanism (we will show below how, in our theory, this could be done).

The Seventh Principle of Change. *For every $s \in \Psi(\textsc{Temp})$ and for every $i \in \mathbb{N}$ $(i > 1)$, if Cassie determines that s persists at $*\mathsf{NOW}_i$, then there is some $s' \in \Psi(\textsc{Temp})$ such that Cassie determines that s' starts, or ceases, at $*\mathsf{NOW}_i$.*

The gist of the above statement is that determining persistence cannot, by itself, be responsible for the movement of NOW. We have informally stated that NOW moves only if Cassie determines that a state starts or ceases to hold. But the second principle of change (Axiom 7.4) makes a weaker claim: two consecutive NOW-MTFs are not epistemically-equivalent. Theoretically, this may be satisfied in a situation where Cassie determines that some state persists (not starts or ceases). Practically, however, this cannot be the case. In particular, consider what is involved in determining that a state persists. First, there are only three means by which Cassie may determine that a state, s, persists: inference, perception, and direct assertion (proprioception being ruled out by the sixth principle of change). Now, let us take a careful look at each of these.

- **Inference.** Suppose that Cassie determines that s persists at $*\mathsf{NOW}_i$ through inference. According to Definition 8.3, this means that, at $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$ and, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$ and $\beta \nvdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$. Note that, not only do we require that, at $[\![*\mathsf{NOW}_{i-1}]\!]$, Cassie does not have any explicit beliefs about whether s holds, but that she *cannot* infer whether it does, given everything that she knows. How, then, can such an inference be possible at $*\mathsf{NOW}_i$? Something must have changed. In particular, some set of propositions, P, must have been added to β (assuming monotonicity) so that what was not possible at $*\mathsf{NOW}_{i-1}$ is possible at $*\mathsf{NOW}_i$. Naturally, P must have been added to β by some means other than inference. For if it was merely inferred, then it cannot be the case that at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$. Therefore, P must originate from perception, proprioception, or direct assertion. Now, if P includes a proposition of the form $\mathsf{Holds}(s',*\mathsf{NOW}_i)$, for some $s' \in \Psi(\textsc{Temp})$, then Cassie determines that s' holds at $*\mathsf{NOW}_i$. If this is a determination of s' starting to hold, then the seventh principle of change is satisfied; if it is a determination of the mere persistence of s', then it must have been achieved through perception or direct assertion—the two cases we review below. If, on the other hand, P does not include any propositions of the form $\mathsf{Holds}(s', {*}\mathsf{NOW}_i)$, then elements of P can only get to β through direct assertion (since one can only perceive or feel what is present). Since such propositions cannot be reporting a (current) change in the environment, we assume that their assertion is not responsible for the movement of NOW (more on this below).

- **Perception.** Suppose that Cassie determines that s persists at ${*}\mathsf{NOW}_i$ through perception. As per Definition 8.3, at ${*}\mathsf{NOW}_{i-1}$, Cassie does not have any beliefs (implicit or explicit) about whether s holds, whereas, at ${*}\mathsf{NOW}_i$, she can actually perceive it. But, for this to happen, something else must have changed. In particular, there are two general scenarios. The first involves a change in one or more of Cassie's bodily states. In this scenario, Cassie changes the location, or direction, of her sensors in such a way that they come to have access to whichever environmental phenomena constitute the state s. In the situation of Figure 12, for instance, not only does Cassie perceive that the walk-light is on, but, crucially, she turns her head toward the walk-light. If s is the state of the walk-light's being on, then, in that situation, Cassie determines that s persists at ${*}\mathsf{NOW}_i$ through (visual) perception. In addition, she determines that s' starts at ${*}\mathsf{NOW}_i$, where s' is the bodily state of Cassie's looking toward the walk-light. In fact, this change in s' is causally-responsible for Cassie's determining that s persists. It is hard to come up with similar examples for auditory perception. One possibility, however, is Cassie's coming to perceive a certain sound as a result of her *moving* closer to its source. The second scenario involves the removal of a barrier, thereby allowing Cassie's sensors to have access to the portion of the environment including what Cassie can recognize as the state s holding. An example is when Cassie opens the door of a room to see the room lights on, or to hear the radio therein playing. Note that, in this scenario, determining that a state s persists through perception is accompanied by (indeed, causally-dependent on) determining, again through perception, that some state s' starts (or ceases, depending on the agent's perspective).[29] The above discussion indicates that, in general, when Cassie determines that a state starts (or ceases), it may also be the case that she determines that one or more states persist. In particular, this happens in two cases corresponding to the two scenarios discussed above. First, the state change determined by Cassie involves the removal of a barrier to perception. Second, the state change is that of a bodily state that uses a modality capable of conveying perceptual

[29]In addition to the two scenarios dicussed here, there is actually a third one whereby Cassie determines that a state persists through perception. This is when Cassie starts to *attend* to a particular state of the environment that is already accessible by her sensors. For example, Cassie may be looking towards the walk-light, with no barriers in between, but only realizes that the walk-light is on when she starts paying attention to it. Being attentive to a particular state is itself a temporary state and it is a change of that state that allows Cassie to determine the persistence of other states. Since we have not formalized this notion of attending in our theory, we do not discuss this issue any further.

information. Technically, this is a modality to which there are corresponding variables in both $\mathcal{M}_{\text{prop}}$ and $\mathcal{M}_{\text{per}}$.

- **Direct assertion.** In the cases of inference and perception, we have argued, from an empirical point of view, for the validity of the seventh principle of change. As the reader must have noted, the discussion of these two cases, in many respects, resembled a proof of the statement of the seventh principle of change. The same could be done here; we can argue that, empirically, the seventh principle of change is a plausible statement for the case of direct assertion of persistence. However, the situation is less idealistic than in the cases of inference and perception; our *theory* of direct assertion is still not developed enough to reflect empirical intuitions.[30] In particular, our account of direct assertion treats it as an inspiration-like phenomenon, a mysterious activity resulting in new beliefs appearing, out of the blue, in Cassie's mind. Actually, this is fine and is indeed the traditional model of adding information to a knowledge base. However, empirically, this is not how it exactly works. An assertion does not make its way magically into one's head; rather, there is an assertion *event* that results in one's acquiring a new belief. Forgetting aside, an agent is aware of such an assertion event, when it happened, and who its agent was. Our theory, however, does not include an account of these assertion events.[31] That is, when told something, Cassie does not have any beliefs about the state changes associated with that *telling*, she is only aware of the asserted proposition, not of the assertion event proper. To see why this is problematic for a proof-like justification of the seventh principle of change, consider the following example. Suppose Cassie determines, through direct assertion, that the walk-light is on. In other words, someone, Stu for example, says: "Cassie, the walk-light is now on". Assuming that nothing else changes in the environment, one thing that has certainly changed is that a new sentence has been uttered (or typed, if you will). More precisely, a state of Stu's uttering something starts, persists for a while, and then ceases, resulting in Cassie's determining the persistence of the walk-light's on-state. Thus, the seventh principle of change readily holds; for Cassie's determining that s (the walk-light's being on) persists is accompanied by (and, again, causally-dependent on) the cessation of some s' (Stu's uttering something). However, since our theory does not account for assertion events, the above scenario cannot be used to empirically justify the seventh principle of change for the case of direct assertion of persistence. Rather, the principle has to be interpreted as a constraint that we impose, by fiat, on our system. The constraint simply entails that asserting persistence does not, in itself, consti-

[30] The reader should, nevertheless, note that linguistic forms of communication with Cassie, particularly in natural language, have been thoroughly studied in the past (Shapiro, 1989; Ali and Shapiro, 1993; Shapiro and Rapaport, 1995; Shapiro, 1996; Shapiro, 2000, for instance).

[31] Though, currently, the point of making Cassie aware of the *source* of each assertion is under investigation (Johnson and Shapiro, 2000a).

tute a change. But note what this means. Suppose that Cassie determines that s persists at $*\mathsf{NOW}_i$ through direct assertion. Further, suppose that $*\mathsf{NOW}_i$ is the value of NOW between the *real* clock-times 12 p.m. and 12:10 p.m., and that the proposition $\mathsf{Holds}(s,*\mathsf{NOW})$ is asserted at 12:05 p.m. To Cassie, $[\![s]\!]$ holds throughout $[\![*\mathsf{NOW}_i]\!]$, but, to us, this might not be the case; $[\![s]\!]$ may have started to hold at 12:04 p.m. To some, this may look like a major problem; for, apparently, Cassie simply has a false belief. Nevertheless, we do not think that there are any major problems, and, even if there is some problem, nothing major seems at stake. To Cassie, the only reason why s's not holding between 12 and 12:04 would matter is if clock times have any significance to her. If they do not, then there is no problem. If they do, then she would have to be aware of them, and NOW will necessarily move with every tick of the clock. In that case, however, the problem disappears, since $*\mathsf{NOW}_i$ cannot extend over a period in which there are salient changes. Note that this follows from the first principle of change (Axiom 7.3). In general, *any* assertion that does not result in Cassie's determining that a state starts (or ceases) also does not result in NOW moving. Note that, were we to actually provide an account of assertion events, direct assertion would cease to stand out as an independent method by which Cassie may acquire new beliefs. In particular, the assertion event itself would be *perceived* (through audition, for example) just as any other external event. What would be needed then is an axiom to the effect that whenever an assertion event takes place, Cassie should believe the content of the assertion (more on this in Section 12.2).

Similar to the seventh principle of change, we have the following principle.

The Eighth Principle of Change. *For every $s \in \Psi(\textsc{Temp})$ and for every $i \in \mathbb{N}$ ($i > 1$), if Cassie starts, or ceases, to perceive s at $*\mathsf{NOW}_i$, then there is some $s' \in \Psi(\textsc{Temp})$ such that Cassie determines that s' starts, or ceases, to hold at $*\mathsf{NOW}_i$.*

Justifying this principle follows the same line of thought introduced above for the perception case of the seventh principle of change. Cassie starts to perceive s due to the removal of a barrier or the repositioning, or re-orienting, of her sensors. Similarly, if Cassie ceases to perceive s, then this must be the result of either the introduction of a barrier or of Cassie's changing the position or orientation of her sensors.

Whether the seventh and eighth principles of change are empirically-justified or theoretically-motivated, we uphold them as guiding principles: mere determination of persistence, or change in perception, does not move NOW. This being said, let us now turn to what exactly *is* involved in the movement of NOW.

8.6 Proceduralizing Belief Update

With the principles discussed above in mind, algorthim `state_change` in Figure 13 presents a general outline of the circumstances surrounding the movement of NOW. Of course, every step of this algorithm needs to be carefully explained. First, however, we need to be explicit about *when* the algorithm is executed and what the arguments, $S^{\uparrow}$ and $S^{\downarrow}$, signify. Algorithm `state_change` gets executed whenever Cassie's body and/or sensors detect a change regarding some state (actually, a set thereof). The types of *change* referred to here will be made explicit. But note that we have excluded information about states reported by direct assertion. This is because of the peculiar features of direct assertion discussed in Section 8.5. We choose to have a separate set of algorithms for handling direct assertion; these will be presented below. Thus, algorithm `state_change` is only concerned with changes in perception or proprioception. The types of change are reflected in the nature of the arguments, $S^{\uparrow}$ and $S^{\downarrow}$, which is totally revealed by the preconditions. Both arguments are *sets* of states. Members of $S^{\uparrow}$ are states that *start* to be detected as holding by Cassie's perceptual and proprioceptual systems. More precisely, these are states that either start to hold or just start to be perceived. This is reflected by the first and second pre-conditions. The first precondition states that, just prior to executing the algorithm, modality variables do not reflect anything about members of $S^{\uparrow}$. The second pre-condition, on the other hand, states that the states denoted by members of $S^{\uparrow}$ actually occupy proprioceptual modalities or are perceived via perceptual modalities. Part of the function of the algorithm is to update modality variables to reflect the new situation (see Theorem 8.1 below). It should be noted that, given the fourth, seventh, and eighth principles of change, there must be some $s \in S^{\uparrow}$ such that, prior to executing the algorithm, $\beta \vdash \neg\mathsf{Holds}(s,{*}\mathsf{NOW})$. That is, some state in $S^{\uparrow}$ must actually be starting to hold. Members of $S^{\downarrow}$ are states that *cease* to be detected as holding by Cassie's perceptual and proprioceptual systems. This may happen if a state ceases or just ceases to be perceived. Again, this is reflected by the third and fourth pre-conditions. We shall, henceforth, assume that, whenever there are states satisfying its pre-conditions, algorithm `state_change` gets executed.

Let us now take a careful look at each step of the algorithm. Step 1 initializes a variable, P_{new}, that is used to collect newly introduced propositions about states holding (see step 4d). Step 3 checks if any of the members of $S^{\uparrow}$ is a state that is already believed to hold. If this is the case, then such a state must have made it into $S^{\uparrow}$ because it started to be perceived or proprioceived.[32] This initiates algorithm `start_ceive`, shown in Figure 14.[33] The first step of `start_ceive` initi-

[32] Recall our discussion in Section 8.5 regarding how Cassie may have a belief about a bodily state that she does not feel.

[33] Here, we are alluding to the term "ception" used in (Talmy, 2000). Though Talmy uses the term to cover perception, proprioception, and conception; we only use it to cover the first two.

Pre-Conditions:

1. *For every* $s \in S^{\uparrow}$*, there is no* $t \in \Psi(\mathcal{T})$ *such that, for some* $\mu \in \mathcal{M}_{\text{prop}}$*,* $*\mu = \mathsf{MHolds}(s,t)$ *or, for some* $\mu \in \mathcal{M}_{\text{per}}$*,* $\mathsf{MHolds}(s,t) \in *\mu$*.*

2. *For every* $s \in S^{\uparrow}$*, there is some* $\mu \in \mathcal{M}_{\text{prop}}$ *such that* $[\![s]\!]$ *occupies the modality corresponding to* μ*, or there is some* $\mu \in \mathcal{M}_{\text{per}}$ *such that* $[\![s]\!]$ *is perceived via the modality corresponding to* μ*.*

3. *For every* $s \in S^{\downarrow}$*, there is some* $t \in \Psi(\mathcal{T})$ *such that either, for some* $\mu \in \mathcal{M}_{\text{prop}}$*,* $*\mu = \mathsf{MHolds}(s,t)$ *or, for some* $\mu \in \mathcal{M}_{\text{per}}$*,* $\mathsf{MHolds}(s,t) \in *\mu$*.*

4. *For every* $s \in S^{\downarrow}$*, there is no* $\mu \in \mathcal{M}_{\text{prop}}$ *such that* $[\![s]\!]$ *occupies the modality corresponding to* μ *and there is no* $\mu \in \mathcal{M}_{\text{per}}$ *such that* $[\![s]\!]$ *is perceived via the modality corresponding to* μ*.*

Algorithm `state_change`$(S^{\uparrow} \subseteq \text{Temp}, S^{\downarrow} \subseteq \text{Temp})$

1. $P_{\text{new}} \longleftarrow \{\}$.
2. For all $s_i \in S^{\uparrow}$
   3. If $\beta \vdash \mathsf{Holds}(s_i, \text{*}\mathsf{NOW})$ then `start_ceive`(s_i, t_i),
      where t_i is the situation interval associated with s_i such that $\beta \vdash \text{*}\mathsf{NOW} \sqsubset t_i$
   4. else
      4a. Pick some $t_i \in \mathcal{T}$, such that $t_i \notin \Psi(\mathcal{T})$.
      4b. If $\beta \vdash \neg\mathsf{Holds}(s_i, \text{*}\mathsf{NOW})$ then `state_start`(s_i, t_i).
      4c. Else, `state_persist`(s_i, t_i).
      4d. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{Mholds}(s_i, t_i)\}$.
5. For all $s_i \in S^{\downarrow}$, `cease_perceive`(s_i).
6. `setup_new_MTF`.
7. `Forward`(P_{new}).

Figure 13: Algorithm `state_change`.

ates algorithm `start_proprioceive` (see Figure 15) which updates proprioception modality variables. Note that the algorithm does not check if its argument, s, actually occupies any modalities; it merely checks if it is a bodily state. The reason is that, given pre-condition 2, $[\![s]\!]$ is either perceived or proprioceived. By Axiom 8.5, a bodily state cannot be perceived. Therefore, if s is a bodily state, then it must be proprioceived. If the state is *not* a bodily state, then it must be perceived via some perception modality. Updating the appropriate perception modality variable is the function of algorithm `start_perceive` of Figure 16. Note that, were we to opt for an implementation in which Cassie may *not* hold beliefs about bodily states that she does not feel, step 3 of algorithm `state_change` should initiate algorithm `start_perceive` directly.

It should also be noted that, unlike perception modality variables, the contents of proprioception modality variables are *overwritten.* Thus, the state previously occupying a proprioceptual modality no longer have a proposition in the corresponding modality variable (see the third principle of change in Section 8.5). This means that algorithm `setup_new_MTF` (Figure 10) would not include such a state in the new MTF. In fact, given Axiom 8.1 and the pre-conditions, all the states displaced by the newly starting bodily states are the bodily states in $S^{\downarrow}$.

If, in step 3 of `state_change`, the state is not already believed to be holding, then there are two possibilities: (i) the state has just started to hold or (ii) it has just started to be perceived, with Cassie having no beliefs about it. In both cases, a new situation interval needs to be associated with the state. This is achieved by step 4a. If the state is believed to be *not* holding, then it must have just started; otherwise, it has just been perceived to persist. In the first case, algorithm `state_start` gets initiated. The algorithm is shown in Figure 17. The only difference between this and algorithm `start_ceive` is that it, first, adds a new belief about the state holding to the belief space.

If an element of $S^{\uparrow}$ has just started to be perceived with Cassie having no beliefs about whether it holds, then algorithm `state_persist` gets initiated by step 4c of `state_change`. This algorithm is shown in Figure 18. It is very similar to algorithm `start_perceive` with the extra step of introducing a new belief into Cassie's belief space. Note that, given the sixth principle of change, the algorithm assumes that its state-argument does not occupy any proprioceptual modalities.

Step 5 of `state_change` processes members of the set $S^{\downarrow}$. Regardless of whether some $s_i \in S^{\downarrow}$ has actually ceased or simply ceased to be perceived, perception modality variables must be updated to reflect the lack of perceptual information about s_i (as indicated by pre-condition 4). This is the function of algorithm `cease_perceive` of Figure 19.

Step 6 initiates algorithm `setup_new_MTF` of Figure 10 which is responsible

Algorithm `start_ceive`(s, t)

1. `start_proprioceive`(s, t).
2. `start_perceive`(s, t).

Figure 14: Algorithm `start_ceive`.

Algorithm `start_proprioceive`(s, t)

1. If $\text{Mod}_{\text{prop}}(s) \neq \{\}$, then
   for all $\mu \in \text{Mod}_{\text{prop}}(s)$, $\mu \longleftarrow \mathsf{MHolds}(s, t)$.

Figure 15: Algorithm `start_proprioceive`.

Algorithm `start_perceive`(s, t)

1. If $\text{Mod}_{\text{per}}(s)$ is defined, then
   $\text{Mod}_{\text{per}}(s) \longleftarrow *\text{Mod}_{\text{per}}(s) \cup \{\mathsf{MHolds}(s, t)\}$.

Figure 16: Algorithm `start_perceive`.

Algorithm `state_start`(s, t)

1. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s, t)\}$.
2. `start_ceive`(s, t).

Figure 17: Algorithm `state_start`.

Algorithm `state_persist`(s, t)

1. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s, t)\}$.
2. `start_perceive`(s, t).

Figure 18: Algorithm `state_persist`.

Algorithm `cease_perceive`(s)

1. If $\mathrm{Mod}_{\mathrm{per}}(s)$ is defined, then
   $\mathrm{Mod}_{\mathrm{per}}(s) \longleftarrow *\mathrm{Mod}_{\mathrm{per}}(s) \setminus \{\mathsf{MHolds}(s,t)\}$

Figure 19: Algorithm `cease_perceive`.

for the actual movement of NOW and the construction of the new NOW-MTF.[34] The final step of the algorithm initiates forward inference on all of the new assertions of states holding. See Section 6 for the reasoning behind this step.

In addition to its pre-conditions, algorithm `state_change` has a number of significant post-conditions.

Theorem 8.1 *For every $s \in$ TEMP and $i \in \mathbb{N}$ $(i > 1)$, if, at $[\![*\mathsf{NOW}_{i-1}]\!]$, algorithm* `state_change` *gets initiated with $s \in S^{\uparrow}$, then, following the execution of the algorithm, there is some $t \in \Psi(\mathcal{T})$ such that $\beta \vdash \mathsf{MHolds}(s,t)$, and either, for some $\mu \in \mathcal{M}_{\mathrm{prop}}$, $*\mu = \mathsf{MHolds}(s,t)$ or, for some $\mu \in \mathcal{M}_{\mathrm{per}}$, $\mathsf{MHolds}(s,t) \in *\mu$.*

Proof. Let s be an arbitrary member of $S^{\uparrow}$. At $[\![*\mathsf{NOW}_{i-1}]\!]$, it is either the case that (i) $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$, (ii) $\beta \vdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$, or (iii) $\beta \not\vdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$ and $\beta \not\vdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$. Consider each case.

(i) Suppose that, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$. By step 3 and Axiom 7.1, there is some $t \in \Psi(\mathcal{T})$ such that $\beta \vdash \mathsf{MHolds}(s,t)$. This proves the first conjunct in the consequent of the statement of the theorem. Step 3 initiates algorithm `start_ceive` with arguments s and t. By pre-condition 2, either there is some $\mu \in \mathcal{M}_{\mathrm{prop}}$ such that $[\![s]\!]$ occupies the modality corresponding to μ, or there is some $\mu \in \mathcal{M}_{\mathrm{per}}$ such that $[\![s]\!]$ is perceived via the modality corresponding to μ. In the first case, following the execution of algorithm `start_proprioceive`, there is some $\mu \in \mathcal{M}_{\mathrm{prop}}$ such that $*\mu = \mathsf{MHolds}(s,t)$. By Axiom 8.1, there is no $s' \in S^{\uparrow}$ $(s' \neq s)$ occupying the modality corresponding to μ. Therefore, subsequent initiations of `start_proprioceive` through the end of `state_change` do not change the value of μ. In the second case, following the execution of algorithm `start_perceive`, there is some $\mu \in \mathcal{M}_{\mathrm{per}}$ (namely $\mathrm{Mod}_{\mathrm{per}}$, as per Axiom 8.4) such that $\mathsf{MHolds}(s,t) \in *\mu$. The only place in `state_change` where the proposition $\mathsf{MHolds}(s,t)$ may be removed from $*\mu$ is algorithm `cease_perceive` in step 5. But, given the pre-conditions, $S^{\uparrow}$ and $S^{\downarrow}$ are disjoint, and algorithm `cease_perceive` never gets applied to s. Therefore $\mathsf{MHolds}(s,t)$ continues to be a member of $*\mu$ through the end of `state_change`.

(ii) Suppose that, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \vdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$. By step 4b, algorithm

[34] In Section 11, we will introduce a revised version of algorithm `setup_new_MTF` which guarantees that, if any of the members of $S^{\downarrow}$ has actually ceased, then Cassie's beliefs would reflect this fact.

state_start gets initiated with s and t as arguments, where t is the new interval introduced in step 4a. By step 1 of state_start, $\beta \vdash \mathsf{MHolds}(s, t)$, which proves the first conjunct of the consequent of the theorem. The proof of the second conjunct follows that of part (i) above.

