A General Framework for Expressing Preferences in Causal Reasoning and Planning

James P. Delgrande
School of Computing Science
Simon Fraser University
Burnaby, B.C.
Canada V5A 1S6
jim@cs.sfu.ca

Torsten Schaub*
Institut für Informatik
Universität Potsdam
Postfach 90 03 27
D–14439 Potsdam, Germany
torsten@cs.uni-potsdam.de

Hans Tompits
Institut für Informationssysteme
Technische Universität Wien
Favoritenstraße 9–11
A–1040 Vienna, Austria
tompits@kr.tuwien.ac.at

Abstract

We consider the problem of representing arbitrary preferences in causal reasoning and planning systems. In planning, a preference may be seen as a goal or constraint that is desirable, but not necessary, to satisfy. To begin, we define a very general query language for histories, or interleaved sequences of world states and actions. Based on this, we specify a second language in which preferences are defined. A single preference defines a binary relation on histories, indicating that one history is preferred to the other. From this, one can define global preference orderings on the set of histories, the maximal elements of which are the preferred histories. The approach is very general and flexible; thus it constitutes a "base" language in terms of which higher-level preferences may be defined. To this end, we investigate two fundamental types of preferences that we call *choice* and *temporal* preferences. We consider concrete strategies for these types of preferences and encode them in terms of our framework. We suggest how to express aggregates in the approach, allowing, for example, the expression of a preference for histories with lowest total action costs. Last, our approach can be used to express other approaches, and so serves as a common framework in which such approaches can be expressed and compared. We illustrate this by indicating how an approach due to Son and Pontelli can be encoded in our approach, as well as the language PDDL3.

*Affiliated with the School of Computing Science at Simon Fraser University, Burnaby, Canada, and the Institute for Integrated and Intelligent Systems at Griffith University, Brisbane, Australia.

1 Introduction

Planning, as traditionally formulated, involves attaining a particular goal, given an initial state of the world and a description of actions. A plan succeeds just when it is executable and attains the goal; otherwise it fails. However, in realistic situations, things are not quite so simple. Thus, there may be requirements specifying that a plan should be as short as possible or that total cost, where costs are associated with actions, be minimised. Besides such indicators of *plan quality*, there may be *domain specific* conditions that are desirable yet not necessary to attain. Consider for example an extension of the well-known monkey-and-bananas problem. In addition to the usual information where the monkey can push a box and climb on the box to grasp some bananas, assume that the monkey has a range of choices for food. Perhaps the monkey prefers to have an appetizer before the main course, although this preference need not be satisfied in attaining the overall goal of eating a full meal. Perhaps too the monkey prefers soup to grubs as an appetizer, although either will serve as an appetizer. If the monkey has soup, then it prefers to have a spoon before it has the soup; and if it does not have a spoon before having soup then it prefers to have a spoon as soon as possible after getting the soup. Clearly, such preferences can be arbitrarily complex, and range over temporal constraints as well as relations among fluents and actions. In this setting, the goal of a planning problem now shifts to determining a *preferred* plan, in which a maximal set of preferences is satisfied along with the goal. Such preferences also make sense outside of planning domains, and in fact apply to arbitrary sequences of temporal events. Hence it is perfectly rational to prefer that it rains during the next several work days (since the plants need the water) but that it be sunny for the weekend. While one has no influence over the weather, it is perfectly sensible to prefer one outcome over another.

In this paper, we consider the problem of using general preferences over (fluent and action) formulas to determine preferences among temporal histories, or plans. While we focus on histories as they are used in action description languages [19], our approach is readily applicable to any planning formalism. We begin by specifying a general query language, $\mathcal{Q}_{\Sigma,n}$, for histories of a given maximum length n built from a signature Σ, in which one can determine whether an arbitrary expression is true in a given history. Given this language, we subsequently define a *preference-specification language*, $\mathcal{P}_{\Sigma,n}$, over $\mathcal{Q}_{\Sigma,n}$ that enables the definition of preference relations between histories. Specifically, a preference between two histories H_h and H_l is expressed in terms of a formula ϕ of $\mathcal{P}_{\Sigma,n}$ such that H_h is not less preferred than H_l under ϕ, written $H_l \preceq_\phi H_h$, just if $\langle H_l, H_h\rangle \models_{\mathcal{P}_{\Sigma,n}} \phi$, where the satisfaction relation of $\mathcal{P}_{\Sigma,n}$, $\models_{\mathcal{P}_{\Sigma,n}}$, is defined between pairs of histories and formulas of $\mathcal{P}_{\Sigma,n}$. This gives us a means of expressing arbitrary preferences between histories. Given such (binary) preferences, one can determine the most preferred history by appealing to extant work on preferences in general (like, e.g., the approach proposed by Brewka [6]).

