Reasoning About Change in a Changing World*

Madhura Nirkhe[1,2], Donald Perlis[1], and Sarit Kraus[2,3]

[1]Department of Computer Science
[2]Institute for Advanced Computer Studies
University of Maryland, College Park, MD 20742

[3]Department of Mathematics and Computer Science
Bar Ilan University Ramat Gan, 52900 Israel

e-mail:madhura@cs.umd.edu, perlis@cs.umd.edu, sarit@umiacs.umd.edu

Abstract

We outline a treatment of frame issues of significance to a time-situated reasoning mechanism in a dynamic setting with deadlines. We describe the inference mechanism designed for our step-logic planner and demonstrate how it successfully handles many variants of the Yale shooting problem, involving both projection and explanation scenarios.

1 A changing world

Uncertainty, urgency and unknowns characterize the complex world of a commonsense agent. Dudley, our agent, is in a particularly tight setting with respect to deadlines. He must formulate and enact a plan to save Nell who is tied to the railroad tracks[1]. As he reasons, the world around him continues to change. As the clock ticks a train rushes towards his beloved. Dudley must evidently keep track of the passage of time. Thus, we are no longer interested in a plan computed in a static world, but in a mechanism that is fully time-situated. Not only is the "now" changing, but Dudley's knowledge base is constantly undergoing change as inferences and observations are made. In a world in which only incomplete knowledge is available to Dudley, he must jump to default conclusions, often to realize in the course of time that they are wrong.

The frame problem in AI [?] is the problem of deciding what persists and what changes as actions are performed. Early STRIPS style planners in static settings explicitly modeled the effects of actions in the form of add and delete lists to get around the persistence problem. The world does not change much here, a single model is updated to reflect the results of actions. Two other related frame issues have been well documented in the literature [?]: The qualification problem is the impossibility of accounting for an arbitrarily large number of potential (pre)conditions for an action and the ramification problem is the difficulty in reasoning about the extended effects of actions as time passes. In this paper we show how Dudley's fully deadline-coupled reasoning mechanism offers a real-time solution to the above aspects of the frame problem.

We have formulated a combination of declarative and procedural approaches to model Dudley. The underlying mechanism is that of Elgot-Drapkin's step-logics which are described in section 2. In section 3 we sketch the temporal reasoning inference rules for frame-default reasoning. These form part of the real-time planning infrastructure. As a litmus test, we describe in section 4 how the same agent Dudley would behave in various dynamic versions of the infamous Yale shooting problem [?, ?, ?, ?, ?, ?, ?, ?, ?].

2 One step at a time

Step-logics [?] were introduced to model a commonsense agent's ongoing process of reasoning in a changing world. A *step* is defined as a fundamental unit of inference time. Beliefs are parameterized by the time taken for their inference, and these time parameters can themselves play a role in the specification of the inference rules and axioms. The most obvious way time parameters can enter is via the expression $Now(i)$, indicating the time is now i. Observations are inputs from the external world, and may arise at any step. When an observation appears, it is considered a belief in the same time-step. Each step of reasoning advances i by 1. At each new step i, the only information available to the agent upon which to base his further reasoning is a snap-shot of his deduction process completed up to and including step $i-1$. The agent's world knowledge is in the form of a database of beliefs. These contain domain specific axioms. A number of (domain independent) inference rules con-

*This work was supported in part by NSF grant No. IRI-9123460.

[1]This problem was first mentioned in the context of time-dependent reasoning by McDermott [?], and more recently discussed in [?].

stitute the inference engine. Among them are rules such as Modus Ponens and rules to incorporate new observations into the knowledge base, as well as rules for temporal reasoning. A projection rule extends the known state of the world in the future based in the context of each partial plan.

A partial plan is a temporally ordered list of triplets. Each triplet consists of an *action*, preceded and followed, respectively, by a list of *conditions* and *results*. The conditions may need to be true over all or some of the duration of the action. An action may be complex or primitive (atomic). A primitive action takes one time step to perform. Firing of an inference rule corresponds to a "think" action. Dudley's non-defeasible[2] beliefs are maintained as a set of **Facts**. Observations become facts in the same time step. Theorems whose premises consist of facts alone are also regarded as facts. Dudley simultaneously develops alternative plans towards attaining his goals or subgoals. Each of these partial plans (including the null plan which is a plan with no actions in it) defines a context within which reasoning can be done about the expected state of the world if the plan were to be carried to completion. The agent maintains a **Context_set**(CS) for each plan. The context set changes with time as the plan undergoes modification and as inferences are made in the context of the plan. The context set consists of the set of facts[3], along with the actions in the partial plan to start with. By the application of a restructured modus ponens rule this set is extended to include the results and extended effects of these actions as time progresses. We refer to [?] for a detailed description of the formalism. This paper deals with frame aspects of the same.