(iii) Suppose that, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$ and $\beta \nvdash \neg\mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$. By step 4c, algorithm state_persist gets initiated with s and t as arguments, where t is the new interval introduced in step 4a. By step 1 of state_persist, $\beta \vdash \mathsf{MHolds}(s, t)$, which proves the first conjunct of the consequent of the theorem. Step 2 of algorithm state_persist initiates algorithm start_perceive with s and t as arguments. By the sixth principle of change, $[\![s]\!]$ cannot be occupying a proprioceptual modality. Therefore, given pre-condition 2, there is some $\mu \in \mathcal{M}_{\text{per}}$ such that $[\![s]\!]$ is perceived via the modality corresponding to μ. Thus, following the execution of start_perceive, there is some $\mu \in \mathcal{M}_{\text{per}}$ such that $\mathsf{MHolds}(s, t) \in *\mu$. Following the proof of part (i), $\mathsf{MHolds}(s, t)$ continues to be a member of $*\mu$ through the end of state_change. □

A corresponding result holds for members of $S^{\downarrow}$.

Theorem 8.2 *For every $s \in$*Temp *and $i \in \mathbb{N}$ ($i > 1$), if, at $[\![*\mathsf{NOW}_{i-1}]\!]$, algorithm* state_change *gets initiated with $s \in S^{\downarrow}$, then, following the execution of the algorithm, there is no $t \in \mathcal{T}$ such that for some $\mu \in \mathcal{M}_{\text{prop}}$, $*\mu = \mathsf{MHolds}(s, t)$ or, for some $\mu \in \mathcal{M}_{\text{per}}$, $\mathsf{MHolds}(s, t) \in *\mu$.*

Proof. Let s be an arbitrary member of $S^{\downarrow}$. Given pre-condition 3, suppose that there is some $t \in \mathcal{T}$ and $\mu \in \mathcal{M}_{\text{prop}}$ such that $*\mu = \mathsf{MHolds}(s, t)$. By pre-condition 4 and the third principle of change, there is some $s' \in S^{\uparrow}$ that occupies the modality corresponding to μ. By the proof of Theorem 8.1, a proposition $\mathsf{MHolds}(s', t')$ (for some $t' \in \Psi(\mathcal{T})$) overwrites the contents of μ. On the other hand, suppose that there is some $t \in \mathcal{T}$ and $\mu \in \mathcal{M}_{\text{per}}$ such that $\mathsf{MHolds}(s, t) \in *\mu$. By Axiom 8.4, $\mu = \text{Mod}_{\text{per}}(s)$. Following the execution of algorthim cease_perceive in step 4 with s as an argument, the proposition $\mathsf{MHolds}(s, t)$ gets removed from $*\text{Mod}_{\text{per}}(s)$. Since $S^{\downarrow}$ and $S^{\uparrow}$ are disjoint (given the pre-conditions), then the result of Theorem 8.1 does not apply to any of the members of $S^{\downarrow}$. Therefore, following the execution of state_change there is no $t \in \mathcal{T}$ such that, following the execution of the algorithm, for some $\mu \in \mathcal{M}_{\text{prop}}$, $*\mu = \mathsf{MHolds}(s, t)$ or, for some $\mu \in \mathcal{M}_{\text{per}}$, $\mathsf{MHolds}(s, t) \in *\mu$. □

Theorem 8.3 *For every $s \in$*Temp *and $i \in \mathbb{N}$ ($i > 1$), if, at $[\![*\mathsf{NOW}_{i-1}]\!]$, algorithm* state_change *gets initiated with $s \in S^{\uparrow}$, then, following the execution of the algorithm, $\beta \vdash \mathsf{Holds}(s, *\mathsf{NOW}_i)$.*

Proof. Let s be an arbitrary member of $S^{\uparrow}$. Given the proof of Theorem 8.1, there is some $t \in \mathcal{T}$ such that, just before executing step 6 of algorithm state_change, for some $\mu \in \mathcal{M}_{\text{prop}}$, $*\mu = \mathsf{MHolds}(s,t)$ or, for some $\mu \in \mathcal{M}_{\text{per}}$, $\mathsf{MHolds}(s,t) \in *\mu$. Therefore, by executing algorithm **`setup_new_MTF`**, NOW moves from $*\mathsf{NOW}_{i-1}$ to $*\mathsf{NOW}_i$ and either step 3 or step 6 adds the proposition $*\mathsf{NOW}_i \sqsubset t$ to β. Since, by Theorem 8.1, $\beta \vdash \mathsf{MHolds}(s,t)$, then, given **AS2** and **AS3**, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$. Since s is arbitrary, then the result follows for all $s \in S^{\uparrow}$. $\square$

Given the above theorem, two corollaries readily follow.

Corollary 8.1 *For every $s \in \Psi(\text{TEMP})$ and $i \in \mathbb{N}$ ($i > 1$), if, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \vdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$ and algorithm* **`state_change`** *gets initiated with $s \in S^{\uparrow}$, then Cassie determines that s starts to hold at $*\mathsf{NOW}_i$.*

Proof. Follows directly from Definition 8.2 and Theorem 8.3. $\square$

Corollary 8.2 *For every $s \in$TEMP and $i \in \mathbb{N}$ ($i > 1$), if, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \nvdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$, $\beta \nvdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$, and algorithm* **`state_change`** *gets initiated with $s \in S^{\uparrow}$, then Cassie determines that s persists at $*\mathsf{NOW}_i$.*

Proof. Follows directly from Definition 8.3 and Theorem 8.3. $\square$

The above corollaries show that, as a side-effect, the execution of algorithm **`state_change`** results in Cassie's determining that states start or persist. In what follows, we shall make the following complementary assumption: if, through perception or proprioception, Cassie determines that a state starts or persists, then this may only be an effect of executing algorithm **`state_change`**. In other words, algorithm **`state_change`** is the *only* component of the system responsible for Cassie's determining the start or persistence of states through perception or proprioception.

As pointed out above, separate algorithms are responsible for the direct assertion of state change. Direct assertion of persistence is the responsibility of algorithm **`assert_persist`** of Figure 20. Asserting onsets and cessations are done through the two algorithms shown in Figures 21 and 22, respectively. The algorithms should be self-evident. The only thing to note is that algorithms **`assert_persist`** and **`assert_start`** are quite similar. The crucial differences are (i) the pre-conditions of each algorthim and (ii) the fact that executing **`assert_start`** results in NOW moving while executing **`assert_persist`** does not. Also note step 3 of algorithm **`assert_cease`** where the ceasing states are asserted to have moved into the past. If, in a particular implementation, it is decided that Cassie should filter out any direct assertions about states of her body, then the above algorithms constitute the locations of such filters.

Given these algorithms we can easily prove the following results.

PRE-CONDITIONS:

1. *If* $*\mathsf{NOW} = *\mathsf{NOW}_i$ *for some* $i > 1$, *then, for every* $s \in S$, $\beta \not\vdash \mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$.

2. *If* $*\mathsf{NOW} = *\mathsf{NOW}_i$ *for some* $i > 1$, *then, for every* $s \in S$, $\beta \not\vdash \neg\mathsf{Holds}(s, *\mathsf{NOW}_{i-1})$.

Algorithm `assert_persist`($S \subseteq$ TEMP)

1. $P_{\text{new}} \longleftarrow \{\}$.
2. For all $s_i \in S$
   3. Pick some $t_i \in \mathcal{T}$, such that $t_i \notin \Psi(\mathcal{T})$.
   4. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s_i, t_i), *\mathsf{NOW} \sqsubset t_i\}$.
   5. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{MHolds}(s_i, t_i)\}$.
6. `Forward`(P_{new}).

Figure 20: Algorithm `assert_persist`.

Theorem 8.4 *For every* $s \in$TEMP *and* $i \in \mathbb{N}$ $(i > 1)$, *if algorithm* `assert_persist` *is initiated at* $[\![*\mathsf{NOW}_i]\!]$ *with* $s \in S$, *then Cassie determines that* s *persists at* $*\mathsf{NOW}_i$.

Proof. For every $s \in S$, by step 4 of `assert_persist`, **AS2**, and **AS3**, $\beta \vdash \mathsf{Holds}(s, *\mathsf{NOW}_i)$ at $[\![*\mathsf{NOW}_i]\!]$, following the execution of the algorithm. Given the pre-conditions of `assert_persist` and Definition 8.3, Cassie determines that s persists at $*\mathsf{NOW}_i$. □

Theorem 8.5 *For every* $s \in$TEMP *and* $i \in \mathbb{N}$ $(i > 1)$, *if algorithm* `assert_start` *is initiated at* $[\![*\mathsf{NOW}_{i-1}]\!]$ *with* $s \in S$, *then Cassie determines that* s *starts to hold at* $*\mathsf{NOW}_i$.

Proof. Given the pre-condition of `assert_start` and Definition 8.2, the proof follows that of Theorem 8.4. □

Theorem 8.6 *For every* $s \in$TEMP *and* $i \in \mathbb{N}$ $(i > 1)$, *if algorithm* `assert_cease` *is initiated at* $[\![*\mathsf{NOW}_{i-1}]\!]$ *with* $s \in S$, *then Cassie determines that* s *ceases to hold at* $*\mathsf{NOW}_i$.

Proof. For every $s \in S$, by step 3 of `assert_cease`, $\beta \vdash \neg\mathsf{Holds}(s, *\mathsf{NOW}_i)$ at $[\![*\mathsf{NOW}_i]\!]$, following the execution of the algorithm. Given the pre-condition

PRE-CONDITIONS:

1. *For every* $s \in S$, $\beta \vdash \neg\mathsf{Holds}(s,\text{*}\mathsf{NOW})$.

Algorithm `assert_start`($S \subseteq$TEMP)

1. `setup_new_MTF`.
2. $P_{\text{new}} \longleftarrow \{\}$.
3. For all $s_i \in S$
   4. Pick some $t_i \in \mathcal{T}$, such that $t_i \notin \Psi(\mathcal{T})$.
   5. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s_i, t_i),\text{*}\mathsf{NOW} \sqsubset t_i\}$.
   6. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{MHolds}(s_i, t_i)\}$.
7. `Forward`(P_{new}).

Figure 21: Algorithm `assert_start`.

PRE-CONDITIONS:

1. *For every* $s \in S$, $\beta \vdash \mathsf{Holds}(s,\text{*}\mathsf{NOW})$.

Algorithm `assert_cease`($S \subseteq$ TEMP)

1. *old-now* $\longleftarrow$ *NOW.
2. `setup_new_MTF`.
3. For all $s_i \in S$

   $\beta \longleftarrow \beta \cup \{\neg\mathsf{Holds}(s,\text{*}\mathsf{NOW}), t_i \prec\text{*}\mathsf{NOW}\}$ where t_i is a situation interval associated with s_i such that $\beta \vdash$ *old-now* $\sqsubset t_i$.

Figure 22: Algorithm `assert_cease`.

PRE-CONDITIONS:

1. *For every* $s \in S$, *there is some* $\mu \in \mathcal{M}_{\text{prop}}$ *such that* $[\![s]\!]$ *occupies the modality corresponding to* μ, *or there is some* $\mu \in \mathcal{M}_{\text{per}}$ *such that* $[\![s]\!]$ *is perceived via the modality corresponding to* μ.

Algorithm `initialize`($S \in$ TEMP)

1. $P_{\text{new}} \longleftarrow \{\}$.
2. `initialize_NOW`.
3. For all $s_i \in S$
    4. Pick some $t_i \in \mathcal{T}$ such that $t_i \notin \Psi(\mathcal{T})$.
    5. `state_start`(s_i, t_i).
    6. $\beta \longleftarrow \beta \cup \{\textsf{*NOW} \sqsubset t_i\}$.
    7. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\textsf{MHolds}(s_i, t_i)\}$.
8. `Forward`(P_{new}).

Figure 23: Algorithm `initialize`.

of `assert_cease` and Definition 8.5, Cassie determines that s ceases to hold at $\textsf{*NOW}_i$. $\square$

Similar to the assumption we made regarding algorithm `state_change`, we assume that algorithms `assert_persist`, `assert_start`, and `assert_cease` provide the sole means by which Cassie may determine, through direct assertion, that a state persists, starts, or ceases, respectively. Such assumptions are needed to be able to prove further results about the system (see Sections 10.4 and 12).

Before concluding this section, one final algorithm needs to be introduced. This algorithm takes care of all the initialization procedures required by the system. Figure 23 shows the algorithm which basically sets up the first NOW-MTF. Step 2 initiates algorithm `initialize_NOW` of Figure 4. The rest of the steps update the belief space and modality variables to reflect states of Cassie's body or the perceived environment as Cassie starts operating.

9 The Perception of Time

The moving NOW provides Cassie with some sense of temporal progression. Nevertheless, the perception of time is not only confined to distinguishing a present moment or a chain of NOW-MTFs; it crucially involves a feel of the *amount* of time taken up by different intervals. In this section, we shall examine alternative means by which we can model this aspect of time perception. A particular model, partially-based on pacemaker theories (see Section 5), will be developed in Section 9.3.

9.1 Knowing and Feeling

Before setting out to discuss alternative models, we need to make a distinction between two concepts: the *knowledge* of a duration, and the *feel* thereof. Cassie's knowledge of the duration of some time interval, t, is represented by a proposition, $\mathsf{Dur}(t, q)$, in β, where q is the amount of time representing the duration of t (see Section 7.2). Such knowledge may come from any of the three sources mentioned in Section 6: direct assertion, inference, or bodily-feedback. There are no restrictions on whether t is in a NOW-MTF or a non-NOW-MTF. Thus, Cassie may have knowledge of the durations of events that she did not witness, but merely heard of from a third person. On the other hand, a feel of the duration of an interval requires a first-person experience of that interval. Formally, Cassie has a feel of the duration of an interval t only if every MTF in Span(t) is a NOW-MTF. That is, if every piece of t is in Cassie's chain of consciousness.

If knowledge of durations is represented by conscious beliefs, what is a *feel*, and how could it be represented? An analogy from the good old domain of color cognition may help. What do we know about colors? Two things. First, we have a mental entity that may be associated with a *name* for the color: *"red"*, *"blue"*, *"green"*, etc. Second, if we have *seen* an object of that color, we would have a *feel*, or perceptual experience, associated with the mental entity. In case we do not have a name for a color, the only possible way to express it is to use variations of "*the color of this or that object*". In our formal framework, the name of the color resides at the KL, and the feel, or perceptual experience, exists at the PML. Similarly, we have two corresponding notions for time intervals: mental entities, represented by symbols in $\mathcal{T}$; and perceptual experiences, represented by PML structures. Associations between KL and PML structures are represented by an *alignments set* (Shapiro, 1998). Essentially, the alignment set, $\mathcal{A}$, is a set of pairs, $\langle \tau, \pi(\tau) \rangle$, where τ is some KL term and $\pi(\tau)$ is the PML structure representing the perceptual experience, or feel, associated with τ. Thus, whereas knowledge of the duration of an interval, t, is represented by a belief in β, a feel for that duration is represented by a pair $\langle t, \pi(t) \rangle$ in $\mathcal{A}$. Of course, Cassie may have both a feel, $\langle t, \pi(t) \rangle$, and a belief, $\mathsf{Dur}(t, q)$,

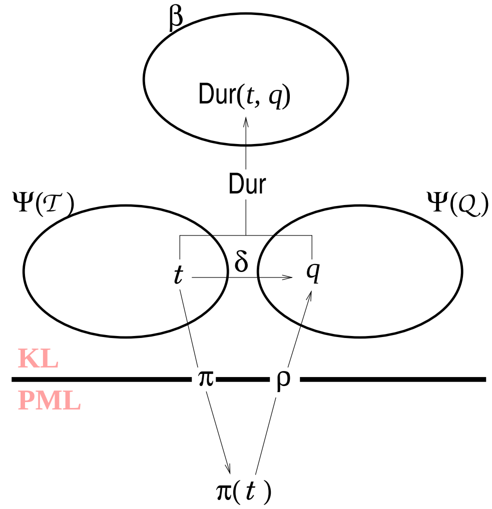

Figure 24: Mappings across the KL-PML interface.

about the duration of t. It is the duty of PML recognition processes to generate the appropriate q, given some $\pi(t)$. We assume the existence of a mapping, ρ, that would map $\pi(t)$ to the appropriate term in $\mathcal{Q}$ (see Figure 24, where δ maps a time interval to the duration associated with it in β). The exact characterization of ρ depends on the nature of PML representations and the interpretation of elements of $\mathcal{Q}$.

9.2 The Contents of an Interval

The basic claim of the cognitive theories of duration perception (Levin and Zakay, 1989, for instance) is that the contents of an interval provides the feel for its duration. The more the events happening within an interval, the longer it feels. The framework developed in the previous section readily provides such a measure; the larger $|\mathrm{Span}(t)|$, the longer t feels. Thus, $|\mathrm{Span}(t)|$ seems to provide a pretty good measure for the duration of t as dictated by cognitive theories. Nevertheless, there are two problems with such an approach. First, the contents of a duration in cognitive theories include both external events as well as internal mental events. According to the Second Principle of Change, NOW moves only when a change that Cassie is *aware* of takes place. As pointed out in Section 7.4, Cassie is not aware of internal changes caused by mental events, such as inference, only of external environmental events. Thus, $|\mathrm{Span}(t)|$ is a very rough and inaccurate measure of duration. Second, even if Cassie is aware of mental events, unlike humans, an artificial agent is not always busy-minded. Unless it is occupied with a particular reasoning task, typically triggered by some external event (including a query by a human operator), an artificial agent is "unconscious" most of the time. Thus, Span(t) would still be rather empty, and an inaccurate measure of the duration of t.

Even more important, note that relative durations are not maintained by $|\text{Span}(t)|$. Suppose that an interval t_1 is actually (i.e., for us, humans) much longer than an interval t_2. Using spans as measures for duration may make Cassie come to believe that t_2 is longer than t_1 just because she became aware of more events during the former. Now one may argue that this is actually reasonable. For what we want to represent is *Cassie's* sense of time, not her knowledge of ours. Thus, for Cassie, t_2 is *actually* longer than t_1. There are however two responses. First, one thing that we would like to represent is temporal regularity: that all occurrences of some event take almost the same amount of time to happen. This kind of regularity is lost if the feel for the durations of those occurrences is dependent on something as irregular as the changes that Cassie becomes aware of during them. Second, artificial agents typically operate in environments where they interact with humans. For this interaction to be effective, something as basic as the sense of time, should be unified. If Cassie were to operate in an environment inhabited by agents with similar perceptual capacities, and where *all* the relevant changes are perceivable, then $|\text{Span}(t)|$ may be a good measure of durations. However, as long as artificial agents operate among us, humans, they need to be adapted to *our* environment.

9.3 The Pacemaker

One problem with using the span of an interval as a measure for its duration, is that spans only represent the number of changes that happened during the interval; they do not carry any information about the amounts of time between those changes. It seems that if there is some means by which we can characterize those amounts, then the duration of an interval is simply the sum of all amounts of time between changes that happened within it. Note that the amount of time between two changes is essentially the duration of an MTF or the smallest element thereof.

Thus, the base case is the feel of durations of atomic intervals. In particular, the durations of NOW-intervals, since one can only have a feel of a duration if all the MTFs in its span are NOW-MTFs. The question then is how to account for this base case. What is it that may provide the feel for the duration of a NOW-interval? Note that one cannot resort to the number of changes, since NOW-intervals are atomic. We are, therefore, left with only one option: some sort of a pacemaker. Our pacemaker is a PML process; essentially, a counter that starts counting once the agent comes to life (i.e., starts operating), and is reset every time NOW moves. More specifically, a PML process periodically increments the integer value of a meta-theoretical (PML) variable, COUNT and a revised version of algorithm `move_NOW` (see Figure 25) resets it. The revised `move_NOW` aligns *NOW with the number of ticks produced by the

Algorithm `move_NOW`

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\beta \longleftarrow \beta \cup \{\textsf{*NOW} \prec t, \textsf{Dur}(\textsf{*NOW}, \rho(\textsf{*COUNT}))\}$.
3. $\mathcal{A} \longleftarrow \mathcal{A} \cup \{\langle \textsf{*NOW,*COUNT} \rangle\}$.
4. $\textsf{COUNT} \longleftarrow 0$.
5. $\textsf{NOW} \longleftarrow t$.

Figure 25: A revised version of `move_NOW` that aligns NOW-intervals with PML structures representing Cassie's sense of their duration.

Algorithm `initialize_NOW`

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\textsf{COUNT} \longleftarrow 0$.
3. $\textsf{NOW} \longleftarrow t$.

Figure 26: A revised version of `initialize_NOW` that resets the PML variable COUNT.

pacemaker since the last time NOW moved (step 3), thus providing a feel for its duration. Figure 26 shows a similarly revised version of algorithm `initialize_NOW`. There are a couple of things to note.

1. Ticks of the pacemaker at the PML, do not correspond to NOW-intervals at the KL. The dynamics of NOW is still governed by Axioms 7.3 and 7.4. The pacemaker merely provides the feel for the duration of NOW-intervals.
2. The rate at which COUNT is incremented is not significant so long as it is (i) constant, and (ii) fast enough to provide different feels for intervals whose durations need to be distinguished as dictated by the domain of application.

Now, let us take a careful look at step 3 of the algorithm. A new pair is added to the alignments set, $\mathcal{A}$, thereby extending it. Recall that a pair $\langle x, y \rangle$ in $\mathcal{A}$ represents an association between a KL term and a PML structure representing its perceptual, or bodily, experience. We assume, and may argue, that if $\langle x, y_1 \rangle \in \mathcal{A}$ and $\langle x, y_2 \rangle \in \mathcal{A}$, then $y_1 = y_2$. Thus, the set $\mathcal{A}$ may be thought of as an extensional representation of a function—a partial function from KL terms to PML structures. In this case,

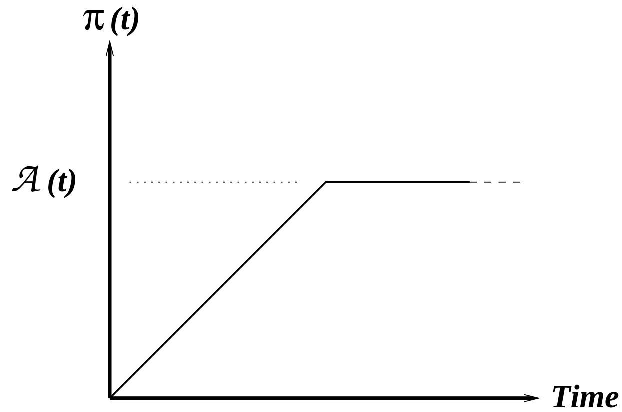

Figure 27: The behavior of $\pi(t)$ over time.

we can use $\mathcal{A}(\tau)$ to refer to the PML structure associated with the term τ in $\mathcal{A}$. But recall that, in Section 9.1 (also see Figure 24), we have indicated that pairs in $\mathcal{A}$ are of the form $\langle \tau, \pi(\tau) \rangle$, where π is a function that maps τ into its perceptual experience. Is there a difference between π and $\mathcal{A}$ (conceived as a function)? For most sorts of KL terms, π and $\mathcal{A}$ are indeed identical, but for terms in $\mathcal{T}$ there is a subtle difference between them. Following standard notation, let $\pi|_{\mathcal{T}}$ and $\mathcal{A}|_{\mathcal{T}}$ be the restrictions of π and $\mathcal{A}$ to $\mathcal{T}$, respectively. Note that $\mathcal{A}|_{\mathcal{T}}$ (and, in general, $\mathcal{A}$) is extensionally-defined, that is, $\mathcal{A}|_{\mathcal{T}}(t)$ is defined *only* if a pair $\langle t, \pi|_{\mathcal{T}}(t) \rangle$ is in $\mathcal{A}$. $\pi|_{\mathcal{T}}$, on the other hand, is intensionally-defined; there is an effective procedure for computing $\pi|_{\mathcal{T}}(t)$, *whenever possible*. By "whenever possible", we are stressing the partiality of $\pi|_{\mathcal{T}}$; recall that $\pi|_{\mathcal{T}}$ is defined for t only if all the members of Span(t) are NOW-MTFs. Thus, for some terms, there are occasions in which $\pi|_{\mathcal{T}}$ is defined and $\mathcal{A}|_{\mathcal{T}}$ is not.