The resulting languages are highly expressive. Indeed, the problem of satisfiability of a query in either language is PSPACE-complete, while the corresponding model-checking problem is polynomial. Arguably, the languages are adequate for expressing qualitative preferences in temporal, causal, and planning frameworks. As well, the approach provides a very general language in which other "higher-level" constructs can be encoded, and in which other approaches can be expressed, and so compared. To illustrate this capability, we identify two basic, but widely applicable types of preferences, called *choice* and *temporal* preferences. In the former case some condition (fluent, action) is preferred over another (for example, having soup to grubs as an appetizer); in the latter case some condition is preferred to take place before another (for example, having an appetizer before a main meal). Given the fundamental importance of these types of preferences, we provide an extended development of how preferred histories can be determined with respect to a set of choice preferences or a set of temporal preferences. We also briefly discuss how aggregates can be defined in our query language. This allows, for example, the specification of the total value of a fluent, such as one corresponding to action costs. From this, using our preference language, we can express, for example, that a history with lowest total action costs is preferred. Last, we claim that one can encode other approaches in our framework. We illustrate this with an encoding of the planning approach by Son and Pontelli [33], as well as a portion of the planning domain definition language PDDL3. Hence, our approach provides a means for both expressing arbitrary preferences and for comparing different approaches.

The next section informally describes basic kinds of preferences in planning, while Section 3 covers related work. Subsequently, Section 4 introduces our preference framework, in which preferences induce a binary (relative) preference relation on histories, followed by a discussion of determining a maximally-preferred history. Section 5 deals with modelling choice and temporal preferences. Section 6 continues with extending our language with aggregate features, and Section 7 discusses how other approaches can be captured within our framework. Section 8 concludes with a brief discussion. This paper subsumes and extends preliminary work reported previously in [7, 8].

2 Preferences in Planning

As mentioned, preferences are a general phenomenon, and are pervasive in realistic domains. Very generally and informally, a preference may in some sense be *absolute* (for example, that one likes sushi) or *relative* (for example, that one prefers sushi to tempura). In either case, a preference may be regarded as providing a criterion for distinguishing states of affairs. In the first case, states of affairs (or possible worlds) in which one has sushi are preferred to those where this does not obtain; in the second case, states of affairs in which one has sushi are preferred to those where one has tempura (with other worlds being incomparable). Given a set of preferences, the goal usually is to determine the most preferred world (or, the state of affairs in which the greatest number of preferences are satisfied).[1] In so doing, one typically relies on determining transitive closure, both with respect to the original set of preferences and with respect to the ordering on possible worlds. However, transitive closure is not expressible in first-order logic. Hence, determining which preferences are applicable in a given situation is typically addressed via meta-level means.

Before reviewing relevant background material and introducing our approach, we briefly consider characteristics of preferences with respect to a set of histories, where a history is a sequence of alternating possible worlds and actions. To this end, we suggest two dimensions whereby preferences may be classified. The first of these, whether preferences are on actions or fluents, we argue is in fact not a meaningful distinction. The second, choice vs. temporal preferences, though, we argue does indeed isolate two useful preference types for reasoning in temporal domains.

[1] To be sure, this is a simplistic characterisation. However, it serves to set the stage for our approach. For a full introduction to preferences, the reader may consult one of the cited works in Section 3.