3 Frame-default inferencing

This section describes Dudley's inference mechanism for temporal reasoning. The temporal persistence rule (TP) tackles the projection frame problem. The Context Set Revision rule (CSR), together with the restructured Modus Ponens rule (RMP) deals with the ramification problem, by allowing the agent to reason about the extended effects of its actions in the context of its plans. The qualification problem is dealt with by a relevance mechanism that brings into short term focus only those conditions that may be relevant to the processing of an action at a current time [?]. For lack of space, we can not provide an example of a treatment of the ramification problem in planning, or of the qualification problem, we only describe a treatment of the projection frame problem here with reference to the YSP.

3.1 Temporal Projection Rule (TP):

At each step i, the context set of each plan as of step $i-1$ is used to obtain a Projection in the context of that plan, based on the default of persistence[4]. A Formula α in the context set corresponding to a predicate X is one of four forms. $X(S:T,\ldots)$ denotes that X holds over interval $S:T$, $\neg X(S:T,\ldots)$ denotes that $\neg X$ holds over $S:T$. Further, $X_c(S:T,\ldots)$ (resp., $\neg X_c(S:T,\ldots)$) are used to denote that not only does the agent believe in X (resp., $\neg X$) over the interval, but that the agent has reason to believe that the time point S could be a possible point of change of the predicate from $\neg X$ to X (resp., from X to $\neg X$) in the event that $\neg X$ (resp., X) holds in the interval ending in S. The formulae corresponding to each predicate X are kept sorted according to their time intervals. Either of $X(S:F,\ldots)$ and $X_c(S:F,\ldots)$ are defined to be in *direct contradiction* with $\neg X(S:F,\ldots)$ or with $\neg X_c(S:F,\ldots)$. A formula $X(S:F,U,\ldots)$ is in *uniqueness contradiction* with $X(S:F,V,\ldots)$ if $X(S:F,U,\ldots) \rightarrow \neg X(S:F,V,\ldots)$ whenever $U \neq V$[5]. A formulae α *du-contradicts* a formula δ, if it is in direct or uniqueness contradiction with δ. In essence, the temporal persistence rule (TP) *smoothes* over time intervals which present gaps in the agent's knowledge. At each step, a **Proj** in the context of a plan p holds the results of the temporal projection rule applied to the context set of the previous step, and is parametrized by the time as well as the name of the partial plan. Our approach can be best described by a term which we call *parallel projection.* That is, the entire known state of the world at one moment is used to determine the (expected) state at the next moment. Since step-logics are built around the idea of specifying what is known (e.g., proven) *so far*, all predicates, and all context sets can be simultaneously reconsidered at each new time step.

Here is a description of the TP rule applied to formulae corresponding to a given predicate X in the context set at step i in order to constitute $\mathbf{Proj}_{i+1}$ for a partial plan. Formulae are kept sorted in the $\mathbf{CS}_i$ according to their time intervals. The formulae in $\mathbf{Proj}_{i+1}$ are strictly those that are obtained by persistence of those in $\mathbf{CS}_i$. TP extends the formulae in $\mathbf{CS}_i$ to form the set $\mathbf{Proj}_{i+1}$. Let α_j and α_{j+1} denote consecutive formulae in this order. The TP rule can then be described as follows[6]:

1. If α_j is of the form $X(S_j:F_j,\ldots)$ and α_{j+1} is of

[2]Strictly speaking though, the agent only has beliefs, never facts, since even observations are not etched in stone, and may very well change over time.

[3]Actually it only consists of the subset of facts that is relevant to the particular partial plan. [?] deals with space bounds on the reasoning and proposes a relevance mechanism to keep the reasoning directed to a particular partial plan for a duration of time.