To see what those *occasions* may be, note that the perceptual experience of an interval, the feel for its duration, evolves over time, as long as the interval has not moved into the past. Thus, $\pi|_{\mathcal{T}}(t)$ increases with time until *NOW is no longer a subinterval of t. At this point, $\pi|_{\mathcal{T}}(t)$ reaches a steady-state value which is the one that gets permanently associated with t. That is, $\mathcal{A}|_{\mathcal{T}}(t)$ *is* the steady state value of $\pi|_{\mathcal{T}}(t)$ (see Figure 27). Note that, just like atomicity, spans, and MTFs, $\pi|_{\mathcal{T}}$ is time-dependent. Actually, the same applies to $\mathcal{A}$, for, with time, the domain of $\mathcal{A}$ gets broader as more entities are perceived. But the dependency of $\pi|_{\mathcal{T}}$ on time is more radical; for a particular $t \in \mathcal{T}$, $\pi(t)$ changes with time as per Figure 27. Note that the horizontal coordinate of the origin in Figure 27 does not correspond to zero-time, the time Cassie comes to life. Rather, it corresponds to the time at which t is first conceived of, moving into $\Psi(\mathcal{T})$. The steady state value is reached once t expires and moves into the past. Note that π(*NOW) *is* *COUNT which increases linearly with time due to the constant rate of the pacemaker. Step (3) in Figure 25 simply associates *NOW with the steady state value of $\pi(\text{*NOW})$. For a non-atomic interval, t, $\pi(t)$ is defined according to the following equation.

$$\pi(t) = \sum_{\Phi(t_i)\in\ \mathrm{Span}(t)} \pi(t_i) \tag{1}$$

As long as *NOW is a subinterval of t, the value computed according to the above equation monotonically-increases with time, reflecting the increase in both $|\mathrm{Span}(t)|$ and $\pi(\text{*NOW})$. Once t has moved into the past $\pi(t)$ reaches its steady state value, $\mathcal{A}(t)$.

Having pointed out the various properties of $\pi|_{\mathcal{T}}$, the mapping from KL terms to PML structures, we now direct our attention to the other direction, the recognition mapping ρ.

9.4 From Feeling to Knowing

First, let us examine what properties of ρ are reasonable for what we take it to represent. For one thing, it is reasonable to assume that ρ is a function, that is, it maps an element of its domain to one and only one element of its range. This is the least we can require of a *recognition* mapping. The question now is whether it is one-to-one or many-to-one. Being one-to-one implies a very *sharp* recognition function, one that allows Cassie to *consciously* distinguish between two durations no matter how similar they feel. Typically, however, conscious knowledge is much coarser than perceptual experience. For example, we use the same word, *"red"*, to refer, not to a single sharp wave length, but to a band thereof. This is not to say that our perception does not distinguish different shades or hues of red, it is just that those distinctions have no impact on reasoning and communication. Therefore, we shall take ρ to be many-to-one, with elements of $\mathcal{Q}$ corresponding to ranges of numbers at the PML. Note that this is not an innovation; it is a standard practice in qualitative physics (Forbus, 1984, for example). This choice is also grounded in our interpretation of elements of $\mathcal{Q}$ as representing intuitions about *typical* amounts. Because of their inherent vagueness, typical amounts are best viewed as ranges, rather than specific values.

Another question is whether ρ should be an onto function. This would imply that every term in $\mathcal{Q}$ corresponds to some perceptual experience. Although, in principle, this may happen to be the case, we opt for the more liberal interpretation, and do not require ρ to be onto. The reason is that Cassie may conceive of amounts that she may *never* experience, due to limitations on perception. For example, even though we can reason about nano-seconds, or construct thought experiments about the speed of light, we cannot hope (at least for a while) to be able to have a direct

Algorithm $\eta(n \in \mathbb{N})$

1. If $n = 0$ then, return 0.
2. Return $1 + \text{round}(\log_{\sqrt{10}}(n))$.

Figure 28: Definition of the function η.

experience of how a nano-second feels, or how traveling at the speed of light may affect our sense of time.

To precisely establish the mapping ρ, we should decide on some partitioning of the set of natural numbers into intervals that correspond to elements of $\mathcal{Q}$. Admittedly, any partitioning would have to be arbitrary unless based on psychological evidence, which, as far as we know, does not exist. Nevertheless, we can still do better than picking some random partition; we partition the natural numbers into *half-orders of magnitude* (Hobbs, 2000). According to Hobbs, half-orders of magnitudes (HOMs) partition the positive reals into intervals geometrically-centered around $\sqrt{10}$, i.e., intervals of the form $[\sqrt{10}^{h-\frac{1}{2}}, \sqrt{10}^{h+\frac{1}{2}}]$, where h is a natural number. Hobbs provides data and argues that HOMs seem to be the right level of granularity for representing typical measures.

> "... I have observed that people find it ... easy to come up with half-order-of-magnitude estimates and that these are very often just as informative as they need to be. ... This suggests that there is some cognitive basis for thinking in terms of half orders of magnitudes. ... For scales that are isomorphic to the integers or the reals, precise values are often not available. We need coarser-grained structures on scales" (Hobbs, 2000, p. 28).

Since each HOM correspond to a unique natural number, h, HOMs are linearly-ordered. Naturally, elements of $\Psi(\mathcal{Q})$ should also be linearly-ordered to reflect the linear hierarchy of HOMs. Order over elements of $\mathcal{Q}$ is established by the function $<_{\mathcal{Q}}$ (see Section 7.2), and beliefs about the relative orders of various typical amounts is induced by the natural order on whole numbers at the PML. A function, η, maps the natural numbers onto half-orders of magnitude. The algorithm that computes η is shown in Figure 28. For some natural number, n, $\eta(n)$ evaluates to a natural number h, such that, if $n = 0$, $h = 0$. Otherwise, n is within the interval $[\sqrt{10}^{(h-1)-\frac{1}{2}}, \sqrt{10}^{(h-1)+\frac{1}{2}}]$. This identifies the HOM to which n belongs.[35] Elements

[35] The only reason we add 1 in step 2 of the algorithm is to distinguish the cases where $n = 1$ and $n = 0$. Note that much hangs on this; η merely generates natural numbers that uniquely correspond to distinct HOMs.

Algorithm $\rho(n \in \mathbb{N})$

1. $h \longleftarrow \eta(n)$.
2. If there is $q \in \Psi(\mathcal{Q})$ such that $\mathcal{A}(q) = h$, then return q.
3. Pick some $q \in \mathcal{Q}$, such that $q \notin \Psi(\mathcal{Q})$.
4. $\mathcal{A} \longleftarrow \mathcal{A} \cup \{\langle q, h \rangle\}$
5. $min \longleftarrow \{q' | \mathcal{A}(q') = h' \wedge h' < h\}$.
6. $max \longleftarrow \{q' | \mathcal{A}(q') = h' \wedge h < h']\}$.
7. If min is not empty, then $\beta \longleftarrow \beta \cup \{q_{gmin} <_{\mathcal{Q}} q\}$, where q_{gmin} is the greatest element of the linearly-ordered poset $\langle min, <_{\mathcal{Q}} \rangle$.
8. If max is not empty, then $\beta \longleftarrow \beta \cup \{q <_{\mathcal{Q}} q_{lmax}\}$, where q_{lmax} is the smallest element of the linearly-ordered poset $\langle max, <_{\mathcal{Q}} \rangle$.
9. return q.

Figure 29: The recognition function ρ.

of $\mathcal{Q}$ are associated in $\mathcal{A}$ with whole numbers corresponding to the HOMs they represent. Figure 29 shows the algorithm that computes the function ρ, given some some positive real number, n. Step (1) computes the HOM of n. Step (2) simply checks if a $\mathcal{Q}$-term corresponding to the HOM of n has already been introduced into $\Psi(\mathcal{Q})$. If not, step (3) introduces a new term. Steps (5) through (8) make sure that the new term is inserted in the appropriate position within the $<_{\mathcal{Q}}$-chain of elements in $\Psi(\mathcal{Q})$. Figure 30 depicts a commuting diagram for time intervals and their durations across the KL-PML interface. Note that, except for η, all the mappings depicted are partial.

The above algorithm guarantees that the image of ρ is linearly-ordered by $<_{\mathcal{Q}}$. That is, symbols denoting durations of intervals for which Cassie has a feel are linearly-ordered. What about intervals for which Cassie has no perceptual experience? For example, suppose that we tell Cassie, by direct assertion, that $\mathsf{Dur}(t, q)$, for some new t. Even though $\mathcal{A}(t)$ is not defined, Cassie may still know the position of q within the $<_{\mathcal{Q}}$-chain. First, q may already be in the chain, which would happen if $\mathcal{A}(q)$ is defined. Second, there could be some q' in the $<_{\mathcal{Q}}$-chain such that $\mathsf{Equiv}(q, q')$ is in β, in which case, both q and q' would occupy the same location in the chain. Otherwise, Cassie would not know the exact position of q unless she is

$$
\begin{tikzcd}
\Psi(\mathcal{T}) \arrow[r, "\delta"] \arrow[d, "\mathcal{A}\mid_{\mathcal{T}}"'] & \Psi(\mathcal{Q}) \arrow[d, "\mathcal{A}\mid_{\mathcal{Q}}"] \\
\mathbb{N} \arrow[r, "\eta"] \arrow[ur, dashed, "\rho"] & \mathbb{N}
\end{tikzcd}
$$

Figure 30: Commuting diagram for time intervals and their durations across the KL-PML interface.

explicitly told.

Now, there are two caveats that should be brought to the reader's attention.

1. Ideally, we should be able to make the strong assertion that the image of ρ is exactly the $<_{\mathcal{Q}}$-chain. That is, Cassie may only know the location of an amount in the $<_{\mathcal{Q}}$-chain if she has a perceptual experience of some interval, t, such that $\mathsf{Dur}(t, q)$. Nevertheless, we opt for a weaker assertion: we assume the image of ρ to be a *subset* of the $<_{\mathcal{Q}}$-chain, while maintaining that only amounts for which $\mathcal{A}$ is defined are in the chain. This implies that there may be amounts for which $\mathcal{A}$ is defined and that, nonetheless, are not in the image of ρ (nor δ, for that matter). These are amounts representing the typical durations of states. More specifically, they are terms in the set $\{q | \exists s[\beta \vdash \mathsf{SDur}(s, q)]\}$. Admittedly, there is no deep theoretical motivation for this step. However, it does give us a technical benefit. Our solution to the problem of the fleeting now (discussed in the next section) involves Cassie's comparison of the durations of intervals in the chain of NOW-MTFs to typical durations of states. This requires the relative orders of these durations to be known. In an ideal world, Cassie would be able to *learn* the typical durations of states by actually experiencing them. This is, indeed, attainable in our model (steps 3 through 8 in Figure 29). For practical purposes, however, we need to be able to *hardwire* associations between terms representing typical durations and HOMs in the alignments set $\mathcal{A}$.

2. For the purpose of this paper, reasoning about amounts is confined to reasoning about their order. Complex arithmetic reasoning within the KL (i.e., not subconscious PML computations) may be useful for a planning agent, for example. However, such reasoning is essentially probabilistic (Hobbs, 2000, pp. 29–31) and is beyond the scope of this work.

10 Seizing the Fleeting Now

10.1 The Lattice of "now"s

Let us now go back to the problem of the fleeting now and try to precisely formalize it within the framework developed in the previous sections. Cassie wonders whether some temporary state, s, holds "now". This wondering may be initiated by various events, including a direct English query by a human operator. "now" (or, in general, the present tense) is interpreted as $*\mathsf{NOW}$. Thus, Cassie's wonder initiates a deductive process for $\mathsf{Holds}(s, *\mathsf{NOW})$. Assuming that NOW is pointing to some term, t_1, then the deduction is actually for $\mathsf{Holds}(s, t_1)$, since $*\mathsf{NOW}$ is but a meta-theoretical shorthand for whichever term NOW happens to be pointing to. In order to figure out whether s holds, Cassie performs some sequence of acts, σ. Since these are acts that Cassie herself performs, she is aware of changes corresponding to their onsets and cessations. Therefore, by the First Principle of Change, NOW moves to some different time, t_2. Cassie determines that s indeed persists. According to the algorithms developed in Section 8.6 (particularly, `assert_persist`, `state_persist`, and `setup_new_MTF`), this is recorded by means of two assertions: $\mathsf{MHolds}(s, t)$ and $*\mathsf{NOW} \sqsubset t$, for some newly-introduced t. By **AS2** and **AS3**, this results in Cassie's coming to believe that $\mathsf{Holds}(s, *\mathsf{NOW})$. This belief *looks* exactly like the one for which the deduction was initiated. Nevertheless, since $*\mathsf{NOW}$ is merely a shorthand for t_2, what is deduced is $\mathsf{Holds}(s, t_2)$, whereas the deduction is for $\mathsf{Holds}(s, t_1)$.

Where exactly does the problem lie? We believe that it lies in the interpretation of "now" as $*\mathsf{NOW}$, the most fine-grained representation of the present time.[36] As pointed out in Section 3, there is generally a collection of intervals representing the present at different levels of granularity. This collection is, in principle, infinite, and Cassie may have a belief to that effect. Nevertheless, at any point, there is only a finite number of those in $\Psi(\mathcal{T})$, and each is introduced by a specific linguistic or reasoning discourse. One such discourse is the query of whether s holds "now". As pointed out in Section 4, "now" should be interpreted as an interval whose length is neither too short, nor too long, for the typical duration of s. This is the basic intuition. How to translate that into a solution to the problem is the subject of the rest of this section. First, a piece of notation.

Definition 10.1 $\lceil . \rceil$ *is a function from the set of* NOW*-intervals to* $\Psi(\mathcal{T})$ *such that, for every* $i \in \mathbb{N}$, $\lceil *\mathsf{NOW}_i \rceil$ *is the set of all reference intervals in* $\Phi(*\mathsf{NOW}_i)$. *That is,* $\lceil *\mathsf{NOW}_i \rceil = \{t | t \in \Phi(*\mathsf{NOW}_i)$ *and* t *is a reference interval*$\}$.

We can immediately make the following observations.

[36] Others may believe that the problem lies elsewhere (see Section 10.5 below), but hopefully our solution would satisfy everybody.

Observation 10.1 *For every $i \in \mathbb{N}$, $\lceil *\mathsf{NOW}_i \rceil$ is a temporal frame.*

Proof. Since, by definition, $\lceil *\mathsf{NOW}_i \rceil$ is a subset of a temporal frame (namely $\Phi(*\mathsf{NOW}_i)$), then, by Definition 7.3 (temporal frame), $\lceil *\mathsf{NOW}_i \rceil$ is a temporal frame. $\square$

Observation 10.2 *For every $i \in \mathbb{N}$, $*\mathsf{NOW}_i \in \lceil *\mathsf{NOW}_i \rceil$.*

Proof. Since $*\mathsf{NOW}_i$ is a reference interval (by Axiom 7.2), then, by Definition 10.1, $*\mathsf{NOW}_i \in \lceil *\mathsf{NOW}_i \rceil$. $\square$

Observation 10.3 *$\lceil . \rceil$ is one-to-one.*

Proof. Let $i, j \in \mathbb{N}$ such that $\lceil *\mathsf{NOW}_i \rceil = \lceil *\mathsf{NOW}_j \rceil$. Assume that $*\mathsf{NOW}_i$ and $*\mathsf{NOW}_j$ are distinct NOW-intervals. From Theorems 7.1 and 7.2, $\Phi(*\mathsf{NOW}_i)$ and $\Phi(*\mathsf{NOW}_j)$ are distinct NOW-MTFs. By Theorem 7.3 and Proposition 7.2,

$$\{*\mathsf{NOW}_i, *\mathsf{NOW}_j\} \subseteq \Phi(*\mathsf{NOW}_i) \triangle \Phi(*\mathsf{NOW}_j).$$

But, by Definition 10.1, $\lceil *\mathsf{NOW}_i \rceil \subseteq \Phi(*\mathsf{NOW}_i)$ and $\lceil *\mathsf{NOW}_j \rceil \subseteq \Phi(*\mathsf{NOW}_j)$. Therefore, by Observation 10.2,

$$\{*\mathsf{NOW}_i, *\mathsf{NOW}_j\} \subseteq \lceil *\mathsf{NOW}_i \rceil \triangle \lceil *\mathsf{NOW}_j \rceil.$$

This means that $\lceil *\mathsf{NOW}_i \rceil$ and $\lceil *\mathsf{NOW}_j \rceil$ are not identical, which leads to a contradiction. Therefore, $*\mathsf{NOW}_i = *\mathsf{NOW}_j$. Since i and j are arbitrary, then the result applies to all members of $\mathbb{N}$. Therefore, $\lceil . \rceil$ is one-to-one. $\square$

Observation 10.4 *For every $i \in \mathbb{N}$ ($i > 0$), $\langle \lceil *\mathsf{NOW}_i \rceil, \lambda x \lambda y(\beta \vdash x \sqsubset y \vee x = y) \rangle$ is a meet-semilattice.*

Proof. Given that $\sqsubset$ is a strict partial order, then, obviously, $\lambda x \lambda y(\beta \vdash x \sqsubset y \vee x = y)$ is a partial order. Since, for every $t \in \lceil *\mathsf{NOW}_i \rceil$, $\beta \vdash *\mathsf{NOW}_i \sqsubset t \vee *\mathsf{NOW}_i = t$, then every two elements of $\lceil *\mathsf{NOW}_i \rceil$ have an infimum, namely $*\mathsf{NOW}_i$. Therefore, $\langle \lceil *\mathsf{NOW}_i \rceil, \lambda x \lambda y(\beta \vdash x \sqsubset y \vee x = y) \rangle$ is a meet-semilattice. $\square$

We will, henceforth, refer to $\lceil *\mathsf{NOW}_i \rceil$ as *the i^{th} lattice of now's* (or simply *the lattice of now's* if the context is clear). It may, or may not, be linearly-ordered to form a *stack* of "now"s (Ismail and Shapiro, 2000b). The reason $\lceil *\mathsf{NOW}_i \rceil$ is not necessarily linear is that different extended "present"s may just overlap and are not always nested. A collection of reference intervals are nested only if they are all

Algorithm `state_query`(s)

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\beta \longleftarrow \beta \cup \{\textsf{*NOW} \sqsubset t, \textsf{Dur}(t, q)\}$, where $\beta \vdash \textsf{SDur}(s, q)$.
3. Initiate deduction for $\textsf{Holds}(s, t)$.

Figure 31: The algorithm `state_query`.

introduced for the first time in the same $\lceil\textsf{*NOW}_i\rceil$. That is, if they represent the *same* present time at different levels of granularity.

To solve the problem of the fleeting now, two points need to be revised:

1. How queries about states holding "now" are represented, and
2. How beliefs about states holding "now" are recorded.

The first point has already been discussed in Section 4. Figure 31 shows the algorithm `state_query` which outlines the steps taken whenever Cassie wonders whether some state, s, holds "now". Basically, a new reference interval is introduced, and its duration is restricted to be within the same HOM as the typical duration of the state s (but see below for a revision of this statement). Note that this new interval is a member of $\lceil\textsf{*NOW}\rceil$.

According to the algorithms of Section 8.6, to record that some state, s, holds, two main assertions are made: $\textsf{MHolds}(s, t)$, for some new t, and $\textsf{*NOW} \sqsubset t$. The point behind the second assertion is to make sure that Cassie believes that the state holds in the present. Given that we have different representations of the present, the second assertion should be replaced by a set of assertions for the *appropriate* elements of $\lceil\textsf{*NOW}\rceil$. In Section 4, we hinted that the appropriate elements of $\lceil\textsf{*NOW}\rceil$ are those whose durations are restricted to be less than, or within, the same HOM as the typical duration of s. Figure 32 outlines an algorithm that, given a state, s, and an interval, t, over which it holds, adds a set of assertions to β, to make sure that t includes all of the appropriate "now"s. Basically, t *bubbles up* the lattice of "now"s as high as it could, incorporating all those "now"s that fall within its extent. Note that, given the first conjunct of the conditional in step 2, including coarse-grained "now"s within the extent of t is, technically, a default assumption. Also note that, because of the same conjunct, algorithm `state_present` has effect only in the case of determining that a state persists, not that it starts.

How exactly are algorithms `state_query` and `state_present` linked to the rest of the system? We will answer this question in Section 10.3. But, first, we need to consider a fine adjustment of the algorithms themselves.

Algorithm `state_present`(s, t)

1. For every $t' \in \lceil$*NOW$\rceil$

   2. If $\beta \not\vdash \neg\mathsf{Holds}(s, t')$
   and $\delta(t') = q$ or $\beta \vdash \delta(t') <_{\mathcal{Q}} q$, where $\beta \vdash \mathsf{SDur}(s, q)$,

      then $\beta \longleftarrow \beta \cup \{t' \sqsubset t\}$.

Figure 32: The algorithm `state_present`.

10.2 The Backward Projection Factor

Algorithm `state_present` (Figure 32) is supposed to capture the intuition that, when one observes a state holding, they assume that it has persisted and will continue to persist for a while. In particular, whenever a state is observed to be holding, not only is it asserted to be holding over *NOW, the most fine-grained representation of the present, but over all representations of the present whose durations fall within the typical duration of the state. By doing so, we are projecting the state backward in time, which reflects the intuition that the state has not just started. This is very crucial for our proposed solution to the problem of the fleeting now, for it is this backward projection that justifies the agent's belief that the state held at some past query-time.

Now, the basic idea of backward projection is plausible, but the exact details of the process, embodied in algorithm `state_present`, may require some fine-tuning. Consider the example of the agent checking whether John is having lunch (see Section 4). The query takes place at 12:15 p.m. and let us say that the typical duration of having lunch is between 15 and 30 minutes. Suppose that, for some reason, the agent could only reach John's office at 12:40 p.m., when he sees John eating. Now, the time between the query and observation events is indeed within the 15-to-30-minutes window of having lunch, but is it reasonable for the agent to assume, by perceiving John having lunch at 12:40, that he was having lunch at 12:15? Our own intuition is: no. Had the walk to John's office taken 5 or 10 minutes, such an assumption would have been valid. Had it taken 15 minutes, the assumption would have still been possible but not as plausible. But 25 minutes, or even 20, is too long a period to safely make that assumption.