Fluent vs. Action Preference. Informally, a preference can be expressed between fluents or between actions. In the first case, one might prefer to have white wine to red; in the latter, one might prefer to travel via transit to driving. This distinction is by no means clear-cut however, and a natural-language assertion representing a preference can often be interpreted as being either between actions or between fluents. For example, preferring to have white wine to red seems to be a preference between fluents, whereas preferring to drink white wine over red seems to be a preference on actions. As well, preferences on actions can easily be specified as preferences on fluents: for each action, simply introduce a new fluent, with the same arguments as the action, indicating that the action has been carried out. Thus, for action $drop(x, y, t)$ indicating that agent x drops object y leading to time point t, one can introduce a new fluent $was\_dropped(x, y)$, that becomes true if and only if the $drop$ action was executed on the same arguments. Hence, one can restrict preferences to being on fluents only with no loss of generality.

Choice vs. Temporal Preference. For *choice preference*, one has preferences concerning how a subgoal is to be attained. For example, a $dinner$ subgoal may be preferentially satisfied by having Japanese food over Chinese food. For *temporal preference*, one has preferences concerning the order in which subgoals are to be achieved. Thus, the subgoal of having had $dinner$ may preferentially be satisfied before that of $movie$. Similarly, choice and temporal preferences involving actions are easily constructed.

As in the preceding distinction between fluent and action preferences, temporal preferences can also be reduced to choice preferences by introducing suitable fluents associated with specific time points. For instance, the preference of having dinner before going to the movies may be represented by using new fluents $dinner\_i$ and $movie\_i$, stating that dinner and movie happens at time point i, respectively, and asserting that $dinner\_i$ is preferred over $movie\_j$ for $i < j$. However, such a rewriting schema yields in general a significant increase in the states of world, which is not desirable for planning. Hence, although strictly speaking temporal preferences can be reduced to choice preferences, nonetheless keeping the distinction between choice and temporal preferences is useful from a representational point of view. Thus, in Section 5, we classify basic preferences as to whether they are a choice preference or a temporal preference between formulas. These are specified via two partial preorders, $\leq_c$ and $\leq_t$, respectively.

There are two other factors concerning preferences that we briefly mention. First, it might seem that there is a distinction between whether a preference is *relative* to others or *absolute*. Thus, in the first case, one might prefer red wine to white, while in the second, one simply has a preference for red wine. One can employ choice preferences to express this latter type of preference, though: that α is (simply) desirable can be expressed by $\neg\alpha \leq_c \alpha$. Consequently, we do not consider absolute preferences to be an independent classification.

Second, we can easily express conditional preferences, or preferences applicable in a given context (for example, in travelling, if the distance is great, transit is preferred over driving). The choice preference that α is preferred to β in the context γ is given by $\gamma \wedge \beta <_c \gamma \wedge \alpha$.

3 Background

Reasoning with preferences is an active area that is receiving increasing attention in AI. The literature is extensive; we mention only Oztürk *et al*. [30] as an introduction to preference modelling in general, and a recent special issue of *Computational Intelligence* [25] for a selection of papers addressing preferences in AI and constraint solving.

There has been some work in expressing procedural control knowledge in planning. For example, Bacchus and Kabanza [2] show how the performance of a forward-chaining planner can be improved by incorporating domain-specific control knowledge; cf. also Doherty and Kvarnstrom [10] and Nau *et al*. [29]. Son *et al.* [32] address domain and procedural control knowledge, as well as ordering constraints, in an action language expressed via an extended logic program. Since these are "hard" constraints, requiring (rather than suggesting) that, e.g., some action precedes another, the goal of such work differs from ours.

Our interests lie with preferences in planning and, more generally, temporal formalisms. Wellman and Doyle [35] suggest that the notion of *goal* is a relatively crude measure for planners to achieve, and instead that a relative preference over possible plan outcomes constitutes (or should constitute) a fundamental objective for planning. They show how to define goals in terms of preferences and, conversely, how to define (incompletely) preferences in terms of sets of goals.

Myers and Lee [28] assume that there is a set of desiderata, such as affordability or time, whereby successful plans can be ranked. A small number of plans is generated, where the intent is to generate divergent plans. The best plan is then chosen, based on a notion of Euclidean distance between these select attributes. In related work, Haddawy and Hanks [23] use a utility function to guide a planner.

Eiter *et al.* [11] describe planning in an answer-set programming framework where action costs are taken into account. The approach allows the specification of desiderata such as computing the shortest plan, or the cheapest plan, or some combination of these criteria. This is achieved by employing weak constraints, which filter answer sets, and thus of plans, based on *quantitative* criteria.