[4]Projections (and persistences) have been studied by numerous authors; see e.g., [?, ?, ?]. Our treatment is along the lines of time-maps of [?].

[5]For example, $at(5, Dudley, home)$ and $at(5, Dudley, railroad)$ are in uniqueness contradiction.

[6]For brevity, we only describe the rule for X. The dual form involving $\neg X$ is similar.

the form $X(S_{j+1} : F_{j+1}, \ldots)$ then **Proj** contains $X(F_j + 1 : S_{j+1} - 1, \ldots)$ whenever $F_j \leq S_{j+1}$.

2. If α_j is of the form $X(S_j : F_j, \ldots)$ and α_{j+1} is of the form $\neg X_c(S_{j+1} : F_{j+1}, \ldots)$ then **Proj** contains $X(F_j + 1 : S_{j+1} - 1, \ldots)$ whenever $F_j \leq S_{j+1}$.

3. If α_j is of the form $X(S_j : F_j, \ldots)$ and α_{j+1} is of the form $\neg X(S_{j+1} : F_{j+1}, \ldots)$ then **Proj** does not speculate over the truth or falsity of X over $F_j + 1 : S_{j+1} - 1$ since the agent has no basis for believing what could be a possible point of change in the value of X.

If any subset interval of α_j and α_{j+1} contradict according to the definition above, the projection is frozen until the contradiction is resolved[7].

3.2 A restructured Modus Ponens (resolution) rule (RMP)

Instead of applying MP in its familiar form : viz. from α and $\alpha \rightarrow \beta$ deduce β, we choose a representation in clause form and apply a restructured MP in accordance with our philosophy to let earlier defaults play out their effects completely. A formula which is a fact has no justification attached to it. A formula which was derived using a projection is only as feasible as the projection. Such a formula is annotated with the projections used in its derivation and is itself classified as a default (in contrast to a fact).

Let $\alpha_1 \vee \alpha_2 \vee \ldots \vee \alpha_n \vee \beta$ be the expression in clause form. Then the process of resolution can be outlined as follows:

- All the $\neg\alpha_j$ that belong to the set of **Facts** are first used to resolve. Subsequently, if the resolution is unfinished[8] members from the context set and the projection which are themselves defaults are next tried. Those from **Proj** with earlier time parameters are used before the ones with later parameters. The result of the resolution β is then annotated with all the projections used, whether directly, or in the annotation of resolving clauses from the context set. The annotations are attached in square brackets to the formulae. This provides a real-time truth maintenance mechanism which is useful in resolving contradictions. The result β may not be annotated with a projection with a later time than the time interval of β, and is discarded, if this is the case.
- If $\beta = X(S : F, \ldots)$ has its time interval $S : F$ such that S is later than the time intervals of all the α_j used in the resolution, then it is marked as $X_c(S : F, \ldots)$ in the context set, to denote that it could be a potential point of inflection in the value of the predicate X.

3.3 Context Set Extension and Revision Rule (CSR):

A restructured resolution rule is used to extend the context set. This allows Dudley to compute the extended effects of actions, thereby addressing the ramification problem. It allows to Dudley to continue to deduce the future consequences of his planning as it interacts possibly with the actions of other agents or with events observed in the world. It allows for reasoning with the current projection. It lets earlier events play out their consequences in an anticipated future before later events. Formulae are annotated by the projections which are used to support them in future conjectures. In the event that the projections cease to hold as of "now", the formulae that are supported by them are dropped from the context set in the revision process. The revision is a kind of real-time truth maintenance.

Following is a description of the rule to revise the Context set of a partial plan at each new time step $i+1$ based on the **CS** at step i and the **Proj** at step i.

1. If two formulae α and δ in $\mathbf{CS}_i$ du-contradict each other, then the following criteria are used in deciding which of them if any are retained in $\mathbf{CS}_{i+1}$[9].
   (a) If α is a fact, while δ is a default (is annotated with a projection), select α and block δ.
   (b) If α and δ are both defaults, select neither[10].

   The formulae selected as above as well as formulae that are not part of a contradiction are inherited subject to the following rules.
2. A formula α from $\mathbf{CS}_i$ which is a fact is inherited to $\mathbf{CS}_{i+1}$.
3. A formula $\alpha[\beta_1, \beta_2, \ldots, \beta_k]$ which is a default is inherited if for any $1 \leq j \leq k$, $\beta_j \notin \mathbf{Proj}_i$.
4. Any new formulae that result from the application of RMP to members of $\mathbf{CS}_i$ are added to $\mathbf{CS}_{i+1}$.