The problem is not whether backward projection per se is valid; we believe it certainly is. The problem is how far in the past one should project a state, how long one should assume that a state, observed to be holding, has been continuously holding. Algorithm `state_present` takes the position that one may project the state backward in time so long as the projection is within the typical duration of the state.

Algorithm `state_present`(s, t)

1. $q' \longleftarrow \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$, where $\beta \vdash \mathsf{SDur}(s, q)$.
2. For every $t' \in \lceil *\mathsf{NOW} \rceil$
   3. If $\beta \nvdash \neg\mathsf{Holds}(s, t')$
      and $\delta(t') = q'$ or $\beta \vdash \delta(t') <_{\mathcal{Q}} q'$,
      then $\beta \longleftarrow \beta \cup \{t' \sqsubset t\}$.

Figure 33: The revised algorithm `state_present`.

Examples like the above, however, renders such a position questionable. Backward projection should be allowed only within some fraction of the typical duration of the state. This is shown in Figure 33, a revision of algorithm `state_present`.

The main difference between this version of the algorithm and that in Figure 32 is the introduction of a duration q', instead of q, the typical duration of the state s, as an upper limit on the span of backward projection. Step (1) introduces q'—obviously a function of s. The computation of q' works as follows:

1. Compute the geometric center of the HOM corresponding to q ($\sqrt{10}^{(\mathcal{A}(q)-1)}$).
2. Raise this value to the value of *bpf*(s). Note that since HOMs are based on geometric means, this corresponds to "multiplying" the value computed in (1) by the factor *bpf*(s). Thus, the old version of algorithm `state_present` (Figure 32) is a special case where the value of this factor is unity.
3. Return the $\mathcal{Q}$-symbol, q', corresponding to the HOM to which the value computed in 2 belongs. Note that since we use the function ρ for this computation, q' may be a newly-introduced term.

Of course the question now is what the function *bpf* is. The function evaluates to a real number representing the *backward projection factor* of its argument state (hence the name). What do we know about this function? Not much; only that it evaluates to a positive real in the interval $(0, 1]$. Actually, we are not even sure if it is necessarily a function of the state, or if it is constant for all states, and in the former case, we do not know exactly how it depends on the state. We believe that the exact definition of the *bpf* function is an empirical question that we pose to psychologists: When observing a state holding, what are the biases of human subjects as to how long the state has been holding, and what are the factors determining those biases? As far as we know, the psychology literature is silent about these issues.

Algorithm `state_query`(s)

1. $q' \longleftarrow \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$, where $\beta \vdash \mathsf{SDur}(s, q)$.
2. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
3. $\beta \longleftarrow \beta \cup \{\mathsf{*NOW} \sqsubset t, \mathsf{Dur}(t, q')\}$.
4. Initiate deduction for $\mathsf{Holds}(s, t)$.

Figure 34: The algorithm `state_query` for the general case of backward projection factor that is not unity.

Modifying algorithm `state_present` as indicated above requires a similar modification to algorithm `state_query`. In particular, the reference interval introduced at the query time intuitively represents the period of time during which observing the state holding would be relevant to the query. As such, the duration of that reference interval should be restricted to q' (as computed above) rather than q, the typical duration of the state. This is shown in Figure 34. For generality, and since there is no decisive way to determine the innards of the *bpf* function, we shall not commit ourselves to any precise definition of *bpf*. For practical purposes, however, we take *bpf* to be the constant $\frac{1}{3}$. Note that this is just a working hypothesis that we will not build into our theory.

10.3 Forward Inference Meets Backward Projection

Let us now return to the question of how algorithms `state_query` and `state_present` fit within the rest of the system. First, consider algorithm `state_query` since there is not much to be said about it. The algorithm is initiated whenever a query is issued about whether some state holds in the present. This can happen in a number of ways. First, it could be a direct query by some other agent (for example, a question posed in English). Second, it could be an *internal* query generated by, for example, the acting system in the process of performing some conditional act.[37] Since we have not presented the complete acting system here, we will assume that algorithm `state_query` is initiated in all of the right places, whenever the need arises to check whether a state holds in the present. In what follows, we shall indicate the initiation of `state_query` with argument s at $[\![\mathsf{*NOW}_i]\!]$ by saying that *a query is issued for s at* $[\![\mathsf{*NOW}_i]\!]$.

Now, let us turn to algorithm `state_present`. Intuitively, the algorithm

[37]By "conditional act", we mean an act that is performed only if some state holds; for example, crossing the street only if the walk-light is on.

should be initiated whenever Cassie determines that a state persists. Thus, one might propose, a call to `state_present` should be added at the appropriate points in algorithms `state_change` and `assert_persist`. Granted, executing these algorithms should also, somehow, result in executing algorithm `state_present`. Nevertheless, note that this only accounts for cases where Cassie determines the persistence of a state through perception, proprioception, or direct assertion; it does not cover the case of inference. To appreciate this point, consider the following argument dismissing it as a problem.

How may Cassie infer that a state, s_1, holds in the present? Typically, this would involve Cassie's having a belief along the following schema, where s_2 is some state different form s_1 and p is a proposition describing background conditions that need to be true for s_2 to entail s_1.

$$\forall Tv[[\mathsf{Holds}(s_2, Tv) \wedge p] \Rightarrow \mathsf{Holds}(s_1, Tv)]$$

The argument now goes like this. Determining that s_1 holds through inference, necessarily involves determining that s_2 holds. Ultimately, this has to be based on determining that some s_n holds through some way other than inference, i.e., perception, proprioception, or direct assertion. Since algorithms `state_change` and `assert_persist` initiate forward inference, then following the above schema, it would be inferred that s_1 holds over all the coarse-grained "now"s over which s_2 holds. Therefore, backward projection is applied to s_1, albeit not directly through algorithm `state_present`.

There are at least two flaws in the above argument.

1. The reason the argument works is that, by performing simple forward inference, s_1 is backward-projected by virtue of s_2's projection. The problem, however, is that, inspecting the above schema, it is very possible that the typical duration of s_1 is longer than that of s_2. Thus, s_1 would not be projected into the past as far as *its* typical duration allows, only as far as s_2's typical duration does.[38]

2. Even if we choose to ignore the issue raised in 1, there is a more basic problem. Given the above schema, Cassie may infer that s_1 holds, not only by determining that s_2 holds, but also by coming to believe the proposition p (or an instance of the schema itself, for that matter). Even more dramatic, although determining that s_1 holds through inference typically involves a belief along the above schema, this is only typical, not necessary. What seems necessary is that the variable Tv be mentioned in the antecedent. For example, a possible translation of (3) appears below in (4), where Tc_1 and Tc_2 denote "3" and "5", respectively.

[38] This is similar to the problem of the persistence of derived information raised by (Myers and Smith, 1988) (also known as the problem of "dependent fluents" (Giunchiglia and Lifschitz, 1995)).

(3) Stu is home from 3 to 5.

(4) $\forall Tv[[Tc_1 \prec Tv \wedge Tv \prec Tc_2] \Rightarrow \mathsf{Holds}(\mathrm{At}(\mathrm{STU}, \mathrm{HOME}), Tv)]$.

By telling Cassie (4), she may determine that Stu is now home if she already believes that "now" is between 3 and 5. Note that this does not involve Cassie's determining that any other state holds.[39] Thus, should algorithm `state_present` be initiated only within `state_change` and `assert_persist`, Cassie may determine that the state At(STU, HOME) persists without the application of backward projection.

It, therefore, seems that backward projection should be built into the very process of inference. In particular, whenever Cassie infers that a state holds "now", backward projection should be applied. This may be achieved by introducing a time-sensitive version of the forward inference algorithm `Forward`. This is shown in Figure 35. The algorithm takes a set of propositions, P, as an argument. It does two main things: (i) applies traditional forward inference on members of P (this is achieved by algorithm `Forward_old` in step 1) and (ii) applies backward projection whenever appropriate. Step 1 initiates traditional forward inference on P. The set P_{inf} is the set of those propositions inferred in the process. The rest of the steps are responsible for backward projection. The conditional in step 3 filters in those members of $P \cup P_{\text{inf}}$ that are about states holding "now". Basically, it then initiates algorithm `state_present`. The only catch is that a situation interval is introduced (step 5a) if one is not associated with the state asserted to be holding "now". Note that, this way, adherence to Axiom 7.1 is built into the algorithms, and the theory builder need not worry about it. Also note that the filtering process in step 3 considers elements of both P_{inf} and P. Since new information is always introduced with forward inference (see Section 6), then backward projection gets applied to all states—those perceived, proprioceived, directly asserted, or inferred.

10.4 Pulling the Rope

By introducing coarse-grained reference intervals in both the querying and the assertion processes, the problem of the fleeting now may be readily solved. For whatever reference interval the querying process introduces would be available at the assertion time if its duration is long enough. Of course, what now remains is an account of how to carry this reference interval from one MTF to the next until it either expires, or the state is observed (*pulling the rope* à la Section 4). The first step is to modify the `setup_new_MTF` algorithm (Figure 10) so that it takes elements of $\lceil$*NOW$\rceil$

[39]It might be possible to rephrase (4) so that the antecedent involves Holds, but as (4) attests, it is also possible not to.

Algorithm $\texttt{Forward}(P \subseteq \Psi(\mathcal{P}))$

1. $P_{\text{inf}} \longleftarrow \texttt{Forward\_old}(P)$.
2. For every $p \in P \cup P_{\text{inf}}$
   3. If $p = \mathsf{Holds}(s,\text{*}\mathsf{NOW})$, for some $s \in \Psi(\text{TEMP})$, then
      4. If there is some $t \in \Psi(\mathcal{T})$ such that $\beta \vdash \mathsf{MHolds}(s,t)$, then $\texttt{state\_present}(s,t)$.
      5. Else
         5a. Pick some $t \in \mathcal{T}$ such that $t \notin \Psi(\mathcal{T})$.
         5b. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s,t), \text{*}\mathsf{NOW} \sqsubset t\}$.
         5c. $\texttt{state\_present}(s,t)$.

Figure 35: Algorithm **Forward**. Building backward projection into forward inference.

into account. Figure 36 outlines the modified algorithm. The algorithm assumes that the new MTF is the i^{th} NOW-MTF. Steps 7 through 9 incorporate members of $\lceil\text{*}\mathsf{NOW}_{i-1}\rceil$ into $\lceil\text{*}\mathsf{NOW}_i\rceil$ just in case the amount of time elapsed since they were introduced (represented by $\pi(t)$ which is computed according to Equation 1) is less than or within the same HOM as their projected durations (otherwise they move into the past as per step 9).

Given what we have presented so far, we can now formally prove the effectiveness of our solution to the problem of the fleeting now. To accomplish this, we need to introduce some notation that should help us proceed through the proofs more conveniently.

Definition 10.2 *For every $i, j \in \mathbb{N}$, the* **temporal distance** *between $\Phi(\text{*}\mathsf{NOW}_i)$ and $\Phi(\text{*}\mathsf{NOW}_j)$, denoted $d_t(\Phi(\text{*}\mathsf{NOW}_i), \Phi(\text{*}\mathsf{NOW}_j))$, is the amount of time that, Cassie feels, separates the start of $[\![\text{*}\mathsf{NOW}_i]\!]$ and the start of $[\![\text{*}\mathsf{NOW}_j]\!]$. More precisely,*

$$d_t(\Phi(\text{*}\mathsf{NOW}_i), \Phi(\text{*}\mathsf{NOW}_j)) = \sum_{k=\min(i,j)}^{\max(i,j)-1} \pi(\text{*}\mathsf{NOW}_k),$$

where $\min(i,j)$ and $\max(i,j)$ are the smaller and larger of i and j, respectively.

It should be clear that d_t is a metric over the space of NOW-MTFs. We shall not attempt to prove this, however, since the proof is (i) obvious and (ii) not instructive

Algorithm `setup_new_MTF`

1. `move_NOW`
2. For all $\mu \in \mathcal{M}_{\text{prop}}$
   3. If there are s and t such that $*\mu = \mathsf{MHolds}(s, t)$ then $\beta \longleftarrow \beta \cup \{*\mathsf{NOW} \sqsubset t\}$.
4. For all $\mu \in \mathcal{M}_{\text{per}}$
   5. For all s and t such that $\mathsf{MHolds}(s, t) \in *\mu$
      6. $\beta \longleftarrow \beta \cup \{*\mathsf{NOW} \sqsubset t\}$.
7. For all $t \in \lceil *\mathsf{NOW}_{i-1} \rceil \setminus \{*\mathsf{NOW}_{i-1}\}$
   8. If $\eta(\pi(t)) \leq \mathcal{A}(\delta(t))$, then $\beta \longleftarrow \beta \cup \{*\mathsf{NOW} \sqsubset t\}$.
   9. Else $\beta \longleftarrow \beta \cup \{t \prec *\mathsf{NOW}\}$.

Figure 36: The modified `setup_new_MTF` algorithm.

in any relevant sense. What should be noted though is that the temporal distance is only defined for NOW-MTFs since the definition primarily depends on their linear order (cf. Theorem 7.4) and the fact that π is defined for NOW intervals but not necessarily for other atomic intervals.

To prove the main result, we need to first prove two lemmas. The first asserts that reference intervals are extended into the appropriate NOW-MTFs.

Lemma 10.1 *For every $t \in \Psi(\mathcal{T})$, $q \in \Psi(\mathcal{Q})$, and $i, n \in \mathbb{N}$ ($i > 0$), if*

1. *t is a reference interval,*
2. *$\beta \vdash \mathsf{Dur}(t, q)$,*
3. *$\mathcal{A}(q)$ is defined,*
4. *for every $\Phi \in \text{Span}(t)$, Φ is a NOW-MTF,*
5. *$\Phi(*\mathsf{NOW}_i)$ is the smallest element of the poset $\langle \text{Span}(t), \text{precedes} \rangle$, and*
6. *$\eta(d_t(\Phi(*\mathsf{NOW}_i), \Phi(*\mathsf{NOW}_{i+n}))) \leq \mathcal{A}(q)$, then*
7. *$\Phi(*\mathsf{NOW}_{i+n}) \in \text{Span}(t)$.*

Proof. We use induction on n to prove the lemma.

Basis. Let $n = 0$. Since $\Phi(\text{*NOW}_i)$ is the smallest element of $\langle \text{Span}(t), \text{precedes} \rangle$, then, trivially, $\Phi(\text{*NOW}_i) \in \text{Span}(t)$.

Induction Hypothesis. Assume that, for every $t \in \Psi(\mathcal{T})$, $q \in \Psi(\mathcal{Q})$, and $i \in \mathbb{N}$ ($i > 0$), the conjunction of statements 1 through 6 implies statement 7, for some $n \in \mathbb{N}$.

Induction Step. We need to show that, for $n + 1$, the conjunction of statements 1 through 6 implies statement 7. By statement 6,

$$\eta(d_t(\Phi(\text{*NOW}_i), \Phi(\text{*NOW}_{i+n+1}))) \leq \mathcal{A}(q).$$

By Definition 10.2, and since $i < i + n + 1$,

$$\eta(\sum_{k=i}^{i+n} \pi(\text{*NOW}_k)) \leq \mathcal{A}(q).$$

Therefore,

$$\eta(\sum_{k=i}^{i+n-1} \pi(\text{*NOW}_k) + \pi(\text{*NOW}_{i+n})) \leq \mathcal{A}(q).$$

By Definition 10.2, and since $i \leq i + n$,

$$\eta(d_t(\Phi(\text{*NOW}_i), \Phi(\text{*NOW}_{i+n})) + \pi(\text{*NOW}_n)) \leq \mathcal{A}(q).$$

Since $\pi(\text{*NOW}_n)$ is a positive quantity, and since η is monotonic,[40] then

$$\eta(d_t(\Phi(\text{*NOW}_i), \Phi(\text{*NOW}_{i+n}))) \leq \mathcal{A}(q).$$

Therefore, using the induction hypothesis, $\Phi(\text{*NOW}_{i+n}) \in \text{Span}(t)$. By Definitions 7.8 (Span) and 10.1 ($\lceil . \rceil$), and statement 1, $t \in \lceil \text{*NOW}_{i+n} \rceil$. As NOW moves from *NOW_{i+n} to *NOW_{i+n+1}, algorithm `setup_new_MTF` gets executed and the conditional in step 8 is applied to t. By Equation 1,

$$\pi(t) = \sum_{\Phi(t_i) \in \ \text{Span}(t)} \pi(t_i)$$

[40]Note that this needs to proved; the proof is obvious though.

Now, note that at the time of evaluation the conditional, the greatest element of $\langle \text{Span}(t), \text{precedes} \rangle$ is $\Phi(\textsf{*NOW}_{i+n})$. Since $\Phi(\textsf{*NOW}_i)$ is the smallest element thereof (statement 5), and since all MTFs in Span(t) are NOW-MTFs (statement 4), then

$$\pi(t) = \sum_{\Phi(t_i) \in \text{ Span}(t)} \pi(t_i) \quad \leq \quad \sum_{k=i}^{i+n} \pi(\textsf{*NOW}_k).^{41}$$

But since

$$\eta(\sum_{k=i}^{i+n} \pi(\textsf{*NOW}_k)) = \eta(d_t(\Phi(\textsf{*NOW}_i), \Phi(\textsf{*NOW}_{i+n+1}))) \leq \mathcal{A}(q),$$

then

$$\pi(t) \leq \mathcal{A}(q) = \mathcal{A}(\delta(t)).$$

Therefore, by step 8 of algorithm `setup_new_MTF`, $\textsf{*NOW}_{i+n+1} \sqsubset t \in \beta$. Thus, by Definition 7.8 (Span), $\Phi(\textsf{*NOW}_{i+n+1}) \in \text{Span}(t)$. $\square$

Inspecting the proof of Lemma 10.1, we can draw the following monotonicity result. We will not show the proof since it could be easily reconstructed from the proof the lemma.

Corollary 10.1 *For all $i, j, k \in \mathbb{N}$, if $i \leq j \leq k$, then $d_t(\Phi(\textsf{*NOW}_i), \Phi(\textsf{*NOW}_j)) \leq d_t(\Phi(\textsf{*NOW}_i), \Phi(\textsf{*NOW}_k))$.*

We now prove the second lemma needed to present the main result. The lemma asserts that, whenever Cassie determines that a state persists, she believes that it holds over all of the appropriate members of the lattice of "now"s.

Lemma 10.2 *For every $s \in \Psi(\text{TEMP})$, $t \in \Psi(\mathcal{T})$, $q_1, q_2 \in \Psi(\mathcal{Q})$, and $i \in \mathbb{N}$ ($i > 1$), if*

1. *$q_1 = \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$, where $\beta \vdash \textsf{SDur}(s, q)$,*
2. *$\beta \vdash \textsf{Dur}(t, q_2)$,*
3. *$q_2 = q_1$ or $\beta \vdash q_2 <_{\mathcal{Q}} q_1$,*
4. *$t \in \lceil \textsf{*NOW}_i \rceil$,*

[41] This is, in fact, an equality given the convexity of t (**AT9**).

5. $\beta \not\vdash \neg\mathsf{Holds}(s, t)$, *and*

6. *Cassie determines that* s *persists at* $*\mathsf{NOW}_i$, *then*

7. $\beta \vdash \mathsf{Holds}(s, t)$.

Proof. By statement 6 and Definition 8.3, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $\beta \not\vdash \mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$ and $\beta \not\vdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i-1})$ and, at $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$. For this to be the case, there must be some $P \subset \mathcal{P}$, such that, at $[\![*\mathsf{NOW}_{i-1}]\!]$, for every $p \in P$, $\beta \not\vdash p$, and, at $[\![*\mathsf{NOW}_i]\!]$, $P \subseteq \beta$. The existence of such P is necessary since, at $[\![*\mathsf{NOW}_{i-1}]\!]$, $*\mathsf{NOW}_i$ was yet to be introduced. Since adding new information to β always initiates forward inference (see Section 6), then algorithm `Forward` of figure 35, gets executed, at $[\![*\mathsf{NOW}_i]\!]$, with P as an argument. Since, at $[\![*\mathsf{NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_i)$, then either $\mathsf{Holds}(s,*\mathsf{NOW}_i) \in P$ or $\mathsf{Holds}(s,*\mathsf{NOW}_i) \in P_{\mathrm{inf}}$ as computed by step 1 of `Forward`. Therefore, $\mathsf{Holds}(s,*\mathsf{NOW}_i) \in P \cup P_{\mathrm{inf}}$. Since, at this point, $*\mathsf{NOW}_i = *\mathsf{NOW}$, then algorithm `state_present` gets executed (by steps 4 or 5c) with s and t' as arguments, where t' is the situation interval associated with s at $[\![*\mathsf{NOW}_i]\!]$. By statements 1 through 5, step 5 of `state_present` adds the proposition $t \sqsubset t'$ to β. By **AS2** and **AS3**, $\beta \vdash \mathsf{Holds}(s, t)$. $\Box$

Given the above results, we can now prove the following theorem which establishes the effectiveness of our solution to the problem of the fleeting now.

Theorem 10.1 *For every* $s \in \Psi(\textsc{Temp})$, $q \in \Psi(\mathcal{Q})$, *and* $i, n \in \mathbb{N}$ $(i > 0)$, *if*

1. $\beta \vdash \mathsf{SDur}(s, q)$,

2. $\mathcal{A}(q)$ *is defined,*

3. *at* $[\![*\mathsf{NOW}_i]\!]$, *a query is issued for* s,

4. n *is the smallest integer such that Cassie determines that* s *persists at* $*\mathsf{NOW}_{i+n}$,

5. $d_t(\Phi(*\mathsf{NOW}_i), \Phi(*\mathsf{NOW}_{i+n})) \leq \sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)}$, *and*

6. *for all* m, $0 \leq m \leq n$, $\beta \not\vdash \neg\mathsf{Holds}(s,*\mathsf{NOW}_{i+m})$, *then*

7. *for all* m, $0 \leq m \leq n$, $\beta \vdash \mathsf{Holds}(s,*\mathsf{NOW}_{i+m})$.

Before stating the proof of the above theorem, we need to point out an issue that would render the proof not as formal as we would like it to be. It should be clear how the proof would proceed. Basically, we shall show that if the reference interval introduced at the time of the query (by algorithm `state_query`) survives until the time of determining the persistence of s, then, by Lemma 10.1, all the intervening "now"s are subintervals thereof. Using Lemma 10.2, we can then show that s holds all those "now"s. The problem, however, lies in statements 4 of Lemma 10.1 and 5 of Lemma 10.2 which should hold for the proof to follow. Statement 4 of Lemma 10.1 requires *all* MTFs in the span of the reference interval to be NOW-MTFs. Unfortunately, the theory presented here does not explicitly, and formally, specify that this is the case. For all we know, a user of the system (maybe a human operator interacting with Cassie) may make an assertion involving the reference interval, thereby adding some non-NOW-MTF to its span. Nevertheless, it should be noted that a valid assumption is that this would not happen (unless somebody is going out of their way to mess with Cassie's mind). Why? Because reference intervals introduced as extended representations of the present (in particular, those introduced by algorithm `state_query`) are, in a sense, *private* symbols of Cassie's mental language. More specifically, users of the system do not know anything about these internally-generated terms. More importantly, we envision assertions to be made to the system in natural language, and there would be no way to thus refer to these reference intervals. Therefore, in the following proof, we shall make the tacit assumption that reference intervals introduced as extended representations of the present contain only NOW-MTFs in their spans. In fact, we shall make a stronger (yet, still reasonable) assumption: the only assertions involving reference intervals introduced by algorithm `state_query` are those made by algorithms `state_query`, `state_present`, and `setup_new_MTF`. In particular, note in assertions of the form $t' \sqsubset t$, where t is a reference interval introduced by algorithm `state_query`, t' is always a NOW-interval. In addition, successive assertions of that form, involve successive values of NOW. Thus, the smallest element of $\langle \text{Span}(t), \text{precedes} \rangle$ is $\Phi(\text{*NOW}_i)$, where *NOW_i is the value of NOW at the time algorithm `state_query` gets executed.