One approach to preferences in planning has been developed by Son and Pontelli [33], where a language for specifying preferences between histories is presented. This language is an extension of action language $\mathcal{B}$ [19] and is subsequently compiled into logic programs under the answer-set semantics [18]. The notion of preference explored is based on so-called *desires* (what we call *absolute preferences*), expressed via formulas built by means of propositional as well as temporal connectives such as *always*, *until*, etc. From desires, preferences among histories are induced as follows: Given a desire ϕ, a history H is preferred to H' if $H \models \phi$ but $H' \not\models \phi$. We discuss this language in more detail in Section 7.1.

Similarly, Bienvenu, Fritz, and McIlraith [3] address planning with preferences in the situation calculus. Preferences are founded on the notion of *basic desire formulas*, whose members are somewhat analogous to formulas in our language $\mathcal{Q}_{\Sigma,n}$. These formulas in turn are used in the composition of *atomic preference formulas* (essentially chains of preferences) and *general preference formulas*. As the authors note, this approach extends and modifies that of Son and Pontelli [33] although expressed in terms of the situation calculus rather than an action language. Based on a concrete means of explicitly combining preferences, a best-first planner, PPLAN, is given. As well, Fritz and McIlraith [15] employ a subset of this language in an approach to compile preferences into DT-Golog.

Brafman and Chernyavsky [5] propose a constraint-based approach to planning with goal preferences. This approach is restricted to preferences on the final state of a history. These preferences are expressed by means of so-called TCP-nets, which makes them easily integratable into CSP-based planning approaches. Feldmann, Brewka, and Wenzel [12] also address planning with prioritized goals, using *ranked knowledge bases* to express qualitative preferences among goals. The accompanying algorithm attempts to determine a sequence of improved plans which converges toward an optimal plan, by repeated calls to an (arbitrary) underlying classical planner.

We note that strong and soft constraints on plan trajectories and problem goals have recently been included into the *Planning Domain Definition Language* (PDDL) [20]. We discuss this language in more detail in Section 7.2.

Last, the goal of our research is an approach for addressing preferences in temporal and planning settings. A prerequisite to the success of this endeavour is the development of languages for combining preference relations; for work in AI to this end, see, for example, Boutilier *et al.* [4] and Brewka [6].

4 Expressing Preferences on Histories

4.1 The Approach

Our central notion is that of a *preference frame*, consisting of a pair $\langle \boldsymbol{H}, \boldsymbol{P} \rangle$, where

- $\boldsymbol{H}$ is a set of *histories* and
- $\boldsymbol{P}$ is a set of *preferences* on histories.

A history is a sequence of states and named transitions between states, for example representing a plan that accomplishes a given goal, or, more generally, some evolution of the world. A preference specifies a criterion for distinguishing between histories. We use the syntax that a history $H \in \boldsymbol{H}$ is a sequence $(s_0, a_1, s_1, a_2, s_2, \ldots, s_{n-1}, a_n, s_n)$, where s_0 is the initial state of H, and the subsequence s_i, a_{i+1}, s_{i+1} indicates that action a_{i+1} takes the world from state s_i to s_{i+1}. This notation is for convenience only; we could as well have based our approach on, for example, *situations* [26] or any other notation that carries the same information. $\boldsymbol{H}$ can be equated with a complete description of a planning problem, with members of $\boldsymbol{H}$ corresponding to complete plans.

We define a preference between two histories directly in terms of a formula ϕ and a satisfaction relation $\models$. That is, we define that H_h is not less preferred than H_l under ϕ, written $H_l \preceq_\phi H_h$, just if $\langle H_l, H_h \rangle \models \phi$. In other words, the formula ϕ expresses a preference condition between two histories, and $H_l \preceq_\phi H_h$ holds if ϕ is true by evaluating it with respect to H_l and H_h. The formulas used for expressing preferences are composed of

- Boolean combinations of fluents and actions indexed by time points in a history, and with respect to a history, and
- quantifications over time points.