[7] As described in the next section, some of the α_j are facts while others are weaker since they are based on projections. The projection is frozen for subsequent α_{j+2}, α_{j+3} , ... unless a fact regarding X is encountered in the chain, at which point it is resumed. An example of this is in the projection of predicate *Alive* in Development 1 of the detective scenario, step 9.

[8] The resolution is said to yield the conclusion β through this restructured MP rule if it is the only unresolved formula residual.

[9] Note that we do not encounter situations in which facts (direct or indirect descendents of observations alone) contradict each other. Observations with different time intervals involving the same predicate may well disagree though.

[10] Where both are defeasible beliefs, a working strategy is to not inherit either of them, and to continue the reasoning to see if one of them will reappear in the face of stronger evidence.

5. Any new facts that arise from observations in step $i+1$, and are relevant to the partial plan are added to $\mathbf{CS}_{i+1}$.

4 Time-situated variations of the YSP

We have applied the above mechanisms to the Yale Shooting Problem [?] in several variations appropriate to step logics. The first is a *witness* scenario where Dudley is a witness to the scene of the crime. Here we show how Dudley draws the intuitive conclusion that Fred must be dead on observing the shoot action. The second is a *killer* scenario where Dudley formulates a plan to kill Fred by a certain deadline and reasons that Fred is expected to be dead in the context of his plan to carry out a shoot action. The third scenario is a *detective* scenario where Dudley must offer a reasonable explanation about actions in the past, to fit his present observations; on seeing Fred alive at a later time, the same mechanism allows him to continue to perform belief revision to account for "why things went wrong"[?]. In the interest of space, we provide an account of a witness scenario and two different detective scenarios involving explanations. The killer version is a direct extension of the planning mechanism.

In the *classical* YSP problem, there is a certain ambiguity about the role of the reasoner. There the reasoning is itself timeless, presumably it takes place after all the events in question.

Witness: We suppose that the reasoner is an eyewitness on the scene of the crime: Robbie sees Fred and sees a loaded gun at time 0, but no action is observed. Then, following a wait period during which nothing happens, at time 4 Robbie sees the gun being fired at Fred, but cannot see what happens to Fred after that. Robbie is then supposed to draw the commonsense conclusion that Fred has been killed. Here we are mimicking Baker's simplified version of the YSP (no loading action[?] is required, and the *wait* action occurs between steps 0 and 4).

Here we sketch the steps in the reasoning.

Axioms: $\neg Loaded(T) \vee \neg Shoot(T) \vee \neg Alive(T+1)$; $\neg Alive(T) \vee Alive(0:T)$

0: $\mathbf{Facts}(0, \{Alive(0)_{obs}, Loaded(0)_{obs}\})$,
$\mathbf{CS}(0, null, \{Alive(0)_{obs}, Loaded(0)_{obs}\})$,
$\mathbf{Proj}(0, null, \{\})$

1: $\mathbf{Facts}(1, \{Alive(0)_{obs}, Loaded(0)_{obs}\})$,
$\mathbf{CS}(1, null, \{Alive(0)_{obs}, Loaded(0)_{obs}\})$,
$\mathbf{Proj}(1, null, \{Alive(1:\infty), Loaded(1:\infty)\})$

2: $\mathbf{Facts}(2, \{Alive(0)_{obs}, Loaded(0)_{obs}\})$,
$\mathbf{CS}(2, null, \{Alive(0)_{obs}, Loaded(0)_{obs}\})$,
$\mathbf{Proj}(2, null, \{Alive(1:\infty), Loaded(1:\infty)\})$

3: $\mathbf{Facts}(3, \{Alive(0)_{obs}, Loaded(0)_{obs}\})$,
$\mathbf{CS}(3, null, \{Alive(0)_{obs}, Loaded(0)_{obs}\})$,
$\mathbf{Proj}(3, null, \{Alive(1:\infty), Loaded(1:\infty)\})$

4: $\mathbf{Facts}(4, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}\})$,
$\mathbf{CS}(4, null, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}\})$,
$\mathbf{Proj}(4, null, \{Alive(1:\infty), Loaded(1:\infty)\})$