Given this assumption, statement 5 of Lemma 10.2 follows from statement 6 of the theorem. Statement 5 of Lemma 10.2 requires that Cassie does not (implicitly or explicitly) believe that s does not hold over the reference interval, t, introduced by `state_query`. The problem is that, without the above assumption, anybody can assert anything about reference intervals, and it would be impossible to prove statement 5. How does the above assumption save the situation? To answer this question, consider how Cassie may come to believe that s does not hold over t. There are two possibilities. First, someone directly asserts $\neg\text{Holds}(s, t)$. This, however, is dismissed by our assumption. The only assertions involving reference intervals are those made by algorithms `state_query`, `state_present`, and `setup_new_MTF`. Evidently, these assertions are about temporal parthood; they do not mention any states. Second, Cassie believes that s does not hold over some sub-interval of t. However, given our assumption that the only sub-intervals of t are NOW-intervals, this possibility is also dismissed by statement 6 of the theorem.

Proof. By statement 3, algorithm `state_query` gets executed at $[\![\text{*NOW}_i]\!]$ (see the discussion in Section 10.3). Step 2 introduces a new reference interval, t', into

$\Psi(\mathcal{T})$. Step 3 makes the two assertions: $\mathsf{*NOW} \sqsubset t'$ and $\mathsf{Dur}(t', q')$, where $q' = \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$. Since, by statement 5,

$$d_t(\Phi(\mathsf{*NOW}_i), \Phi(\mathsf{*NOW}_{i+n})) \leq \sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)},$$

then by the monotonicity of η and the definition of ρ,

$$\eta(d_t(\Phi(\mathsf{*NOW}_i), \Phi(\mathsf{*NOW}_{i+n}))) \leq \mathcal{A}(q').$$

In fact, by Corollary 10.1 and the monotonicity of η, for every m, $0 \leq m \leq n$

$$\eta(d_t(\Phi(\mathsf{*NOW}_i), \Phi(\mathsf{*NOW}_{i+m}))) \leq \mathcal{A}(q').$$

Therefore, the following statements are true.

1. t' is a reference interval.
2. $\beta \vdash \mathsf{Dur}(t', q')$.
3. $\mathcal{A}(q') = \sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)}$ is defined.
4. For every $\Phi \in \mathrm{Span}(t')$, Φ is a NOW-MTF (which follows from the assumption discussed above).
5. $\Phi(\mathsf{*NOW}_i)$ is the smallest element of the poset $\langle \mathrm{Span}(t'), \mathrm{precedes} \rangle$ (which follows from the same assumption).
6. $d_t(\Phi(\mathsf{*NOW}_i), \Phi(\mathsf{*NOW}_{i+m})) \leq \mathcal{A}(q')$, for every m, $0 \leq m \leq n$.

Therefore, by Lemma 10.1, , for every m, $0 \leq m \leq n$, $\Phi(\mathsf{*NOW}_{i+m}) \in \mathrm{Span}(t')$. By Definition 7.8 (Span) and 10.1 ($\lceil . \rceil$), it follows that, for every m, $0 \leq m \leq n$, $t' \in \lceil \mathsf{*NOW}_{i+m} \rceil$ and $\mathsf{*NOW}_{i+m} \sqsubset t'$.

Now, by statement 6 of the theorem and the assumption discussed above, $\beta \not\vdash \neg\mathsf{Holds}(s, t')$. Therefore, the following statements are true.

1. $q' = \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$, where $\beta \vdash \mathsf{SDur}(s, q)$.
2. $\beta \vdash \mathsf{Dur}(t', q')$.
3. $t' \in \lceil \mathsf{*NOW}_{i+n} \rceil$.
4. $\beta \not\vdash \neg\mathsf{Holds}(s, t')$.
5. Cassie determines that s persists at $\mathsf{*NOW}_{i+n}$ (by statement 4 of the theorem).

Therefore, by Lemma 10.2, $\beta \vdash \mathsf{Holds}(s, t')$. By the divisitivity of states (**AS2**), it follows that, for every m, $0 \leq m \leq n$, $\beta \vdash \mathsf{Holds}(s, \mathsf{*NOW}_{i+m})$. $\square$

10.5 Where Exactly Does the Problem Lie?

All of the above being said, we should now consider possible objections and suspicions that some might raise against our solution—indeed, our interpretation—of the problem of the fleeting now. As pointed out in Section 3, the problem, we believe, lies in an intrinsic vagueness of the concept of "now": "now" refers to an interval that does not have any well-defined, context-independent boundaries. The NOW-intervals of the theory represent the concept of "now" at the finest level of granularity and non-NOW reference intervals stand for coarser representations thereof. Our solution to the problem of the fleeting now is based on interpreting the "now" of the query, not as *NOW, but as a coarser reference interval—a member of $\lceil$*NOW$\rceil$.

The above notwithstanding, it should be noted that, *technically*, nothing much hangs on this assumption. In particular, some may argue that such a move is not motivated, and that "now" is not vague, but always refers to the sharp instant of experience.[42] Even under that assumption, our solution would still work. This would only involve replacing step 3 in algorithm `state_query` (Figure 34) with "Initiate deduction for Holds(s,*NOW)". More precisely, the solution would work as follows.

1. Introduce a new reference interval, t, and restrict its length to a factor of the typical duration of s as determined by the *bpf* function.
2. Assert that t is a super-interval of t_1, the current value of *NOW.
3. Initiate deduction for Holds(s,*NOW).
4. Suppose that s is determined to persist at t_2.
5. If t survives until t_2, then algorithm `state_present` guarantees that we get Holds(s, t).
6. Now, by the divisitivity of states, we also get Holds(s, t_1), and the query is answered positively.

Thus, the only difference between the two scenarios (other than the conceptual one) is the use of the divisitivity of states to *link* determining the persistence of s, at t_2, to the query at t_1. In this case, the only crucial notion is that of backward projection.

But, now, someone may argue that, if backward projection is the key, then the issue is more general. For example, if, at 2 p.m., someone says that John was having lunch at 12:15 p.m. Should we assume that John was having lunch sharply at 12:15? Or should we employ backward projection and project the state of having lunch several minutes prior to 12:15? In more abstract terms, suppose that, at

[42] For example, Antony Galton, in personal communication.

$[\![*\mathsf{NOW}_i]\!]$, Cassie is told that a state, s, held at $*\mathsf{NOW}_{i-n}$. Shouldn't the backward projection mechanism be applied in this case too, asserting s to hold over all members of $\lceil *\mathsf{NOW}_{i-n} \rceil$ whose durations fall within $\mathrm{bpf}(s)$? Intuitively, this seems reasonable. Therefore, backward projection is not peculiar to the present, it may also be applied to assertions about the past, and, thus, the problem of the fleeting now has nothing to do with "now" per se!

Although, technically, the above argument may have some merit, there are reasons why we believe that it does not provide an adequate explanation of the problem of the fleeting now. The main point is that, though it might appear that backward projection is all that is needed to account for the problem, we believe that this is only the case at the technical level, not at the deeper, more fundamental, conceptual level. This belief stems from our concern, not only with explaining the reasoning aspects of the problem of the fleeting now, but with what it reveals about our very conceptualization of the present, in particular, as revealed by our use of language.

Consider, in more detail, the example of John's lunch. Suppose that Stu asks Cassie:

(5) Is John having lunch (now)?

Cassie walks to John's office and finds him eating his lunch. How should she reply to Stu?, Intuitively, she should say:

(6) Yes, he is.

What is interesting here is that *both* the question and the reply are in the present tense. How can that be, given that (5) and (6) are uttered, strictly speaking, at different times? The only way that this may be possible is if there is an interpretation of the present that encompasses the times of both utterances. This is exactly the gist of our solution to the problem of the fleeting now: interpreting the "now" of the question as a coarse reference interval. If this interval persists until the time of the reply (where the persistence is determined by contextual factors), then both the question and the reply fall within the same "now" and may, thus, be both expressed in the present tense.

Someone might claim that the fact that both (5) and (6) are expressed in the present tense is a mere peculiarity of language. In particular, since (6) is a reply to (5), it uses the same tense. However, this is obviously not true. Consider the situation where Cassie walks to John's office to find that he has just finished his lunch. In this case, (6) would not be an appropriate answer to (5). Rather, (7) (or a variant thereof) seems to be the only reasonable thing to say. The important thing to note here is that an affirmative answer to (5) would have to be expressed in the past tense. (Note that simply "He *was*" with the appropriate intonation is equally plausible.)

(7) He was when you asked, but not any more.

Another objection may be that (6) is expressed in the present tense simply because John is indeed having lunch at the time of its utterance. But consider the situation where Cassie walks to John's office to find him just about to start eating his lunch. Possible reasonable replies may be:

(8) He has just started.

(9) He *is* now, but not when you asked.

Or even:

(10) He is *now* having lunch.

Although all of these replies are in the present tense, they all (implicitly or explicitly) make it clear that, at the time of the question, John was not having lunch. Our ability to make these distinctions and yet find (6) to be a plausible reply in the original situation (where Cassie finds John in the midst of having lunch) can only be explained by the vagueness of the present.

11 Another Note on the Frame Problem

In Section 8.2, we proposed a mechanism for temporally-projecting perceived and bodily states. This proposal provides an elegant account of the *current* persistence of such states as time goes by—the variant of the frame problem that emerges in our theory. Nevertheless, the model of time perception developed in Section 9 has some further implications for the frame problem. Research on solving the frame problem is dominated by the important insight that, except for very few exceptions, the world is generally stable. Indeed, it must be, or otherwise we will not be able to function appropriately. (Shanahan, 1997) summarizes this view in what he calls "the common sense [sic] law of inertia": "Inertia is normal. Change is exceptional." (Shanahan, 1997, p. 18). It is this basic idea that underlies most of the solutions proposed to the frame problem. Now, as put by (Shanahan, 1997), the commonsense law of inertia indeed makes sense. However, what often lies behind most of the proposed solutions to the problem is a stronger version of the law of inertia. Typically, the assumption is that, unless it is known that something has happened which would cause some state to cease to hold, then the state still holds. Again, such an assumption might be reasonable as long as one is willing to be open-minded in interpreting "something has happened". In particular, there is something that is always happening, and yet has only been considered by a handful of authors (McDermott, 1982; Dean and Kanazawa, 1988; Lifschtiz and Rabinov, 1989; Kanazawa, 1992, for example) in the literature on the frame problem, namely *the passage of time*. Time is always moving and, as it does, it is reasonable to assume that some states have ceased to hold, based on intuitions about their typical durations. For example, if one sees a cat sitting on a mat, a green traffic-light, or a smiling face; it is not reasonable to assume that, an hour later, the cat is *still* on the mat, the light is *still* green, and the person is *still* smiling. In fact, (Stein, 1990, p. 372) advises that states "that are 'known' to change without warning" should be excluded from the scope of the commonsense law of inertia. What Stein fails to tell us is how such knowledge may come about.

Similar to our notion of the typical duration of a state, (McDermott, 1982) introduces the notion of a *life-time* associated with a *persistence*. Roughly, a persistence is the uninterrupted holding of a state over a period of time, which weakly corresponds to our situation intervals (in McDermott's terminology, our "states" are "facts"). The life-time associated with a persistence is the duration of the period of time throughout which a state may be assumed to persist unless otherwise known (note the non-monotonicity involved).

(Dean and Kanazawa, 1988) (followed by (Kanazawa, 1992)) present similar notions within a probabilistic framework. Using a discrete time line, Dean and Kanazawa present a system to compute the probability of a given fluent holding at a given time (their system does much more than that, but this is the aspect that concerns us here). Given a time, t, and a fluent, P, if no knowledge is available about any events causing P to start or cease between t and $t + \Delta$, then the probability of P holding over $t + \Delta$ is given by the following formula ($\langle P, t\rangle$ is their way of saying that P holds at t).

$$p(\langle P, t + \Delta\rangle) = p(\langle P, t + \Delta\rangle | \langle P, t\rangle)\ p(\langle P, t\rangle).$$

In the particular examples discussed in (Dean and Kanazawa, 1988), $p(\langle P, t+\Delta\rangle | \langle P, t\rangle) = e^{-\lambda\Delta}$. Thus, unless otherwise known, fluents (or states, in our theory) *decay* as time passes by—the same point made by (McDermott, 1982). Corresponding to our typical durations, the exponent λ is determined by the rate of decay (and, hence, the life-time) of the given fluent/state.

Although not introduced to present inherent temporal constraints on the persistence of states, Shanahan's notion of *trajectories* (Shanahan, 1997, ch. 13) may be used in that venue. In Shanahan's system, a formula "Trajectory(f1,t1,f2,d)" roughly means that if fluent f1 is initiated at time t1, then fluent f2 holds at time t1+d.[43] If the logic allows fluents of the form "not(f1)" to represent the fluent that holds whenever f1 does not (cf. (Allen, 1984)), then the formula "Trajectory(f1,t1,not(f1),d)" may be used to state that d is the life-time/typical duration of the fluent f1.

In all of the above-discussed systems, knowledge of how much time has passed is crucial. Where would this knowledge come from? It should be obvious that our model readily provides the answer: the pacemaker. Using Cassie's perception of time as provided by the pacemaker, she may defeasibly reason about the persistence of states. In particular, just as with reference intervals, a situation interval associated with a state for which there is a known typical duration may be excluded from a new MTF if the amount of time that has elapsed since the state was determined to be holding is longer than its typical duration. Alternately, a state, that is not known to have ceased, may be assumed to extend into a new MTF so long as it has not exceeded its typical duration. For an agent acting on-line, we believe that, in most cases, this is how it determines the persistence of states, since it seldom comes to know of events causing them to cease.

Consider the following example. Suppose Cassie is in a room with two, red and green, sources of light. Cassie looks at the red light, notices that it is on, and then turns toward the green light. Is the red light still on? The answer would reasonably be "yes" if the amount of time elapsed is within the same HOM as the typical duration of the red light. Otherwise, a suitable answer is "no". But now consider the following situation. Cassie turns back to the red light and finds that it is turned off. Now there are two possibilities. If the amount of time elapsed since she had last observed the light on is longer than its typical duration, then no problem; Cassie would have already assumed that the light is off. On the other hand, if not that much time has passed, then what is observed (the red light is off) contradicts what would be otherwise assumed (the red light is still on). Needless to say, knowledge induced through perception, bodily feedback, or direct assertion (our model for communication) is more credible than assumptions based on intuitions about typical durations of states. Thus, making use of typical durations to extend a situation interval into the new MTF should be a final resort, should only be turned to in case Cassie cannot infer that the situation interval has already moved into the past.

More formally, this can be achieved as per the third (and final) revision of

[43] Trajectories are used in (Shanahan, 1997) in order to represent continuous change. Shanahan's semantics involve a qualification of f1 and f2 that does not concern us here: f1 is a "discrete" fluent and f2 is a "continuous" fluent. Basically, a discrete fluent holds over intervals of non-zero durations, while a continuous fluent may hold instantaneously. The reader may notice that this is the same distinction that (Galton, 1990) makes between "states of motion" and "states of position". However, unlike Galton's notions, a fluent's being discrete or continuous is not an intrinsic property thereof, but an extrinsic one. To take an example form (Shanahan, 1997, p. 260), while a ball is falling, its height is a continuous fluent, but, once it lands, the height becomes a discrete fluent.

Algorithm $\texttt{setup\_new\_MTF}(S^{\downarrow} \subseteq \Psi(\textsc{Temp}))$

1. $\texttt{move\_NOW}$
2. For all $\mu \in \mathcal{M}_{\text{prop}}$
    3. If there are s and t such that $*\mu = \mathsf{MHolds}(s, t)$, then $\beta \longleftarrow \beta \cup \{\text{*}\mathsf{NOW} \sqsubset t\}$.
4. For all $\mu \in \mathcal{M}_{\text{per}}$
    5. For all s and t such that $\mathsf{MHolds}(s, t) \in *\mu$
        6. $\beta \longleftarrow \beta \cup \{\text{*}\mathsf{NOW} \sqsubset t\}$.
7. For all $t \in \lceil \text{*}\mathsf{NOW}_{i-1} \rceil \setminus \{\text{*}\mathsf{NOW}_{i-1}\}$
    8. If $\eta(\pi(t)) \leq \mathcal{A}(\delta(t))$, then $\beta \longleftarrow \beta \cup \{\text{*}\mathsf{NOW} \sqsubset t\}$.
    9. Else $\beta \longleftarrow \beta \cup \{t \prec \text{*}\mathsf{NOW}\}$.
10. For all $s \in S^{\downarrow}$ $\beta \longleftarrow \beta \cup \{\neg\mathsf{Holds}(s, \text{*}\mathsf{NOW})\}$.
11. For all $t \in \Phi(\text{*}\mathsf{NOW}_{i-1}) \setminus \lceil \text{*}\mathsf{NOW}_{i-1} \rceil$,
    12. If $\beta \vdash \text{*}\mathsf{NOW} \sqsubset t$, then $\beta \longleftarrow \beta \cup \{\text{*}\mathsf{NOW} \sqsubset t\}$.
    13. Else, if $\beta \vdash \text{*}\mathsf{NOW} \not\sqsubset t$, then $\beta \longleftarrow \beta \cup \{t \prec \text{*}\mathsf{NOW}\}$.
    14. Else, if $\eta(\pi(t)) \leq \mathcal{A}(q)$, where $\mathsf{SDur}(s, q)$ and $\mathsf{MHolds}(s, t)$, then
        $\beta \longleftarrow \beta \cup \{\text{*}\mathsf{NOW} \sqsubset t\}$.
    15. Else $\beta \longleftarrow \beta \cup \{t \prec \text{*}\mathsf{NOW}\}$.

Figure 37: The modified $\texttt{setup\_new\_MTF}$ algorithm.

algorithm `setup_new_MTF` in Figure 37. This version of the algorithm contains many features not in that of Figure 36 (in particular, steps 10 through 15). First, consider the loop starting in step 11. Step 11 restricts the application of the following steps to the set of situation intervals in the previous MTF. Step 12 includes into the new MTF all those states that hold in the previous MTF and that may be inferred to hold in the new MTF. For example, this includes states that Cassie has just ceased to perceive but that she, nonetheless, has reason to believe that they still hold. Step 13 takes a complementary action; it puts in the past all those situation intervals that may be inferred to *not* be members of the new MTF. These may include intervals associated with bodily states that have just ceased to hold, perceivable states that were just perceived to cease, and members of $S^{\downarrow}$ (to which we turn below). If a state cannot be inferred to hold, or not, in the new MTF, then steps 14 and 15 include or exclude such a state depending on whether it has exceeded its typical duration. Note that this is exactly the defeasible assumption made by (McDermott, 1982).

Now, let us turn to step 10 and the set $S^{\downarrow}$. Members of $S^{\downarrow}$ are states that have just been asserted to cease. That is, $S^{\downarrow}$ is itself the argument S of algorithm `assert_cease` of Figure 22. Step 10 simply makes sure that members of $S^{\downarrow}$ are excluded from the new MTF (and actually moved into the past by step 13). The reason we need to include $S^{\downarrow}$ and step 10 in algorithm `setup_new_MTF` is purely technical. Inspecting Figure 22, algorithm `assert_cease` is responsible for two main tasks: (i) moving NOW (with the construction of the new MTF) and (ii) updating β so that Cassie believes that states in the argument S no longer hold. The problem, however, is that updating β (as per Figure 22) takes place after executing algorithm `setup_new_MTF`, and, indeed, it must, since the new value of NOW needed for the propositions updating β is only introduced then. Why is this a problem? It is a problem, given the current version of `setup_new_MTF`, since step 14 may incorporate some of the states in S into the new MTF. We might, thus, be introducing inconsistencies into the system when this is, clearly, unwarranted. The only way that updating β be performed *after* introducing the new value of NOW but *before* executing step 14 of `setup_new_MTF` is if updating β is itself part of algorithm `setup_new_MTF`. This is exactly what step 10 is responsible for. Note that, this way, algorithm `assert_cease` simply initiates `setup_new_MTF` (see Figure 38).[44]

For ease of reference, Appendix A contains a compilation of all the temporal progression algorithms presented throughout the paper.

[44]It should be noted that, despite this revision of `assert_cease`, Theorem 8.6 still holds.

PRE-CONDITIONS:

1. *For every* $s \in S$, $\beta \vdash \mathsf{Holds}(s, {*}\mathsf{NOW})$.

Algorithm `assert_cease`($S \subseteq$ TEMP)

1. `setup_new_MTF`(S).

Figure 38: Revised version of algorithm `assert_cease`.

12 On Soundness and Completeness

Before concluding our investigation, we should say something about the correctness of the system presented above. As the title of this section shows, the plan is to discuss issues of soundness and completeness. A word of caution though. We are *not* going to provide proofs of soundness and completeness that usually accompany the presentation of a logical system (hence the "on" in the title). Rather, we shall present criteria for soundness and completeness that are slightly different, and more appropriate for our purposes, from the traditional ones, and outline, in as precise terms as possible, proofs that our system observes these criteria.

First of all, let us remind the reader that, as far as our logic is concerned, we do not need to show any proofs of (traditional) soundness and completeness; our logic is a standard first-order logic, no modal operators, no default rules, nothing exotic. What is the problem then? The problem is that traditional soundness and completeness worry about rules of inference, whether one can infer all and only those things that are true in every interpretation in which an initial set of premises is true. Where the premises are coming from is irrelevant to traditional soundness and completeness, and rightfully so. But in our system we need to worry about where some of the premises are coming from. In particular, according to Sections 8.2 and 11, the temporal projection of states is sometimes not the result of logical inference, but of subconscious temporal progression processes. Beliefs generated thus are, technically, premises since their presence in β is not the result of inference. On the other hand, they are not exactly normal premises pushed into β by fiat; they are the result of the interaction of the algorithms, the meta-theoretical axioms, the contents of β, and the rules of inference. Since these constitute our system, Cassie's temporal machinery, we need to show that it does the right job, that Cassie believes (or can infer) all and only those propositions that she is justified to believe based on what she knows, and more importantly, what she feels (which makes our task different from that of autoepistemic logicians (Moore, 1984; Moore, 1985)).

In particular, given the focus of our investigation, we are interested in two types of beliefs:

1. A state s's holding in the present.
2. The event of some state s's holding being in the past.

In the terminology developed above, propositions of type 1 are of the form $*\mathsf{NOW} \sqsubset t$, where t is a situation interval associated with s. Similarly, type-2 propositions are of the form $t \prec *\mathsf{NOW}$. We need to show that our system results in Cassie's believing propositions of type 1 or 2 if and only if it is reasonable for her to do so. The question now is what "reasonable" means. This depends primarily on the type of state in question and how Cassie comes to believe propositions about it.