Indexing with respect to time points and histories is achieved via *labelled atoms* of form $\ell : b(i)$. Here, ℓ is a *label*, either $\boldsymbol{l}$ or $\boldsymbol{h}$, referring to a history which is considered to be lower or higher ranked, respectively, b is an action or fluent name, and i is a time point. The relation $\langle H_l, H_h \rangle \models \boldsymbol{l} : b(i)$ is true if b holds at time point i in history H_l; and analogously for $\langle H_l, H_h \rangle \models \boldsymbol{h} : b(i)$. This is extended to labelled formulas in the expected fashion.[2] For example, we can express that history H_h is preferred to history H_l if fluent f is true at some point in H_h but never true in H_l by the formula

$$\phi = (\boldsymbol{h} : \exists i\, f(i)) \ \wedge \ (\boldsymbol{l} : \forall i\, \neg f(i)) \;, \tag{1}$$

providing $\langle H_l, H_h \rangle \models \phi$ holds.

The binary relation $\preceq_\phi$ induced by a formula $\phi \in \boldsymbol{P}$ has no particular properties (such as transitivity) of course, since any such property will depend on the formula ϕ. Depending on the type of preference encoded in $\boldsymbol{P}$, one would supply a strategy from which a maximally preferred history is selected. Thus, for preferences only of the form (1), indicating which fluents are desirable, the maximally preferred histories might be the ones which were ranked as "preferred" by the greatest number of preferences in $\boldsymbol{P}$.

4.2 Histories and Queries on Histories

In specifying histories, we begin with notation adapted from Gelfond and Lifschitz [19] in their discussion of *transition systems*.

Definition 1 *An* action signature*, Σ, is a triple $\langle V, F, A \rangle$, where V is a set of* value names*, F is a set of* fluent names*, and A is a set of* action *names.*

If $V = \{1, 0\}$, then Σ is called propositional*. If V, F, and A are finite, then Σ is called* finite.

Definition 2 *Let $\Sigma = \langle V, F, A \rangle$ be an action signature.*

A history*, H, over Σ is a sequence*

$$(s_0, a_1, s_1, a_2, s_2, \ldots, s_{n-1}, a_n, s_n),$$

where

- $n \geq 0$,
- *each s_i, $0 \leq i \leq n$, is a mapping assigning each fluent $f \in F$ a value $v \in V$, and*
- $a_1, \ldots, a_n \in A$.

The functions $s_0, \ldots, s_n$ are called states*, and n is the* length *of history H, symbolically $|H|$.*

The states of a history may be thought of as possible worlds, and the actions take one possible world into another. For a propositional action signature $\Sigma = \langle V, F, A \rangle$, fluent $f \in F$ is said to be *true at* state s iff $s(f) = 1$, otherwise f is *false at* s.

[2] We remark that labels, as employed here, serve a similar purpose as labels used in tableau systems [14] or in labelled deduction [16].

In order to express complex conditions that refer to properties of fluents and actions within a history, we first define a query language on histories of maximum length n over an action signature Σ, named $\mathcal{Q}_{\Sigma,n}$. In the next subsection, we extend this language to deal with formulas that refer to more than one history.

Definition 3 *Let $\Sigma = \langle V, F, A \rangle$ be an action signature and $n \geq 0$ a natural number.*
We define the query language $\mathcal{Q}_{\Sigma,n}$ as follows:

1. *The alphabet of $\mathcal{Q}_{\Sigma,n}$ consists of*
   - *(a) a set $\mathcal{V}$ of* time-stamp variables*, or simply* variables*,*
   - *(b) the set $\{0, \ldots, n\}$ of natural numbers,*
   - *(c) the arithmetic function symbols '$+$' and '$\cdot$',*
   - *(d) the arithmetic relation symbol '$<$',*
   - *(e) the equality symbol '$=$',*
   - *(f) the set $A \cup F$ of action and fluent names,*
   - *(g) the sentential connectives '$\neg$' and '$\supset$',*
   - *(h) the quantifier symbol '$\exists$', and*
   - *(i) the parentheses '$($' and '$)$'.*
2. *A* time term *is an arithmetic term recursively built from variables and numbers in $\mathcal{V} \cup \{0, \ldots, n\}$, employing $+$ and $\cdot$ (as well as parentheses) in the usual manner.*

   A time atom *is an arithmetic expression of form $(t_1 < t_2)$ or $(t_1 = t_2)$, where t_1, t_2 are time terms.*

   We use TT_n to refer to the set of time terms.
3. *A* fluent atom *is an expression of form $(f(t) = v)$, where $f \in F$, $t \in TT_n$, and $v \in V$.*