5: $\mathbf{Facts}(5, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}\})$,
$\mathbf{CS}(5, null, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}, \neg Alive_c(5)[loaded(4)]\})$,
$\mathbf{Proj}(5, null, \{Alive(1:\infty), Loaded(1:\infty)\})$

6: $\mathbf{Facts}(6, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}\})$,
$\mathbf{CS}(6, null, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}, \neg Alive_c(5)[loaded(4)]\})$,
$\mathbf{Proj}(6, null, \{Alive(1:4), \neg Alive(6:\infty), Loaded(1:\infty)\})$

⋮

The witness version of the YSP gives the intuitive answer: In the context of the *null* plan, Fred must have died at step 5 as a result of the shooting, proviso of course, that the default regarding the gun staying loaded up until step 4 is indeed true. The projection reflects the presence of this conclusion in the context set. Note, that the set of facts does not include this conclusion, since $\neg Alive(5)[loaded(4)]$ is still defeasible, and hence can not be regarded as a fact.

To illustrate the real-time nature of our reasoning process, consider now two different developments of the witness scenario following step 6. In one of them, at step 7, Fred is seen to be alive and walking about, in another, at step 7, the gun is examined, and found to be in fact unloaded. Suppose for the moment that shooting does not unload the gun. What conclusions can Robbie make in each of these developments of the former scenario?

Development 1: But Fred is alive!

7: $\mathbf{Facts}(7, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}, Alive(7)_{obs}\})$,
$\mathbf{CS}(7, null, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}, \neg Alive_c(5)[loaded(4)], Alive(7)_{obs}\})$,
$\mathbf{Proj}(7, null, \{Alive(1:4), \neg Alive(6:\infty), Loaded(1:\infty)\})$

8: $\mathbf{Facts}(8, \{Alive(0)_{obs}, Alive(1:6), Loaded(0)_{obs}, Shoot(4)_{obs}, Alive(7)_{obs}\})$,
$\mathbf{Proj}(9, null, \{Alive(8:\infty), Loaded(1:\infty)\})$

10: $\mathbf{Facts}(10, \{Alive(0)_{obs}, Alive(1:6), Loaded(0)_{obs}, \neg Loaded(4), Shoot(4)_{obs}, Alive(7)_{obs}\})$,
$\mathbf{CS}(10, null, \{Alive(0)_{obs}, Alive(1:6), Loaded(0)_{obs}, \neg Loaded(4), Shoot(4)_{obs}, Alive(7)_{obs}\})$,
$\mathbf{Proj}(10, null, \{Alive(8:\infty), \neg Loaded(5:\infty)\})$

Development 2: Look at this gun, it is not loaded!

7: $\mathbf{Facts}(7, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}, \neg Loaded(7)_{obs}\})$,
$\mathbf{CS}(7, null, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}, \neg Alive_c(5)[loaded(4)], \neg Loaded(7)_{obs}\})$,
$\mathbf{Proj}(7, null, \{Alive(1:4), \neg Alive(6:\infty), Loaded(1:\infty)\})$

8: $\mathbf{Facts}(8, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}, \neg Loaded(7)_{obs}\})$,
$\mathbf{CS}(8, null, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}, \neg Alive_c(5)[loaded(4)], \neg Loaded(7)_{obs}\})$,
$\mathbf{Proj}(8, null, \{Alive(1:4), Alive(6:\infty), \neg Loaded(8:\infty)\})$

9: $\mathbf{Facts}(9, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}, \neg Loaded(7)_{obs}\})$,
$\mathbf{CS}(9, null, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}, \neg Loaded(7)_{obs}\})$,
$\mathbf{Proj}(9, null, \{Alived(1:4), \neg Alive(6:\infty), \neg Loaded(8:\infty)\})$

10: $\mathbf{Facts}(10, \{Alive(0)_{obs}, Alive(1:6), Loaded(0)_{obs}, Shoot(4)_{obs}, \neg Loaded(7)_{obs}\})$,
$\mathbf{CS}(10, null, \{Alive(0)_{obs}, Loaded(0)_{obs}, Shoot(4)_{obs}, \neg Loaded(7)_{obs}\})$,
$\mathbf{Proj}(10, null, \{Alive(1:\infty), \neg Loaded(8:\infty)\})$