12.1 Feeling is Believing: A "Completeness" Result for Bodily States

For a bodily state, s, characterizing when it is reasonable for Cassie to believe that s holds in the present is clear-cut—when, and only when, the state occupies some proprioceptual modality. Similarly, Cassie should believe that an event of s holding has moved into the past when, and only when, the state stops occupying proprioceptual modalities. Note that nowhere here have we required s to be actually holding, in the first case, or have actually ceased to hold in the second. All that concerns us is whether Cassie *feels* that it does. For example, Cassie might be holding an object but feels that she is empty-handed due to some glitch at the SAL. Yet, believing that she is empty-handed is justified by what Cassie feels.[45] Thus, we do not care whether what Cassie believes is in fact knowledge—justified true belief; we only care about its being justified. That is, we only require Cassie to be *rational*.

The main result of this section is Theorem 12.1 below. The theorem simply states that, by the system presented above (that is, the logic, the axioms, and the PML algorithms), whenever a bodily state occupies some proprioceptual modality, Cassie has a belief that that state holds. Note what this means. It means that bodily sensations, or feelings, purely PML phenomena, are aligned with conscious beliefs at the KL. Thus, in a sense, the theorem states a completeness result. In order to prove the theorem, we need to prove a lemma first. Informally, the lemma states that the values of modality variable are faithful to what Cassie feels.

Lemma 12.1 *For every $i \in \mathbb{N}$ $(i > 0)$, $s \in$ TEMP such that $\mathrm{Mod}_{\mathrm{prop}}(s) \neq \{\}$, and $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$, at $[\![*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ occupies the modality corresponding to μ if and only if there is some $t \in \Psi(\mathcal{T})$ such that $*\mu = \mathsf{MHolds}(s, t)$ and $\beta \vdash \mathsf{MHolds}(s, t)$.*

Proof. Let s be a bodily state and let μ be in $\mathrm{Mod}_{\mathrm{prop}}(s)$. We use induction on i to prove the lemma.

[45] Unless Cassie believes that there is some SAL problem.

Basis. Let $i = 1$. $\Phi(*\mathsf{NOW}_1)$ is established by initiating algorithm `initialize`. Suppose that $[\![s]\!]$ occupies the modality corresponding to μ. Therefore, s is a member of S, the argument of `initialize`. By executing algorithm `state_start` in step 5, the proposition $\mathsf{MHolds}(s,t)$ gets added to β (step 1) and algorithm `state_proprioceive` gets executed with arguments s and t (step 2), where t is the interval introduced by step 4 of `initialize`. Since $[\![s]\!]$ occupies the modality corresponding to μ, then, by step 2 of `state_proprioceive`, $*\mu = \mathsf{MHolds}(s,t)$, and, trivially, $\beta \vdash \mathsf{MHolds}(s,t)$ since $\mathsf{MHolds}(s,t) \in \beta$. Conversely, suppose that $[\![s]\!]$ does not occupy the modality corresponding to μ. By Axiom 8.1, some other state, $[\![s']\!]$ occupies that modality. Following the same argument outlined above, algorithm `initialize` results in assigning μ a proposition $\mathsf{MHolds}(s',t')$, for some $t' \in \Psi(\mathcal{T})$. Therefore, $*\mu \neq \mathsf{MHolds}(s,t)$ for any $t \in \mathcal{T}$.

Induction Hypothesis. Assume that, at $[\![*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ occupies the modality corresponding to μ if and only if there is some $t \in \Psi(\mathcal{T})$ such that $*\mu = \mathsf{MHolds}(s,t)$ and $\beta \vdash \mathsf{MHolds}(s,t)$.

Induction Step. Consider the situation at $[\![*\mathsf{NOW}_{i+1}]\!]$. We will break the proof of the induction step into four sub-proofs: sub-proofs (a) and (b) cover two complementary cases of the *if* part of the theorem, sub-proofs (c) and (d) cover two similar cases of the *only-if* part.

(a). Suppose that $[\![s]\!]$ occupies the modality corresponding to μ at both $[\![*\mathsf{NOW}_{i+1}]\!]$ and $[\![*\mathsf{NOW}_i]\!]$. By the induction hypothesis, at $[\![*\mathsf{NOW}_i]\!]$, μ is set to $\mathsf{MHolds}(s,t)$, for some $t \in \Psi(\mathcal{T})$. In addition, $\beta \vdash \mathsf{MHolds}(s,t)$. Thus, we only need to show that μ retains its value over $[\![*\mathsf{NOW}_{i+1}]\!]$. The value of NOW can change from $*\mathsf{NOW}_i$ to $*\mathsf{NOW}_{i+1}$ only if one of algorithms `state_change`, `assert_start`, or `assert_cease` is executed. Since `assert_start` and `assert_cease` do not involve any steps that set proprioception modality variables, then if NOW moves due to the execution of either algorithm, μ would retain its value. Now, suppose that NOW moves due to the execution of algorithm `state_change`. The only places in `state_change` where proprioception modality variables are set are steps 3 and 4b (through the execution of algorithm `state_proprioceive`). But since $[\![s]\!]$ occupies the modality corresponding to μ at both $[\![*\mathsf{NOW}_{i+1}]\!]$ and $[\![*\mathsf{NOW}_i]\!]$, then by Axiom 8.1 and pre-condition 2, there is no $s' \in S^{\uparrow}$ such that $[\![s']\!]$ occupies the modality corresponding to μ. Thus, μ never gets changed by algorithm `state_proprioceive`. Therefore, at $[\![*\mathsf{NOW}_{i+1}]\!]$, following the execution of `state_change`, $*\mu = \mathsf{MHolds}(s,t)$.

(b). Suppose that, at $[\![*\mathsf{NOW}_{i+1}]\!]$, but not at $[\![*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ occupies the modality corresponding to μ. By Axiom 8.2, at $[\![*\mathsf{NOW}_i]\!]$, there is no $\mu' \in \mathcal{M}_{\text{prop}}$ such that $[\![s]\!]$ occupies the modality corresponding to μ'. Further, by Axiom 8.5, there is no $\mu' \in \mathcal{M}_{\text{per}}$ such that $[\![s]\!]$ is perceived via the modality corresponding to μ'. By the induction hypothesis, at $[\![*\mathsf{NOW}_i]\!]$, there is no $t \in \mathcal{T}$ and $\mu' \in \mathcal{M}$

such that $*\mu' = \mathsf{MHolds}(s,t)$. Therefore, as $[\![s]\!]$ undergoes the transition from not occupying the modality corresponding to μ to occupying it, s satisfies pre-conditions 1 and 2 of algorithm `state_change`, and the algorithm gets initiated with $s \in S^{\uparrow}$. By the proof of Theorem 8.1, at $[\![\text{*NOW}_{i+1}]\!]$, for some $t \in \mathcal{T}$, $*\mu = \mathsf{MHolds}(s,t)$ and $\beta \vdash \mathsf{MHolds}(s,t)$.

(c). Suppose that, at $[\![\text{*NOW}_i]\!]$, but not at $[\![\text{*NOW}_{i+1}]\!]$, $[\![s]\!]$ occupies the modality corresponding to μ. By Axiom 8.1, there is some s' ($s' \neq s$) such that, at $[\![\text{*NOW}_{i+1}]\!]$, but not at $[\![\text{*NOW}_i]\!]$, $[\![s']\!]$ occupies the modality corresponding to μ. Thus, following the proof of part (b) above (switching s' and s), at $[\![\text{*NOW}_{i+1}]\!]$, $*\mu = \mathsf{MHolds}(s',t')$, for some $t' \in \mathcal{T}$. Therefore, at $[\![\text{*NOW}_{i+1}]\!]$, $*\mu \neq \mathsf{MHolds}(s,t)$, for any $t \in \mathcal{T}$.

(d). Suppose that, at neither $[\![\text{*NOW}_i]\!]$ nor $[\![\text{*NOW}_{i+1}]\!]$, does $[\![s]\!]$ occupy the modality corresponding to μ. By Axiom 8.1, there are s' and s'', different from s, such that, at $[\![\text{*NOW}_i]\!]$, $[\![s']\!]$ occupies the modality corresponding to μ and, at $[\![\text{*NOW}_{i+1}]\!]$, $[\![s'']\!]$ occupies the modality corresponding to μ. If $s'' = s'$, then, following the proof of part (a) above, at $[\![\text{*NOW}_{i+1}]\!]$, $*\mu = \mathsf{MHolds}(s'',t'')$, for some $t'' \in \mathcal{T}$. If, on the other hand, $s'' \neq s'$, then following the proof of part (b) above, at $[\![\text{*NOW}_{i+1}]\!]$, $*\mu = \mathsf{MHolds}(s'',t'')$, for some $t'' \in \mathcal{T}$. Since $s \neq s''$, then in either case, at $[\![\text{*NOW}_{i+1}]\!]$, $*\mu \neq \mathsf{MHolds}(s,t)$, for any $t \in \mathcal{T}$.

From (a), (b), (c), and (d) the induction step follows. Since s and μ are arbitrary, the lemma follows. □

Given the above result, we are now ready to prove the main theorem for this section.

Theorem 12.1 *For every $i \in \mathbb{N}$ ($i > 0$), $s \in$ TEMP such that $\mathrm{Mod}_{\mathrm{prop}}(s) \neq \{\}$, and $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$, at $[\![\text{*NOW}_i]\!]$, if $[\![s]\!]$ occupies the modality corresponding to μ, then $\beta \vdash \mathsf{Holds}(s, \text{*NOW}_i)$.*

Proof. Pick some $i \in \mathbb{N}$ ($i > 0$), $s \in$TEMP, and $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$. Suppose that, at $[\![\text{*NOW}_i]\!]$, $[\![s]\!]$ occupies the modality corresponding to μ. By Lemma 12.1, at $[\![\text{*NOW}_i]\!]$, $*\mu = \mathsf{MHolds}(s,t)$ and $\beta \vdash \mathsf{MHolds}(s,t)$ for some $t \in \Psi(\mathcal{T})$. Step 3 of algorithm `setup_new_MTF` results in adding the proposition $p = \text{*NOW}_i \sqsubset t$ to β. Given p, **AS2**, and **AS3**, $\beta \vdash \mathsf{Holds}(s, \text{*NOW}_i)$. Since i, s, and μ are arbitrary, then the result applies to all $i \in \mathbb{N}$, $s \in$TEMP, and $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$. □

12.2 Is Believing Feeling?

The above theorem is essentially a completeness-like result. Similarly, we may prove the following soundness-like Theorem.

Theorem 12.2 *For every* $i \in \mathbb{N}$ $(i > 0)$, $s \in$ TEMP *such that* $\mathrm{Mod}_{\mathrm{prop}}(s) \neq \{\}$, *and* $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$, *at* $[\![\text{*NOW}_i]\!]$, *if* $\beta \vdash \mathsf{Holds}(s,\text{*NOW}_i)$, *then* $[\![s]\!]$ *occupies the modality corresponding to* μ.

Proof. We prove the theorem using contraposition. Pick some $i \in \mathbb{N}$ $(i > 0)$, $s \in$TEMP, and $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$. Suppose that, at $[\![\text{*NOW}_i]\!]$, $[\![s]\!]$ does not occupy the modality corresponding to μ. By Axioms 8.1 and 8.2, there is some $s' \in$ TEMP such that $s' \neq s$, $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s')$, and, at $[\![\text{*NOW}_i]\!]$, $[\![s']\!]$ occupies the modality corresponding to μ. By Theorem 12.1, at $[\![\text{*NOW}_i]\!]$, $\beta \vdash \mathsf{Holds}(s',\text{*NOW}_i)$. Given Axiom 8.3, at $[\![\text{*NOW}_i]\!]$, $\beta \vdash \neg\mathsf{Holds}(s,\text{*NOW}_i)$. Since i, s, and μ are arbitrary, then the result applies to all $i \in \mathbb{N}$, $s \in$TEMP, and $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$. $\square$

Having proved the above theorem, we feel obliged to raise a number of issues that may render the reader a little skeptic about Theorem 12.2 (and Theorem 12.1, for that matter). This, however, is not due to a problem with our system or our mathematical skills, but due to a number of reasons that have to do with gullibility, astuteness, and the need for future research.

First of all, recall what we mean by "soundness" here. We do not merely mean the property of a *logical system* that allows only valid *inferences*, we mean that feature of a *cognitive agent* that allows it to have only *justified beliefs*. As per Theorem 12.1, bodily sensations justify beliefs about bodily states, but there could be other justifications for such beliefs too. For example, a human (or otherwise) agent may *tell* Cassie about states of her own body, for example that she is holding a fruit (recall our discussion in Section 8.5 in relation to the sixth principle of change). If Cassie feels that she is holding a fruit, then no problem. But what if she does not? Should she *believe* that she is holding a fruit? It depends. If Cassie is a highly-gullible agent (which most AI systems are), then she should believe whatever she is told, and that would be a counter-example for the *believing-is-feeling* hypothesis for bodily states. But this is only considering one extreme of the scale of agent-gullibility. On the other extreme, Cassie always ignores, and never believes, anything she is *told* about her bodily states; if she feels them, she will believe in them, according to Theorem 12.1. But this too is an extreme position. In fact, in order for Cassie to initially learn what different bodily sensations mean, some outside help is inevitable. A moderate position is, therefore, required. A possibility would be for Cassie to believe those assertions about states of her body that do not contradict beliefs invoked by what she feels. For example, feeling and believing that she is holding an apple she should reject an outside statement indicating that she is empty-handed. However, given the way we have developed our system (in particular, Axioms 8.1 and 8.3), unless Cassie is told exactly what she feels, the asserted information would have to be contradictory to Cassie's beliefs. Thus, in this case too, Cassie would reject anything that she is told about bodily states she does not feel. The most lenient position, the one we tacitly adopt, is to allow Cassie to hold beliefs contradictory to what she feels. There are justifications for this. For suppose that the holding-an-apple bodily sensations are mere hallucinations due to some SAL problem. In that case, Cassie may learn about such a problem from the contradictory external input as long as it is coming from some credible source (for example, a human supervisor, not a fellow robot). What this boils down to, then, is the issue of belief revision in a multi-source environment. Progress on this front is currently underway but still not complete enough to be integrated into our theory (see (Johnson and Shapiro, 2000a) for a preliminary report). Were Cassie to choose to believe that she is holding an apple (because the source of the "empty-handed"-assertion is less credible than proprioception) than no problem; Theorems 12.1 and 12.2 are upheld. On the other hand, if Cassie were to believe that she is empty-handed, then both theorems would be violated. Thus, Theorems 12.1 and 12.2 are valid as long as Cassie does not choose to believe, as mandated by belief revision principles, that a bodily state holds when (Cassie feels) that it does not occupy any modalities.

12.3 Seeing is Believing: A "Completeness" Result for Perceived States

Similar to what we did in Section 12.1, we shall show that our system guarantees that, whenever Cassie is perceiving a state holding, then she may syntactically-infer that it does. We first prove a lemma similar to Lemma 12.1

Lemma 12.2 *For every $i \in \mathbb{N}$ ($i > 0$) and $s \in$ TEMP such that $\mathrm{Mod}_{\mathrm{per}}(s)$ is defined, at $[\![\mathsf{*NOW}_i]\!]$, $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$ if and only if there is some $t \in \Psi(\mathcal{T})$ such that $\mathsf{MHolds}(s,t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$ and $\beta \vdash \mathsf{MHolds}(s,t)$.*

Proof. Pick some $s \in$ TEMP such that $\mathrm{Mod}_{\mathrm{per}}(s)$ is defined. We use induction on i to prove the lemma.

Basis. Let $i = 1$. $\Phi(\mathsf{*NOW}_1)$ is established by initiating algorithm `initialize`. Suppose that $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$. Therefore, s is a member of S, the argument of `initialize`. By executing algorithm `state_start` in step 5, the proposition $\mathsf{MHolds}(s,t)$ gets added to β (step 1) and algorithm `state_perceive` gets executed with arguments s and t (step 2), where t is the interval introduced by step 4 of `initialize`. Since $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$, then, by algorithm `state_perceive`, $\mathsf{MHolds}(s,t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$. Conversely, suppose that $[\![s]\!]$ is not perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$. By Axiom 8.4 and pre-condition 1, $s \notin S$. Thus, the execution of algorithm `state_perceive` would not result in adding any propositions concerning s to $*\mathrm{Mod}_{\mathrm{per}}(s)$. Since, for $[\![\mathsf{*NOW}_1]\!]$, algorithm `initialize` is the only place where modality variables are set, then, at $[\![\mathsf{*NOW}_1]\!]$, $*\mathrm{Mod}_{\mathrm{per}}(s) \neq \mathsf{MHolds}(s,t)$ for any $t \in \mathcal{T}$.

Induction Hypothesis. Assume that, at $[\![*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$ if and only if there is some $t \in \Psi(\mathcal{T})$ such that $\mathsf{MHolds}(s, t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$ and $\beta \vdash \mathsf{MHolds}(s, t)$.

Induction Step. Consider the situation at $[\![*\mathsf{NOW}_{i+1}]\!]$. Similar to the proof of Lemma 12.1, we break down the proof into four parts.

(a). Suppose that $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$ at both $[\![*\mathsf{NOW}_{i+1}]\!]$ and $[\![*\mathsf{NOW}_i]\!]$. By the induction hypothesis, at $[\![*\mathsf{NOW}_i]\!]$, $\mathsf{MHolds}(s, t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$ and $\beta \vdash \mathsf{MHolds}(s, t)$, for some $t \in \Psi(\mathcal{T})$. Thus, we only need to show that the proposition $\mathsf{MHolds}(s, t)$ remains a member of $*\mathrm{Mod}_{\mathrm{per}}(s)$ over $[\![*\mathsf{NOW}_{i+1}]\!]$. The value of NOW can change from $*\mathsf{NOW}_i$ to $*\mathsf{NOW}_{i+1}$ only if one of algorithms `state_change`, `assert_start`, or `assert_cease` is executed. Since `assert_start` and `assert_cease` do not involve any steps that set perception modality variables, then if NOW moves due to the execution of either algorithm, $\mathrm{Mod}_{\mathrm{per}}(s)$ would not be changed. Now, suppose that NOW moves due to the execution of algorithm `state_change`. The only place in `state_change` where propositions get removed from perception modality variables is step 5, through the execution of algorithm `cease_perceive`. But since $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$ at both $[\![*\mathsf{NOW}_{i+1}]\!]$ and $[\![*\mathsf{NOW}_i]\!]$, then, by pre-condition 4 of `state_change`, $s \notin S^{\downarrow}$. Thus, algorithm `cease_perceive` does not get applied to s, and the proposition $\mathsf{MHolds}(s, t)$ never gets removed from $\mathrm{Mod}_{\mathrm{per}}(s)$ by algorithm `state_change`. Therefore, at $[\![*\mathsf{NOW}_{i+1}]\!]$, following the execution of `state_change`, $\mathsf{MHolds}(s, t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$ and $\beta \vdash \mathsf{MHolds}(s, t)$.

(b). Suppose that, at $[\![*\mathsf{NOW}_{i+1}]\!]$, but not at $[\![*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$. By the induction hypothesis, at $[\![*\mathsf{NOW}_i]\!]$, there is no $t \in \mathcal{T}$ such that $\mathsf{MHolds}(s, t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$. Thus, by Axiom 8.4, s satisfies pre-conditions 1 and 2 of algorithm `state_change`, and the algorithm gets initiated with $s \in S^{\uparrow}$. By the proof of Theorem 8.1, at $[\![*\mathsf{NOW}_{i+1}]\!]$, $\mathsf{MHolds}(s, t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$, for some $t \in \Psi(\mathcal{T})$.

(c). Suppose that, at $[\![*\mathsf{NOW}_i]\!]$, but not at $[\![*\mathsf{NOW}_{i+1}]\!]$, $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$. By Axiom 8.4, at $[\![*\mathsf{NOW}_{i+1}]\!]$, there is no $\mu \in \mathcal{M}_{\mathrm{per}}$ such that $[\![s]\!]$ is perceived via the modality corresponding to μ, and, by Axiom 8.5, $[\![s]\!]$ does not occupy any proprioception modality. Given the induction hypothesis, at $[\![*\mathsf{NOW}_i]\!]$, $\mathsf{MHolds}(s, t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$, for some $t \in \Psi(\mathcal{T})$. Therefore, s satisfies pre-conditions 3 and 4 of algorithm `state_change`, and the algorithm gets initiated with $s \in S^{\downarrow}$. By Theorem 8.2, at $[\![*\mathsf{NOW}_{i+1}]\!]$, after executing the algorithm, $\mathsf{MHolds}(s, t) \notin *\mathrm{Mod}_{\mathrm{per}}(s)$, for any $t \in \mathcal{T}$.

(d). Suppose that, at neither $[\![*\mathsf{NOW}_i]\!]$ nor $[\![*\mathsf{NOW}_{i+1}]\!]$, is $[\![s]\!]$ perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$. By the induction hypothesis, at $[\![*\mathsf{NOW}_i]\!]$, $\mathsf{MHolds}(s, t) \notin *\mathrm{Mod}_{\mathrm{per}}(s)$, for any $t \in \mathcal{T}$. Therefore, we only need to show that $*\mathrm{Mod}_{\mathrm{per}}(s)$ continues to not include any propositions concerning s over $[\![*\mathsf{NOW}_{i+1}]\!]$.

The value of NOW can change from $^*\mathsf{NOW}_i$ to $^*\mathsf{NOW}_{i+1}$ only if one of algorithms `state_change`, `assert_start`, or `assert_cease` is executed. Since `assert_start` and `assert_cease` do not involve any steps that set perception modality variables, then if NOW moves due to the execution of either algorithm, $\mathrm{Mod}_{\mathrm{per}}(s)$ would not be changed. Now, suppose that NOW moves due to the execution of algorithm `state_change`. Propositions may only be added to perception modality variables through the execution of algorithm `state_perceive`. In algorithm `state_change`, this may happen in steps 3, 4b, and 4c. However, since s violates pre-condition 2 of the algorithm, $s \notin S^{\uparrow}$, and `state_perceive` never gets applied to it. Therefore, at $[\![^*\mathsf{NOW}_{i+1}]\!]$, following the execution of `state_change`, $\mathsf{MHolds}(s,t) \notin *\mathrm{Mod}_{\mathrm{per}}(s)$, for any $t \in \mathcal{T}$.

From (a), (b), (c), and (d) the induction step follows. Since s is arbitrary, the lemma follows. $\square$

Given the above result, we can now prove a theorem corresponding to Theorem 12.1.