   An action atom *is an expression of form $a(t)$, where $a \in A$ and $t \in TT_n$.*

   The set of atoms *is made up of the set of time atoms, fluent atoms, and action atoms. An atom containing no variables is* ground.
4. *A* literal *is an atom, or an atom preceded by the sign $\neg$.*
5. *A* formula *is a Boolean combination of atoms, along with quantifier expressions of form $\exists i$, for $i \in \mathcal{V}$, formed in the usual recursive fashion.*

   A formula containing no free (time-stamp) variables is closed.
6. *A* query *is a closed formula.*

Besides the primitive symbols of $\mathcal{Q}_{\Sigma,n}$, we define the operators $\wedge$, $\vee$, and $\leq$, and the universal quantifier $\forall$, in the usual way. We sometimes drop parentheses in formulas if no ambiguity arises, and we may write quantified formulas like $\mathsf{Q}i_1\mathsf{Q}i_2\,\alpha$ as $\mathsf{Q}i_1, i_2\,\alpha$, for $\mathsf{Q} \in \{\forall, \exists\}$. As usual, a literal of form $\neg(f(t) = v)$ may also be written as $(f(t) \neq v)$. For formula α, variables $i_1, \ldots, i_k$, and numbers $m_1, \ldots, m_k$, by $\alpha[i_1/m_1, \ldots, i_k/m_k]$ we denote the result of uniformly substituting i_j by m_j in α, for each $j \in \{1, \ldots, k\}$. Thus, if $i_1, \ldots, i_k$ are the free variables in α, then $\alpha[i_1/m_1, \ldots, i_k/m_k]$ is a closed formula. For ground time term t, $val(t)$ is the value of t according to standard integer arithmetic.

Variables range over time points, and so quantification applies to time points only. Atoms consist of action or fluent names indexed by a time point, or of a predicate on arithmetic (time point) expressions. Atoms are used to compose formulas in the standard fashion, and queries consist of closed formulas. This means that we remain within the realm of propositional logic, since quantified expressions $\forall i$ and $\exists i$ can be replaced by the conjunction or disjunction (respectively) of their instances.

As an example, let $red \in V$, $colour \in F$, $pickup \in A$, and $i, j \in \mathcal{V}$. Then,

$$pickup(4), \quad colour(i+j) = red, \quad i < j + 2$$

are atoms. As well,

$$(colour(j) = red) \land \big(\forall k\ (k < j) \supset \neg(colour(k) = red)\big)$$

is a formula, and

$$\exists i, j \big((i + 2 < j) \land pickup(i) \land \neg(colour(j) = red)\big)$$

is a closed formula and so a query. This last formula is true in a history in which $pickup$ is true at some time point and three or more time points later $colour$ does not have value red.

The definition of truth of a query is as follows.

Definition 4 *Let $H = (s_0, a_1, s_1, \ldots, a_k, s_k)$ be a history over Σ of length $k \leq n$, and let Q be a query of $\mathcal{Q}_{\Sigma,n}$.*

We define $H \models_{\mathcal{Q}_{\Sigma,n}} Q$ recursively as follows:

1. *If Q is a ground time atom, then $H \models_{\mathcal{Q}_{\Sigma,n}} Q$ iff Q is true according to the rules of integer arithmetic.*

2. *If $Q = (f(t) = v)$ is a ground fluent atom, then $H \models_{\mathcal{Q}_{\Sigma,n}} Q$ iff $s_m(f) = v$, where $m = \min(val(t), n)$.*

3. *If $Q = a(t)$ is a ground action atom, then $H \models_{\mathcal{Q}_{\Sigma,n}} Q$ iff $a = a_m$, where $m = \min(val(t), n)$.*

4. *If $Q = \neg\alpha$, then $H \models_{\mathcal{Q}_{\Sigma,n}} Q$ iff $H \not\models_{\mathcal{Q}_{\Sigma,n}} \alpha$.*

5. *If $Q = (\alpha \supset \beta)$, then $H \models_{\mathcal{Q}_{\Sigma,n}} Q$ iff $H \not\models_{\mathcal{Q}_{\Sigma,n}} \alpha$ or $H \models_{\mathcal{Q}_{\Sigma,n}} \beta$.*