Both the above are explanation scenarios. The elegant recovery of the agent is to be attributed to the non-monotonic inference process that initiates changes to restore consistency. In the first development, when Fred is found to be alive, and his being alive past the time of the shooting is regarded as a fact, Dudley can successfully conclude that the gun must have been unloaded between the loading and the time of the shooting. Furthermore, he can successfully change his on-going model to reflect these changes. In the second development, when the gun is found to be unloaded at a later time, under the assumption that shooting does not unload the gun, Dudley chooses to be ambivalent about Fred's death. Even though he had infered earlier that Fred would be dead by default, now he no longer has the same confidence in the projection, it must be changed, and the conclusion about Fred's death refrained from making, until more is known about the time of the unloading.

5 Conclusion

Our design and implementation of Dudley shows he has the flexibility to adapt his reasoning to changes in his environment in a non-monotonic manner, perhaps much in the way that human commonsense reasoners do. For example, added evidence resulting in apparent contradictions do not swamp Dudley with an infinite confusion of wffs, unlike traditional logics. He can sort his way through the initial contradiction with the help of his *time-evolving* non-monotonic inferences, as in [?]. We have shown that he can successfully handle temporal explanation problems with the same degree of ease as temporal projection. We have provided some suitable examples that cover both projection and explanation scenarios.

References

[1] Andrew B. Baker. A simple solution to the Yale Shooting Problem. In *Proceedings, KR89*, pages 11–20, 1989.

[2] P. Cohen and H. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.

[3] T. Dean and D. McDermott. Temporal data base management. *Artificial Intelligence*, 32(1):1–55, 1987.

[4] J. Elgot-Drapkin and D. Perlis. Reasoning situated in time I: Basic concepts. *Journal of Experimental and Theoretical Artificial Intelligence*, 2(1):75–98, 1990.

[5] Michael Gelfond. Autoepistemic logic and formalization of common-sense reasoning. In *Proceedings of the second international workshop on nonmonotonic reasoning*, 1988. Munich 88.

[6] Michael P. Georgeff. Many agents are better than one. In F. M. Brown, editor, *Proceedings of the 1987 Workshop on The Frame Problem*, pages 59–75x. Morgan Kaufmann, 1987. Lawrence, Kansas.

[7] M. L. Ginsberg and D. E. Smith. Reasoning about action I: A possible worlds approach. In F. M. Brown, editor, *Proceedings of the 1987 Workshop on The Frame Problem*, pages 233–258, Lawrence, KS, 1987. Morgan Kaufmann.

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

[9] H. Kautz. The logic of persistence. In *Proceedings, AAAI-86*, pages 401–405, 1986.

[10] Vladimir Lifschitz. Formal theories of action. In *The Frame Problem in Artificial Intelligence*, pages 35–57, Los Altos, CA, 1987. Morgan-Kaufmann.

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

[12] D. McDermott. Planning and acting. *Cognitive Science*, 2:71–109, 1978.

[13] D. McDermott. Nonmonotonic logic and temporal projection. *Artificial Intelligence*, 33:379–412, 1987.

[14] L. Morgenstern and L. A. Stein. Why things go wrong: A formal theory of causal reasoning. In *Proceeding, AAAI-88*, pages 518–523, 1988.

[15] P. H. Morris. The anomalous extension problem in default reasoning. *Artificial Intelligence*, 35:383–399, 1988.

[16] M. Nirkhe, S. Kraus, and D. Perlis. Fully deadline-coupled planning: One step at a time. In *Proceedings of the sixth international symposium on methodologies for intelligent systems*, Charlotte, NC, 1991.

[17] M. Nirkhe, S. Kraus, and D. Perlis. Time-situated reasoning within tight deadlines and realistic space and computation bounds. In *Proceedings of the Second Symposium On Logical Formalizations Of Commonsense Reasoning*, Austin, TX, 1993.

[18] Judea Pearl. On logic and probability. *Computational Intelligence*, 4:99–103, 1988.

[19] Yoav Shoham. Chronological ignorance: Experiments in nonmonotonic temporal reasoning. *Artificial Intelligence*, 36:279–331, 1988.

[20] R. Wilensky. *Planning and understanding.* Addison Wesley, Reading, Mass, 1983.