Theorem 12.3 *For every $i \in \mathbb{N}$ $(i > 0)$ and $s \in$ TEMP such that $\mathrm{Mod}_{\mathrm{per}}(s)$ is defined, if, at $[\![^*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$, then $\beta \vdash \mathsf{Holds}(s,^*\mathsf{NOW}_i)$.*

Proof. Pick some $i \in \mathbb{N}$ $(i > 0)$ and $s \in$TEMP such that $\mathrm{Mod}_{\mathrm{per}}(s)$ is defined. Suppose that, at $[\![^*\mathsf{NOW}_i]\!]$, $[\![s]\!]$ is perceived via the modality corresponding to $\mathrm{Mod}_{\mathrm{per}}(s)$. By Lemma 12.2, at $[\![^*\mathsf{NOW}_i]\!]$, $\mathsf{MHolds}(s,t) \in *\mathrm{Mod}_{\mathrm{per}}(s)$ and $\beta \vdash \mathsf{MHolds}(s,t)$, for some $t \in \Psi(\mathcal{T})$. Steps 4 through 6 of algorithm `setup_new_MTF` result in adding the proposition $p =^*\mathsf{NOW}_i \sqsubset t$ to β. Given p, **AS2**, and **AS3**, $\beta \vdash \mathsf{Holds}(s,^*\mathsf{NOW}_i)$. Since i and s are arbitrary, then the result applies to all $i \in \mathbb{N}$ $(i > 0)$ and $s \in$TEMP such that $\mathrm{Mod}_{\mathrm{per}}(s)$ is defined. $\square$

Similar to Theorem 12.1, the above theorem is a completeness-like results for perceivable states. However, unlike with Theorem 12.2, we cannot prove a corresponding soundness-like result. Note that a soundness result would state that, whenever Cassie believes that a perceivable state holds, then she perceives it. Evidently, such a result cannot be justified on any empirical grounds; Cassie may be told that a perceivable state holds even though she does not perceive it. Unlike the case with bodily states, one cannot argue that Cassie should not believe any such assertions if she does not actually perceive the state. For example, there is nothing wrong in Cassie's believing Stu's assertion that the walk-light is on when she is not looking towards the walk-light.

12.4 Persistence Through Time-Perception

Theorems 12.1 and 12.3 have illustrated that Cassie may syntactically infer that a state holds based on the non-syntactic, PML phenomena of perception and proprioception. As pointed out in Section 11, Cassie may also believe that a state continues to hold based on her knowledge of the typical duration of the state and her sense of how much time has passed. The latter is another PML phenomenon that falls outside the bounds of the logical theory. In this section, we prove that, in the appropriate circumstances, Cassie's feel for how much time has passed results in the appropriate beliefs about the persistence of states.

Theorem 12.4 *For every* $t \in \Psi(\mathcal{T})$, $s \in \Psi(\textsc{Temp})$, $q \in \Psi(\mathcal{Q})$, *and* $i, n \in \mathbb{N}$ *(*$i > 0$*), if*

1. $\beta \vdash \mathsf{MHolds}(s,t)$,
2. $\beta \vdash \mathsf{SDur}(s,q)$,
3. $\mathcal{A}(q)$ *is defined,*
4. *for every* $\Phi \in \mathrm{Span}(t)$, Φ *is a* NOW*-MTF,*
5. $\Phi(*\mathsf{NOW}_i)$ *is the smallest element of the poset* $\langle \mathrm{Span}(t), \mathrm{precedes} \rangle$,
6. $\eta(d_t(\Phi(*\mathsf{NOW}_i), \Phi(*\mathsf{NOW}_{i+n}))) \leq \mathcal{A}(q)$, *and*
7. *for every* $m \in \mathbb{N}$, $0 < m \leq n$, *Cassie does not determine that* s *ceases to hold at* $*\mathsf{NOW}_{i+m}$, *then*
8. $\Phi(*\mathsf{NOW}_{i+n}) \in \mathrm{Span}(t)$.

The reader should note that the statement of the theorem is similar to the statement of Lemma 10.1; the proof closely follows that of the lemma.

Proof. We use induction on n to prove the lemma.

Basis. Let $n = 0$. Given statement 5, trivially, $\Phi(*\mathsf{NOW}_i) \in \mathrm{Span}(t)$.

Induction Hypothesis. Assume that, for every $t \in \Psi(\mathcal{T})$, $s \in \Psi(\textsc{Temp})$, $q \in \Psi(\mathcal{Q})$, and $i \in \mathbb{N}$ ($i > 0$), the conjunction of statements 1 through 7 implies statement 8, for some $n \in \mathbb{N}$.

Induction Step. We need to show, for $n + 1$, that the conjunction of statements 1 through 7 implies statement 8. Given statement 6,

$$\eta(d_t(\Phi(*\mathsf{NOW}_i), \Phi(*\mathsf{NOW}_{i+n+1}))) \leq \mathcal{A}(q).$$

By Corollary 10.1,

$$\eta(d_t(\Phi(^*\mathsf{NOW}_i), \Phi(^*\mathsf{NOW}_{i+n}))) \leq \mathcal{A}(q),$$

which, by the induction hypothesis, implies that $\Phi(^*\mathsf{NOW}_{i+n}) \in \mathrm{Span}(t)$. By Definition 7.8 (Span), $t \in \Phi(^*\mathsf{NOW}_{i+n})$. As NOW moves from $^*\mathsf{NOW}_{i+n}$ to $^*\mathsf{NOW}_{i+n+1}$, algorithm `setup_new_MTF` gets executed. Since t is a situation interval in $\Phi(^*\mathsf{NOW}_{i+n})$, then we are only concerned with steps 11 through 15 of the algorithm.

If $\beta \vdash ^*\mathsf{NOW}_{i+n+1} \sqsubset t$, then the conditional of step 12 is satisfied, and the proposition $^*\mathsf{NOW}_{i+n+1} \sqsubset t$ is added to β. It follows by Definitions 7.4 (MTFs) and 7.8 that $\Phi(^*\mathsf{NOW}_{i+n+1}) \in \mathrm{Span}(t)$.

On the other hand, if $\beta \not\vdash ^*\mathsf{NOW}_{i+n+1} \sqsubset t$, then control flows to step 13 of `setup_new_MTF`. Since $t \in \Phi(^*\mathsf{NOW}_{i+n})$, then $\beta \vdash ^*\mathsf{NOW}_{i+n} \sqsubset t$. Thus, by **AS2** and **AS3**, $\beta \vdash \mathsf{Holds}(s, ^*\mathsf{NOW}_{i+n})$. Given statement 7, Cassie does not determine that s ceases to hold at $^*\mathsf{NOW}_{i+n+1}$. Therefore, by Definition 8.5, $\beta \not\vdash ^*\mathsf{NOW}_{i+n+1} \not\sqsubset t$. Thus, the conditional in step 13 of algorithm `setup_new_MTF` fails and control flows to step 14. Following the same reasoning in the proof of Lemma 10.1, the conditional in 14 is true with respect to t, and the proposition $^*\mathsf{NOW}_{i+n+1} \sqsubset t$ gets added to β. By Definitions 7.4 and 7.8, $\Phi(^*\mathsf{NOW}_{i+n+1}) \in \mathrm{Span}(t)$. □

12.5 Past States

We have shown that, at any time, if it is reasonable for Cassie to believe that some state holds (based on PML phenomena), then she indeed believes (at least implicitly) that it holds. In this section, we show that, at any time, Cassie's beliefs also reasonably reflect that the event of a state holding has moved into the past.

First, we prove the following post-condition of algorithm `setup_new_MTF`.

Lemma 12.3 *For every $t \in \Psi(\mathcal{T})$, $s \in \Psi(\textsc{Temp})$, $q \in \Psi(\mathcal{Q})$, and $i \in \mathbb{N}$ $(i > 1)$, if*

1. $\beta \vdash \mathsf{MHolds}(s, t)$,
2. $^*\mathsf{NOW}_{i-1} \sqsubset t \in \beta$,
3. $\beta \vdash \mathsf{SDur}(s, q)$,
4. *$\mathcal{A}(q)$ is defined, and*
5. *for every $\Phi \in \mathrm{Span}(t)$, Φ is a* NOW-*MTF, then*
6. *following the $i - 1^{st}$ execution of algorithm* `setup_new_MTF`, $\beta \vdash ^*\mathsf{NOW}_i \sqsubset t$ *or* $\beta \vdash ^*\mathsf{NOW}_i \not\sqsubset t$.

Proof. Consider the $i-1^{st}$ execution of algorithm `setup_new_MTF`. Note that this is the execution responsible for the transition from $*\mathsf{NOW}_{i-1}$ to $*\mathsf{NOW}_i$. By statement 2 and Definition 7.4 (MTFs), $t \in \Phi(*\mathsf{NOW}_{i-1})$ and, by statement 1, t is a situation interval. Thus, we only need to consider steps 11 through 15 of the algorithm. If, at step 12, $\beta \vdash *\mathsf{NOW}_i \sqsubset t$, then the algorithm ends, and, trivially, $\beta \vdash *\mathsf{NOW}_i \sqsubset t$ following its execution. On the other hand, if the conditional of step 12 is not true, then step 13 is executed. If, at step 13, $\beta \vdash *\mathsf{NOW}_i \not\sqsubset t$, then, trivially, $\beta \vdash *\mathsf{NOW}_i \not\sqsubset t$ following the execution of the algorithm. Otherwise, step 14 gets executed. By statement 5, $\pi(t)$ is defined. Given statement 4, the conditional in 14 may be evaluated. If $\eta(\pi(t)) \leq \mathcal{A}(q)$, then step 14 guarantees that, following the execution of the algorithm, $\beta \vdash *\mathsf{NOW}_i \sqsubset t$. Otherwise, the proposition $t \prec \mathsf{NOW}_i$ gets added to β by step 15. By **TT1**, following the execution of the algorithm, $\beta \vdash *\mathsf{NOW}_i \not\sqsubset t$. $\square$

We now prove the main result.

Theorem 12.5 *For every $t \in \Psi(\mathcal{T})$, $s \in \Psi(\text{TEMP})$, $q \in \Psi(\mathcal{Q})$, and $i \in \mathbb{N}$ $(i > 1)$, if*

1. $\beta \vdash \mathsf{MHolds}(s, t)$,
2. $*\mathsf{NOW}_{i-1} \sqsubset t \in \beta$.
3. $\beta \vdash \mathsf{SDur}(s, q)$.
4. *$\mathcal{A}(q)$ is defined.*
5. *for every $\Phi \in \mathrm{Span}(t)$, Φ is a* NOW*-MTF, and*
6. *Cassie determines that s ceases to hold at $*\mathsf{NOW}_i$, then*
7. *for every $n \in \mathbb{N}$, $\beta \vdash t \prec *\mathsf{NOW}_{i+n}$.*

Proof. By Lemma 12.3, at $[\![*\mathsf{NOW}_i]\!]$ (following the $i-1^{st}$ execution of algorithm `setup_new_MTF`), either $\beta \vdash *\mathsf{NOW}_i \sqsubset t$ or $\beta \vdash *\mathsf{NOW}_i \not\sqsubset t$. Given statement 6 and Definition 8.5, the former cannot be the case.[46] Therefore, following the $i-1^{st}$ execution of algorithm `setup_new_MTF`, $\beta \vdash *\mathsf{NOW}_i \not\sqsubset t$. Following the proof of Lemma 12.3, this may happen either through step 13 or step 15. In either case, the proposition $t \prec *\mathsf{NOW}_i$ gets added to β. By Theorem 7.4 and the transitivity of $\prec$ (**AT2**), for every $n \in \mathbb{N}$, $\beta \vdash t \prec *\mathsf{NOW}_{i+n}$. $\square$

[46] Recall that we are dismissing the possibility of contradictory information about current states (see Section 8.1). If this were not the case, then statement 6 may still be accommodated if belief revision algorithms choose to retract the proposition $*\mathsf{NOW}_i \sqsubset t$ (especially if introduced as a default assumption by step 14 of `setup_new_MTF`) in favor of a newly derived/asserted contradictory proposition.

13 Conclusions and Implications

So what are the results of our investigation? Two main results, a couple of implications for research on temporal reasoning, and lots of details about the epistemology of time. First, let us go over the two main results. The primary result of this investigation is providing a solution to the problem of the fleeting now: a problem that faces agents interleaving reasoning and acting while maintaining a sense of the present time. Basically, the agent is interested in whether some state holds "now". However, since reasoning and sensory acts take time, whatever conclusion it makes will be, strictly speaking, about a different "now". Our solution is based on the simple intuition that the concept of "now" is vague as to the size of the interval it represents. The agent wonders whether the state holds, not at the sharp moment of experience, but over a broader "now", an interval whose duration is comparable to the typical duration of the state. Such an interval may still be "now", relative to some coarse level of granularity, at the time of the conclusion. Whether this is the case depends on the amount of time it takes the agent to reach a conclusion. To formalize these intuitions, we developed a theory of subjective time, which is the second main result of this investigation.

An agent's subjective sense of time involves two main components: a representation of "now" that continuously changes reflecting the progression of time, and a feel for how much time has passed. Due to the granular vagueness of the concept of "now", a mereological meet semi-lattice of intervals represents the agent's sense of "now" at different levels of granularities. A meta-logical variable, NOW, that assumes values from amongst the set of time-denoting terms of the logic, represents the agent's notion of "now" at the finest level of granularity; its value is the smallest element of the lattice of "now"s at any time. Values of NOW form a chain ordered by the temporal precedence relation, and the progression of time is modeled by introducing a new term at the end of the chain and making it the new value of NOW.

The agent's sense of how much time has passed is based on an internal subconscious clock, the pacemaker, that primarily gives the agent a feel for the duration of NOW-intervals. Two durations may *feel* somewhat different but, nevertheless, be indistinguishable at the level of conscious reasoning. Therefore, a quantized representation of durations was introduced to ground the amount-denoting terms of the logic. Amounts are quantized into intervals each corresponding to one half-order of magnitude (Hobbs, 2000). Both the use of half-order of magnitudes for quantization and the pacemaker for the sense of time are not central to the theory; they could be replaced by other, more suitable alternatives if needed.

The theory makes the reasonable assumption that the value of NOW changes (reflecting the progression of time) when and only when there is a detectable state change in the agent's environment. At the current state of the theory, the environment does not include states of the agent's mind; thus, pure reasoning does not take time as far as the theory is concerned. However, reasoning may take time if it involves performing some actions (mainly, sensory actions) in order to add a missing link to a chain of inference. For NOW to change, then, the agent must perceive, proprioceive, or be told about some state change. We have outlined a number of principles that govern state change as determined by each of these means (in addition to inference). Based on these principles, a collection of algorithms have been introduced to account for what happens when time passes. In addition to changing NOW, other components of the system need to be updated—most importantly, the agent's belief space, since the passage of time always involves new beliefs about the detected changes. It should also be noted that the algorithms are general enough to account for multiple simultaneous changes including onsets and/or cessations of various states.

In addition to updating NOW and the agent's belief space, a set of modality registers are updated whenever a change is determined through perception or proprioception. Modality registers contain propositions representing what each agent modality is being used for or what it perceptually conveys about the environment. The main utility of these registers is in the smooth, reasoning-free projection of bodily and perceived states as time passes by. We have outlined an algorithm (`setup_new_MTF`) that is invoked every time NOW changes and that is responsible for extending continuing states over the new NOW. The algorithm makes use of the agent's knowledge of how much time has passed and typical durations of states in order to defeasibly determine whether a state continues to persist. A number of results were presented demonstrating the agent's *temporal rationality*—that its beliefs about whether a state holds "now", or whether it has moved into the past, are justified by what it knows, what it perceives, what it proprioceives, and what it feels regarding how much time has passed.

With respect to typical durations and issues of persistence, the theory presented here raises a number of questions to psychologists of time:

1. What kinds of mental representations and processes are involved in reasoning about typical durations of states?

2. As regards the backward projection factor, what are the biases of human subjects as to how long a perceived state has been holding?

3. What are the factors determining those biases?

The results of our investigation indeed go beyond a solution to the problem of the fleeting now. For example, an agent that knows the typical durations required by certain activities to achieve their goals, may use its sense of time to consciously detect failure. In fact, the use of durations in planning and plan execution has been used at a "subconscious" level in various AI systems (Vere, 1983; Ambros-Ingerson and Steel, 1988, for example). In general, using its sense of time, the agent may impose certain constraints on when it should achieve some goals, and an action may be determined by the agent to have failed if it continues past a dead-line. How to effectively integrate the sense of time with action execution is worth further investigation.

Future research may also proceed in other directions. For example, reasoning about durations as presented here is confined to reasoning about their order with respect to the relation $<_Q$. This is sufficient for the purposes of this paper. However, representing and reasoning about durations have a broader range of applications in temporal reasoning (Allen and Kautz, 1988, for example). For example, consider an agent that knows that a bus arrives every 10 minutes and waits at the bus stop for 3 minutes. Given that a bus was at the bus stop at some time, t_1, was there a bus at t_2? In general, to answer this question, the agent needs to know the amount of time separating t_1 and t_2. Such knowledge may be explicitly conveyed to the agent by direct assertion. But, in general, it need not. For example, the agent might know that there is some t_3 such that $t_1 \prec t_3 \prec t_2$, and that the amount of time between t_1 and t_3 is 5 minutes and that between t_3 and t_2 is also 5 minutes. In such a situation, the agent should be capable of answering the question. This may only be achieved, however, if the agent can engage in some arithmetical reasoning about durations—something that is not accommodated by our theory.

Even if one were to dismiss the kind of reasoning required to cope with the situation presented above, consider a case where both t_1 and t_2 are NOW-intervals. In such a case, the duration between t_1 and t_2 is given by $d_t(\Phi(t_1), \Phi(t_2))$ (see Definition 10.2). This quantity exists at the PML and may be efficiently computed there without the need for any arithmetical reasoning. What is needed then is to allow such a computation as part of the reasoning process. This may indeed be achieved if one thinks of performing this computation as a sensory act that may be executed to add a missing link to a chain of reasoning. How this exactly is to be done requires further investigation.

14 Acknowledgements

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. Comments by William J. Rapaport and Frances L. Johnson on earlier drafts are highly appreciated. The authors also appreciate the valuable comments of three anonymous reviewers for IJCAI-01 and two anonymous reviewers of Commonsense 2001 on some of the ideas presented in this paper.

This work was supported in part by ONR under contract N00014-98-C-0062.

A Temporal Progression Algorithms

Algorithm $\eta(n \in \mathbb{N})$

1. If $n = 0$ then, return 0.
2. Return $1 + \text{round}(\log_{\sqrt{10}}(n))$.

Algorithm $\rho(n \in \mathbb{N})$

1. $h \longleftarrow \eta(n)$.
2. If there is $q \in \Psi(\mathcal{Q})$ such that $\mathcal{A}(q) = h$, then return q.
3. Pick some $q \in \mathcal{Q}$, such that $q \notin \Psi(\mathcal{Q})$.
4. $\mathcal{A} \longleftarrow \mathcal{A} \cup \{\langle q, h\rangle\}$
5. $min \longleftarrow \{q' | \mathcal{A}(q') = h' \wedge h' < h\}$.
6. $max \longleftarrow \{q' | \mathcal{A}(q') = h' \wedge h < h']\}$.
7. If min is not empty, then $\beta \longleftarrow \beta \cup \{q_{gmin} <_{\mathcal{Q}} q\}$, where q_{gmin} is the greatest element of the linearly-ordered poset $\langle min, <_{\mathcal{Q}}\rangle$.
8. If max is not empty, then $\beta \longleftarrow \beta \cup \{q <_{\mathcal{Q}} q_{lmax}\}$, where q_{lmax} is the smallest element of the linearly-ordered poset $\langle max, <_{\mathcal{Q}}\rangle$.
9. return q.

Algorithm `initialize_NOW`

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\mathsf{COUNT} \longleftarrow 0$.
3. $\mathsf{NOW} \longleftarrow t$.

Algorithm `move_NOW`

1. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
2. $\beta \longleftarrow \beta \cup \{\textsf{*NOW} \prec t, \textsf{Dur}(\textsf{*NOW}, \rho(\textsf{*COUNT}))\}$.
3. $\mathcal{A} \longleftarrow \mathcal{A} \cup \{\langle\textsf{*NOW},\textsf{*COUNT}\rangle\}$.
4. $\textsf{COUNT} \longleftarrow 0$.
5. $\textsf{NOW} \longleftarrow t$.

Algorithm `setup_new_MTF`$(S^{\downarrow} \subseteq \Psi(\textsc{Temp}))$

1. `move_NOW`
2. For all $\mu \in \mathcal{M}_{\text{prop}}$
   3. If there are s and t such that $*\mu = \textsf{MHolds}(s,t)$, then $\beta \longleftarrow \beta \cup \{\textsf{*NOW} \sqsubset t\}$.
4. For all $\mu \in \mathcal{M}_{\text{per}}$
   5. For all s and t such that $\textsf{MHolds}(s,t) \in *\mu$
      6. $\beta \longleftarrow \beta \cup \{\textsf{*NOW} \sqsubset t\}$.
7. For all $t \in \lceil \textsf{*NOW}_{i-1} \rceil \setminus \{\textsf{*NOW}_{i-1}\}$
   8. If $\eta(\pi(t)) \leq \mathcal{A}(\delta(t))$, then $\beta \longleftarrow \beta \cup \{\textsf{*NOW} \sqsubset t\}$.
   9. Else $\beta \longleftarrow \beta \cup \{t \prec \textsf{*NOW}\}$.
10. For all $s \in S^{\downarrow}$ $\beta \longleftarrow \beta \cup \{\neg\textsf{Holds}(s,\textsf{*NOW})\}$.
11. For all $t \in \Phi(\textsf{*NOW}_{i-1}) \setminus \lceil \textsf{*NOW}_{i-1} \rceil$,
    12. If $\beta \vdash \textsf{*NOW} \sqsubset t$, then $\beta \longleftarrow \beta \cup \{\textsf{*NOW} \sqsubset t\}$.
    13. Else, if $\beta \vdash \textsf{*NOW} \not\sqsubset t$, then $\beta \longleftarrow \beta \cup \{t \prec \textsf{*NOW}\}$.
    14. Else, if $\eta(\pi(t)) \leq \mathcal{A}(q)$, where $\textsf{SDur}(s,q)$ and $\textsf{MHolds}(s,t)$, then $\beta \longleftarrow \beta \cup \{\textsf{*NOW} \sqsubset t\}$.
    15. Else $\beta \longleftarrow \beta \cup \{t \prec \textsf{*NOW}\}$.

Algorithm `state_present`(s,t)

1. $q' \longleftarrow \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$, where $\beta \vdash \mathsf{SDur}(s,q)$.
2. For every $t' \in \lceil \mathsf{*NOW} \rceil$
   3. If $\beta \nvdash \neg\mathsf{Holds}(s,t')$
      and $\delta(t') = q'$ or $\beta \vdash \delta(t') <_{\mathcal{Q}} q'$,
      then $\beta \longleftarrow \beta \cup \{t' \sqsubset t\}$.

Algorithm `state_query`(s)

1. $q' \longleftarrow \rho(\sqrt{10}^{bpf(s)(\mathcal{A}(q)-1)})$, where $\beta \vdash \mathsf{SDur}(s,q)$.
2. Pick some $t \in \mathcal{T}$, such that $t \notin \Psi(\mathcal{T})$.
3. $\beta \longleftarrow \beta \cup \{\mathsf{*NOW} \sqsubset t, \mathsf{Dur}(t,q')\}$.
4. Initiate deduction for $\mathsf{Holds}(s,t)$.