6. *If $Q = \exists i\, \alpha$, then $H \models_{\mathcal{Q}_{\Sigma,n}} Q$ iff, for some $0 \leq m \leq n$, $H \models_{\mathcal{Q}_{\Sigma,n}} \alpha[i/m]$.*

If $H \models_{\mathcal{Q}_{\Sigma,n}} Q$ holds, then H *satisfies* Q. For simplicity, if $\mathcal{Q}_{\Sigma,n}$ is unambiguously fixed, we also write $\models$ instead of $\models_{\mathcal{Q}_{\Sigma,n}}$.

The rationale for taking $m = \min(val(t), n)$ in Items 2 and 3 of Definition 4 is to take into account that a time term may refer to a time point which lies beyond the interval determined by the length n of a given history. Intuitively, if $val(t)$ is greater than the length n of history H, then a ground atomic query $b(t)$, where b is either an action name or a fluent name, is satisfied by H if it is satisfied at the last state of H.

For propositional action signatures, it is convenient to define the following abbreviations. To begin with, fluent atoms of form $f(t) = 1$ may be written simply as $f(t)$. Hence, since for any history H we have that $H \models (f(t) = 0)$ iff $H \models \neg(f(t) = 1)$, we may write an atom of form $f(t) = 0$ as a literal $\neg f(t)$. Furthermore, the following operators, which basically correspond to similar well-known operators from linear temporal logic (LTL), can be defined:

- $\Box b = \forall i\, b(i)$;
- $\Diamond b = \exists i\, b(i)$; and
- $(b \,\mathsf{U}\, g) = \exists i \big(g(i) \wedge \forall j((j < i) \supset b(j))\big)$.

Here, b and g are fluent or action names. Informally, $\Box b$ expresses that b *always* holds, $\Diamond b$ that b holds *eventually*, and $b \,\mathsf{U}\, g$ that b holds continually *until* g holds. Other temporal operators are likewise expressible.

We have the following results concerning complexity in $\mathcal{Q}_{\Sigma,n}$.

Theorem 1 *Let Σ be an action signature and $n \geq 1$ a natural number.*

1. *Deciding whether $H \models_{\mathcal{Q}_{\Sigma,n}} Q$ holds, given a history $H = (s_0, a_1, s_1, \ldots, a_n, s_n)$ over Σ of length n and a query $Q = (\mathsf{Q}_1 i_1)(\mathsf{Q}_2 i_2)\ldots(\mathsf{Q}_m i_m)\, C$ of $\mathcal{Q}_{\Sigma,n}$, where $\mathsf{Q}_j \in \{\forall, \exists\}$, $1 \leq j \leq m$, and C is a formula containing no quantifiers, can be determined in $O(|C|^m)$ time.*
2. *Deciding whether there is a history H over Σ of length n such that $H \models_{\mathcal{Q}_{\Sigma,n}} Q$ holds, given a query Q of $\mathcal{Q}_{\Sigma,n}$, is PSPACE-complete.*

The proof of the first part is straightforward, since for each quantifier expression $(\mathsf{Q}i)\alpha$, one needs to test the n substitution instances of α conjunctively (for universal quantification) or disjunctively (for existential quantification). For the second part, for showing that the problem is at least in PSPACE, the reduction is from satisfiability of quantified Boolean formulas to formulas of $\mathcal{Q}_{\Sigma,1}$; for showing that it is no worse than PSPACE, the reduction is from formulas of $\mathcal{Q}_{\Sigma,n}$ to formulas of linear time temporal logic (LTL).

4.3 Expressing Preferences between Histories

We define a preference between two histories, H_l and H_h, directly in terms of a formula ϕ and a satisfaction relation $\models$:

$$H_l \preceq_\phi H_h \quad \text{iff} \quad \langle H_l, H_h \rangle \models \phi \,. \tag{2}$$

The intent is that ϕ expresses a condition in which H_h is at least as preferred as H_l. This requires that we be able to talk about the truth values of fluents and actions in H_l and H_h. In the previous subsection, we defined a query language on histories, $\mathcal{Q}_{\Sigma,n}$, and a notion of truth in a history for a query. Given these definitions, we are now in a position to introduce a definition of preference between histories, as in (2). To do this, we introduce a *preference-specification language* $\mathcal{P}_{\Sigma,n}$ based on $\mathcal{Q}_{\Sigma,n}$.