Algorithm `state_change`$(S^{\uparrow} \subseteq \textsc{Temp}, S^{\downarrow} \subseteq \textsc{Temp})$

1. $P_{\text{new}} \longleftarrow \{\}$.
2. For all $s_i \in S^{\uparrow}$
   3. If $\beta \vdash \mathsf{Holds}(s_i, \mathsf{*NOW})$ then `start_ceive`(s_i,t_i), where t_i is the situation interval associated with s_i such that $\beta \vdash \mathsf{*NOW} \sqsubset t_i$
   4. else
      4a. Pick some $t_i \in \mathcal{T}$, such that $t_i \notin \Psi(\mathcal{T})$.
      4b. If $\beta \vdash \neg\mathsf{Holds}(s_i, \mathsf{*NOW})$ then `state_start`(s_i,t_i).
      4c. Else, `state_persist`(s_i,t_i).
      4d. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{Mholds}(s_i,t_i)\}$.
5. For all $s_i \in S^{\downarrow}$, `cease_perceive`(s_i).
6. `setup_new_MTF`$(\{\})$.
7. `Forward`(P_{new}).

Algorithm `start_ceive`(s,t)

1. `start_proprioceive`(s,t).
2. `start_perceive`(s,t).

Algorithm `start_proprioceive`(s, t)

1. If $\mathrm{Mod}_{\mathrm{prop}}(s) \neq \{\}$, then
   for all $\mu \in \mathrm{Mod}_{\mathrm{prop}}(s)$, $\mu \longleftarrow \mathsf{MHolds}(s, t)$.

Algorithm `start_perceive`(s, t)

1. If $\mathrm{Mod}_{\mathrm{per}}(s)$ is defined, then
   $\mathrm{Mod}_{\mathrm{per}}(s) \longleftarrow *\mathrm{Mod}_{\mathrm{per}}(s) \cup \{\mathsf{MHolds}(s, t)\}$.

Algorithm `state_start`(s, t)

1. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s, t)\}$.
2. `start_ceive`(s, t).

Algorithm `state_persist`(s, t)

1. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s, t)\}$.
2. `start_perceive`(s, t).

Algorithm `cease_perceive`(s)

1. If $\mathrm{Mod}_{\mathrm{per}}(s)$ is defined, then
   $\mathrm{Mod}_{\mathrm{per}}(s) \longleftarrow *\mathrm{Mod}_{\mathrm{per}}(s) \setminus \{\mathsf{MHolds}(s, t)\}$

Algorithm `assert_persist`$(S \subseteq \textsc{Temp})$

1. $P_{\mathrm{new}} \longleftarrow \{\}$.
2. For all $s_i \in S$
   3. Pick some $t_i \in \mathcal{T}$, such that $t_i \notin \Psi(\mathcal{T})$.
   4. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s_i, t_i), {*}\mathsf{NOW} \sqsubset t_i\}$.
   5. $P_{\mathrm{new}} \longleftarrow P_{\mathrm{new}} \cup \{\mathsf{MHolds}(s_i, t_i)\}$.
6. `Forward`(P_{new}).

Algorithm `assert_start`($S \subseteq$ TEMP)

1. `setup_new_MTF`($\{\}$).
2. $P_{\text{new}} \longleftarrow \{\}$.
3. For all $s_i \in S$
   4. Pick some $t_i \in \mathcal{T}$, such that $t_i \notin \Psi(\mathcal{T})$.
   5. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s_i, t_i), \mathsf{*NOW} \sqsubset t_i\}$.
   6. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{MHolds}(s_i, t_i)\}$.
7. `Forward`(P_{new}).

Algorithm `assert_cease`($S \subseteq$ TEMP)

1. `setup_new_MTF`(S).

Algorithm `initialize`($S \in$ TEMP)

1. $P_{\text{new}} \longleftarrow \{\}$.
2. `initialize_NOW`.
3. For all $s_i \in S$
   4. Pick some $t_i \in \mathcal{T}$ such that $t_i \notin \Psi(\mathcal{T})$.
   5. `state_start`(s_i, t_i).
   6. $\beta \longleftarrow \beta \cup \{\mathsf{*NOW} \sqsubset t_i\}$.
   7. $P_{\text{new}} \longleftarrow P_{\text{new}} \cup \{\mathsf{MHolds}(s_i, t_i)\}$.
8. `Forward`(P_{new}).

Algorithm `Forward`($P \subseteq \Psi(\mathcal{P})$)

1. $P_{\text{inf}} \longleftarrow$ `Forward_old`(P).
2. For every $p \in P \cup P_{\text{inf}}$
   3. If $p = \mathsf{Holds}(s, \mathsf{*NOW})$, for some $s \in \Psi$(TEMP), then
      4. If there is some $t \in \Psi(\mathcal{T})$ such that $\beta \vdash \mathsf{MHolds}(s, t)$, then `state_present`($s, t$).
      5. Else
         5a. Pick some $t \in \mathcal{T}$ such that $t \notin \Psi(\mathcal{T})$.
         5b. $\beta \longleftarrow \beta \cup \{\mathsf{MHolds}(s, t), \mathsf{*NOW} \sqsubset t\}$.
         5c. `state_present`(s, t).

B Proofs of Theorems TT1 and TT2

We will break the proof of **TT1** down by proving two simpler lemmas.

- **LT1.** $Tv_1 \prec Tv_2 \Rightarrow \neg[Tv_2 \sqsubseteq Tv_1]$

Proof.

1. $Tv_1 \prec Tv_2$ (Assumption)
2. $Tv_2 \sqsubseteq Tv_1$ (Assumption)
3. $Tv_1 \prec Tv_2 \wedge Tv_2 \sqsubseteq Tv_1$ (1, 2, $\wedge$-introduction)
4. $[Tv_1 \prec Tv_2 \wedge Tv_2 \sqsubseteq Tv_1] \Rightarrow Tv_2 \prec Tv_2$ (**AT7**, $\forall$-elimination)
5. $Tv_2 \prec Tv_2$ (3, 4, $\Rightarrow$-elimination)
6. $\mathsf{Equiv}(Tv_2, Tv_2) \Rightarrow \neg[Tv_2 \prec Tv_2]$ (**AT3**, $\forall$-elimination)
7. $\neg[Tv_2 \prec Tv_2]$ (6, reflexivity of Equiv, $\Rightarrow$-elimination)
8. $\neg[Tv_2 \sqsubseteq Tv_1]$ (2, 5, 7, $\neg$-introduction)
9. $Tv_1 \prec Tv_2 \Rightarrow \neg[Tv_2 \sqsubseteq Tv_1]$ (1, 8, $\Rightarrow$-introduction)

Q.E.D.

- **LT2.** $Tv_1 \prec Tv_2 \Rightarrow \neg[Tv_1 \sqsubseteq Tv_2]$

Proof.

1. $Tv_1 \prec Tv_2$ (Assumption)
2. $Tv_1 \sqsubseteq Tv_2$ (Assumption)
3. $Tv_1 \prec Tv_2 \wedge Tv_1 \sqsubseteq Tv_2$ (1, 2, $\wedge$-introduction)
4. $[Tv_1 \prec Tv_2 \wedge Tv_1 \sqsubseteq Tv_2] \Rightarrow Tv_1 \prec Tv_1$ (**AT8**, $\forall$-elimination)
5. $Tv_1 \prec Tv_1$ (3, 4, $\Rightarrow$-elimination)
6. $\mathsf{Equiv}(Tv_1, Tv_1) \Rightarrow \neg[Tv_1 \prec Tv_1]$ (**AT3**, $\forall$-elimination)
7. $\neg[Tv_1 \prec Tv_1]$ (6, reflexivity of Equiv, $\Rightarrow$-elimination)

8. $\neg[Tv_1 \sqsubseteq Tv_2]$ (2, 5, 7, $\neg$-introduction)

9. $Tv_1 \prec Tv_2 \Rightarrow \neg[Tv_1 \sqsubseteq Tv_2]$ (1, 8, $\Rightarrow$-introduction)

Q.E.D.

Given **LT1** and **LT2**, **TT1** readily follows.

Similarly, we prove **TT2** by proving the two lemmas **LT3** and **LT4**.

- **LT3.** $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3] \Rightarrow Tv_2 \sqsubset Tv_1$
  Proof.

  1. $Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3$ (Assumption)
  2. $\neg[Tv_2 \sqsubset Tv_1]$ (Assumption)
  3. $Tv_2 \sqsubseteq Tv_1$ (1, $\wedge$-elimination)
  4. $\mathsf{Equiv}(Tv_2, Tv_1)$ (Follows from 2, 3, and **AT10**)
  5. $Tv_1 \sqsubseteq Tv_2$ (4, **AT6**, $\Rightarrow$-elimination)
  6. $Tv_3 \sqsubseteq Tv_1$ (1, $\wedge$-elimination)
  7. $Tv_3 \sqsubseteq Tv_1 \wedge Tv_1 \sqsubseteq Tv_2$ (5, 6, $\wedge$-introduction)
  8. $Tv_3 \sqsubseteq Tv_2$ (7, **AT5**, $\Rightarrow$-elimination)
  9. $Tv_2 \prec Tv_3$ (1, $\wedge$-elimination)
  10. $\neg[Tv_3 \sqsubseteq Tv_2]$ (9, **LT1**, $\Rightarrow$-elimination)
  11. $Tv_2 \sqsubset Tv_1$ (2, 8, 10, $\neg$-elimination)
  12. $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3] \Rightarrow Tv_2 \sqsubset Tv_1$
  (1, 11, $\Rightarrow$-introduction)

  Q.E.D.

- **LT4.** $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3] \Rightarrow Tv_3 \sqsubset Tv_1$
  Proof.

  1. $Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3$ (Assumption)
  2. $\neg[Tv_3 \sqsubset Tv_1]$ (Assumption)
  3. $Tv_3 \sqsubseteq Tv_1$ (1, $\wedge$-elimination)
  4. $\mathsf{Equiv}(Tv_3, Tv_1)$ (Follows from 2, 3, and **AT10**)
  5. $Tv_1 \sqsubseteq Tv_3$ (4, **AT6**, $\Rightarrow$-elimination)
  6. $Tv_2 \sqsubseteq Tv_1$ (1, $\wedge$-elimination)
  7. $Tv_2 \sqsubseteq Tv_1 \wedge Tv_1 \sqsubseteq Tv_3$ (5, 6, $\wedge$-introduction)

8. $Tv_2 \sqsubseteq Tv_3$ (7, **AT5**, $\Rightarrow$-elimination)

9. $Tv_2 \prec Tv_3$ (1, $\wedge$-elimination)

10. $\neg[Tv_2 \sqsubseteq Tv_3]$ (9, **LT2**, $\Rightarrow$-elimination)

11. $Tv_3 \sqsubset Tv_1$ (2, 8, 10, $\neg$-elimination)

12. $[Tv_2 \sqsubseteq Tv_1 \wedge Tv_3 \sqsubseteq Tv_1 \wedge Tv_2 \prec Tv_3] \Rightarrow Tv_3 \sqsubset Tv_1$
(1, 11, $\Rightarrow$-introduction)

Q.E.D.

TT2 follows directly from **LT3** and **LT4**.

References

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

Allen, J. (1984). Towards a general theory of action and time. *Artificial Intelligence*, 23:123–154.

Allen, J. and Hayes, P. (1985). A common-sense theory of time. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pages 528–531, Los Altos, CA. Morgan Kaufmann.

Allen, J. F. and Kautz, H. A. (1988). A model of naive temporal reasoning. In (Hobbs and Moore, 1988), pages 261–268.

Almeida, M. (1987). *Reasoning about the Temporal Structure of Narratives.* PhD thesis, Department of Computer Science, State University of New York at Buffalo. Techincal report 87–10.

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. C. (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.

Ambros-Ingerson, J. and Steel, S. (1988). Integrating planning, execution and monitoring. In *Proceedings of the 7th National Conference on Artificial Intelligence*, pages 83–88, San Mateo, CA. Morgan Kaufmann.

Artale, A. and Franconi, E. (1998). A temporal description logic for reasoning about actions and plans. *Journal of Artificial Intelligence Research*, 9:463–506.

Bacchus, F., Tenenberg, J., and Koomen, J. (1991). A non-reified temporal logic. *Artificial Intelligence*, 52:87–108.

Barwise, J. and Perry, J. (1983). *Situations and Attitudes.* The MIT Press, Cambridge, MA.

Brown, F. M., editor (1987). *The Frame Problem in Artificial Intelligence : Proceedings of the 1987 Workshop*, Los Altos, CA. Morgan Kaufmann.

Chisholm, R. M. (1990). Events without times: An essay on ontology. *Noûs*, 24:413–428.

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

Dean, T. and Kanazawa, K. (1988). Probabilistic temporal reasoning. In *Proceedings of the 7th National Conference on Artificial Intelligence*, pages 524–528, Los Altos, CA. Morgan Kaufmann.

Dix, J., Kraus, S., and Subrahmanian, V. (2001). Temporal agent programs. *Artificial Intelligence*, 127(1):87–135.

Dowden, B. (1998). What is essential to being a clock. In Fieser, J., editor, *Internet Encyclopedia of Philosophy <http://www.utm.edu/research/iep/t/time>*.

Elgot-Drapkin, J. J. and Perlis, D. (1990). Reasoning situated in time I: Basic concepts. *Journal of Experimental and Theoretical Artificial Intelligence*, 2:75–98.

Eschenbach, C., Habel, C., and Smith, B., editors (1994). *Topological Foundations of Cognitive Science.* Papers from the Workshop at the FISI-CS, Buffalo, NY.

Forbus, K. (1984). Qualitative process theory. *Artificial Intelligence*, 24:85–168.

Friedman, W. (1990). *About Time: Inventing the Fourth Dimension.* MIT Press, Cambridge, MA.

Gabbay, D., Hogger, C. J., and Robinson, J. A., editors (1995). *Epistemic and Temporal Reasoning*, volume 4 of *Handbook of Logic in Artificial Intelligence and Logic Programming.* Clarendon Press, Oxford.

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

Galton, A. (1990). A critical examination of Allen's theory of action and time. *Artificial Intelligence*, 42:159–188.

Galton, A. (1995). Time and change for AI. In (Gabbay et al., 1995), pages 175–240.

Giunchiglia, E. and Lifschitz, V. (1995). Dependent fluents. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, volume 2, pages 1964–1969, San Mateo, CA. Morgan Kaufmann.

Gurney, J., Perlis, D., and Purang, K. (1997). Interpreting presuppositions using active logic: From contexts to utterances. *Computational Intelligence*, 5(4).

Habel, C. (1994). Discreteness, finiteness, and the structure of topological spaces. In (Eschenbach et al., 1994).

Hanks, S. and McDermott, D. (1987). Nonmonotonic logic and temporal projection. *Artificial Intelligence*, 33:379–412.

Hayes, P. (1985). The second naive physics manifesto. In Hobbs, J. and Moore, R., editors, *Formal Theories of the Commonsense World*, pages 1–36. Ablex Publishing Corporation, Norwood, NJ.

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.

Hobbs, J. (2000). Half orders of magnitude. In Obrst, L. and Mani, I., editors, *Papers from the Workshop on Semantic Approximation, Granularity, and Vagueness*, pages 28–38. A Workshop of the Seventh International Conference on Principles of Knowledge Representation and Reasoning.

Hobbs, J. and Moore, R., editors (1988). *Formal Theories of the Commonsense World.* Ablex Publishing Corporation, Norwood, NJ.

Huntley, M. (1984). The semantics of english imperatives. *Linguistics and Philosophy*, 7:103–133.

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

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

Johnson, F. L. and Shapiro, S. C. (2000a). Automatic belief revision in SNePS. In Baral, C. and Truszczynski, M., editors, *Proceedings of the 8th International Workshop on Non-Monotonic Reasoning (NMR2000)*. Unpaginated, 5 pages.

Johnson, F. L. and Shapiro, S. C. (2000b). Formalizing a deductively open belief space. Technical Report 2000-02, Department of Computer Science and Engineering, State University of New York at Buffalo.

Johnson, F. L. and Shapiro, S. C. (2000c). Implementing integrity constraints in an existing belief revision system. In Baral, C. and Truszczynski, M., editors, *Proceedings of the 8th International Workshop on Non-Monotonic Reasoning (NMR2000)*. Unpaginated, 8 pages.

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

Kamp, H. (1979). Some remarks on the logic of change. Part I. In Rohrer, C., editor, *Time, Tense, and Quantifiers*, pages 135–174. Niemeyer, Tübingen.

Kanazawa, K. (1992). *Reasoning about Time and Probability*. PhD thesis, Brown University, Department of Computer Science, Providence, RI. Technical Report CS-92-61.

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

Kautz, H. (1986). The logic of persistence. In *Proceedings of the 5th National Conference on Artificial Intelligence*, pages 401–405, Los Altos, CA. Morgan Kaufmann.

Krifka, M. (1989). Nominal reference, temporal constitution and quantification in event semantics. In Bartsch, R., van Benthem, J., and van Emde Boas, P., editors, *Semantics and Contextual Expression*, pages 75–116. Foris Publications, Dordrecht, Holland.

Ladkin, P. (1986). Primitives and units for time specification. In *Proceedings of the 5th National Conference on Artificial Intelligence*, pages 354–359, Los Altos, CA. Morgan Kaufmann.

Lakoff, G. and Johnson, M. (1999). *Philosophy in the Flesh: The Embodied Mind and its Challenge to Western Thought*. Basic Books, New York, NY.

Lespérance, Y. and Levesque, H. (1994). An argument for indexical representations in temporal reasoning. In Elio, R., editor, *Proceedings of the Tenth Biennial Conference of the Canadian Society for Computational Studies of Intelligence*, pages 271–277, Banff, Canada.

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

Levin, I. and Zakay, D., editors (1989). *Time and Human Cognition: A Life-Span Perspective*. Number 59 in Advances in Psychology. North-Holland, Amsterdam.

Lifschtiz, V. and Rabinov, A. (1989). Things that change by themselves. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 864–867, Los Altos, CA. Morgan Kaufmann.

Link, G. (1998). *Algebraic Semantics in Language and Philosophy*. CSLI Publications, Stanford, CA.

Maida, A. S. and Shapiro, S. C. (1982). Intensional concepts in propositional semantic networks. *Cognitive Science*, 6:291–330.

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.

McTaggart, J. E. (1908). The unreality of time. *Mind*, 68:457–474.

Millikan, R. G. (1990). The myth of the essential indexial. *Noûs*, 24(5):723–734.

Moore, R. C. (1984). Possible-world semantics for autoepistemic logic. In *Proceedings of the 1984 Non-monotonic Reasoning Workshop*, pages 344–354, New Paltz, NY. American Association for Artificial Intelligence.

Moore, R. C. (1985). Semantical considerations on nonmonotonic logic. *Artificial Intelligence*, 25:75–94.

Morgenstern, L. (1996). The problem with solutions to the frame problem. In Ford, K. and Pylyshyn, Z., editors, *The Robot's Dilemma Revisited: The Frame Problem in Artificial Intelligence*. Ablex, Norwood.

Myers, K. and Smith, D. (1988). The persistence of derived information. In *Proceedings of the 7th National Conference on Artificial Intelligence*, pages 496–500, Los Altos, CA. Morgan Kaufmann.

Nirkhe, M., Kraus, S., Miller, M., and Perlis, D. (1997). How to (plan to) meet a deadline between *now* and *then*. *Journal of Logic and Computation*, 7(1):109–156.

Pastor, M. and Artieda, J., editors (1996). *Time, Internal Clocks and Movement*. Number 115 in Advances in Psychology. Elsevier, Amsterdam.

Perlis, D., Elgot-Drapkin, J. J., and Miller, M. (1991). Stop the world—I want to think. *International Journal of Intelligent Systems*, 6:443–456.

Perry, J. (1979). The problem of the essential indexical. *Noûs*, 13(1):3–21.

Pianesi, F. and Varzi, A. (1996). Events, topology, and temporal relations. *The Monist*, 79(1):89–116.

Pinto, J. A. (1994). *Temporal Reasoning in the Situation Calculus*. PhD thesis, Department of Computer Science, University of Toronto, Toronto, ON.

Poynter, D. (1989). Judging the duration of time intervals: A process of remembering segments of experience. In (Levin and Zakay, 1989), pages 305–331.

Prior, A. (1967). *Past, Present and Future*. Clarendon Press, Oxford.

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

Pylyshyn, Z., editor (1986). *The Robot's Dilemma : The Frame Problem in Artificial Intelligence*. Ablex, Norwood, NJ.

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

Rapaport, W. J. (1978). Meinongian theories and a Russellian paradox. *Noûs*, 12:153–180.

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

Reiter, R. (1991). The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In Lifschitz, V., editor, *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, pages 359–380. Academic Press, San Diego, CA.

Reiter, R. (1998). Sequential, temporal GOLOG. In Cohn, A. G., Schubert, L. K., and Shapiro, S. C., editors, *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR '98)*, pages 547–556, San Francisco, CA. Morgan Kaufmann.

Rescher, N. and Garson, J. (1968). Topological logic. *The Journal of Symbolic Logic*, 33(4):537–548.

Rumelhart, D., Lindsay, P., and Norman, D. (1972). A process model for long-term memory. In (Tulving and Donaldson, 1972), pages 198–248.

Schmiedel, A. (1990). A temporal terminological logic. In *Proceedings of the 9th National Conference on Artificial Intelligence*, pages 640–645, San Mateo, CA. Morgan Kaufmann.

Shanahan, M. (1997). *Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia*. MIT Press, Cambridge, MA.

Shapiro, S. C. (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. C. (1991). Cables, paths, and "subconscious" reasoning. In Sowa, J. F., editor, *Principles of Semantic Networks: Explorations in the Representation of knowledge*, pages 137–156. Morgan Kaufmann, San Mateo, CA.

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

Shapiro, S. C. (1996). Formalizing English. *International Journal of Expert Systems*, 9(1):151–171.

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

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 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. C. and Rapaport, W. J. (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. C. and the SNePS Implementation Group (1999). *SNePS 2.5 User's Manual.* Department of Computer Science and Engineering, State University of New York at Buffalo.

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

Shoham, Y. (1986). Chronological ignorance: Time, nonmonotonicity, necessity and causal theories. In *Proceedings of the 5th National Conference on Artificial Intelligence*, pages 389–393, Los Altos, CA. Morgan Kaufmann.

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

Shoham, Y. and McDermott, D. (1988). Problems in formal temporal reasoning. *Artificial Intelligence*, 36:49–61.

Smith, B. C. (1988). The semantics of clocks. In Fetzer, J. H., editor, *Aspects of Artificial Intelligence*, pages 3–31. Kluwer Academic Publishers, Dordrecht.

Stein, L. A. (1990). An atemporal frame problem. *International Journal of Expert Systems*, 3(4):371–381.

Talmy, L. (1988). Force dynamics in language and cognition. *Cognitive Science*, 12:49–100.

Talmy, L. (2000). *Toward a Cognitive Semantics.* MIT Press.

ter Meulen, A. (1997). *Representing Time in Natural Language.* The MIT Press, Cambridge, MA, second edition.

Thielscher, M. (1999). From situation calculus to fluent calculus: State update axioms as a solution to the inferential frame problem. *Artificial Intelligence*, 111:277–299.

Tulving, E. (1972). Episodic and semantic memory. In (Tulving and Donaldson, 1972), pages 382–404.

Tulving, E. and Donaldson, W., editors (1972). *Organization of Memory*. Academic Press, New York, NY.

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

van Benthem, J. (1995). Temporal logic. In (Gabbay et al., 1995), pages 241–350.

Vere, S. A. (1983). Planning in time: Windows and durations for activities and goals. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 5:246–267.

Williams, B. (1986). Doing time: Putting qualitative reasoning on firmer ground. In *Proceedings of the 5th National Conference on Artificial Intelligence*, pages 105–112.