Definition 5 *Let $\Sigma = \langle V, F, A \rangle$ be an action signature and $n \geq 0$ a natural number.*

We define the preference-specification language $\mathcal{P}_{\Sigma,n}$ *over* $\mathcal{Q}_{\Sigma,n}$ *as follows:*

1. *The alphabet of $\mathcal{P}_{\Sigma,n}$ consists of the alphabet of the query language $\mathcal{Q}_{\Sigma,n}$, together with the symbols $\boldsymbol{l}$ and $\boldsymbol{h}$, called* history labels*, or simply* labels.

2. *Atoms of $\mathcal{P}_{\Sigma,n}$ are either time atoms of $\mathcal{Q}_{\Sigma,n}$ or expressions of form $\ell : p$, where $\ell \in \{\boldsymbol{l}, \boldsymbol{h}\}$ is a label and p is an action or fluent atom of $\mathcal{Q}_{\Sigma,n}$.*

   Atoms of form $\ell : p$ are called labelled atoms*, with ℓ being the label* of *$\ell : p$. A labelled atom $\ell : p$ is* ground *iff p is ground.*

3. *Formulas of $\mathcal{P}_{\Sigma,n}$ are built from atoms in a similar fashion to formulas of $\mathcal{Q}_{\Sigma,n}$ and are called* preference formulas.

4. *A preference formula containing no free (time-stamp) variables is* closed.

   A preference axiom*, or simply* axiom*, is a closed preference formula.*

For a formula α of $\mathcal{Q}_{\Sigma,n}$ and a history label $\ell \in \{\boldsymbol{l}, \boldsymbol{h}\}$, by $\ell : \alpha$ we understand the formula resulting from α by replacing each action and fluent atom p of α by the labelled atom $\ell : p$. Informally, a labelled atom $\ell : p$ expresses that p holds in a history associated with label ℓ. The idea is that histories associated with label $\boldsymbol{h}$ are at least as preferred as histories associated with label $\boldsymbol{l}$. This is made precise as follows.

Definition 6 *Let Σ be an action signature and $n \geq 0$. Furthermore, let ϕ be a preference axiom of $\mathcal{P}_{\Sigma,n}$ and let H_l, H_h histories over Σ with $|H_i| \leq n$, for $i = l, h$.*

We define $\langle H_l, H_h \rangle \models_{\mathcal{P}_{\Sigma,n}} \phi$ recursively as follows:

1. *If ϕ is a time atom, then $\langle H_l, H_h \rangle \models_{\mathcal{P}_{\Sigma,n}} \phi$ iff ϕ is true according to the rules of integer arithmetic.*

2. *If $\phi = \ell : p$ is a ground labelled atom, for $\ell \in \{\boldsymbol{l}, \boldsymbol{h}\}$, then $\langle H_l, H_h \rangle \models_{\mathcal{P}_{\Sigma,n}} \phi$ iff*

   (a) $H_l \models_{\mathcal{Q}_{\Sigma,n}} p$, for $\ell = \boldsymbol{l}$, and

   (b) $H_h \models_{\mathcal{Q}_{\Sigma,n}} p$, for $\ell = \boldsymbol{h}$.

3. *If $\phi = \neg\psi$, then $\langle H_l, H_h \rangle \models_{\mathcal{P}_{\Sigma,n}} \phi$ iff $\langle H_l, H_h \rangle \not\models_{\mathcal{P}_{\Sigma,n}} \psi$.*

4. *If $\phi = \psi \supset \delta$, then $\langle H_l, H_h \rangle \models_{\mathcal{P}_{\Sigma,n}} \phi$ iff $\langle H_l, H_h \rangle \not\models_{\mathcal{P}_{\Sigma,n}} \psi$ or $\langle H_l, H_h \rangle \models_{\mathcal{P}_{\Sigma,n}} \delta$.*

5. *If $\phi = \exists i\, \psi$, then $\langle H_l, H_h \rangle \models_{\mathcal{P}_{\Sigma,n}} \phi$ iff, for some $0 \leq m \leq n$, $\langle H_l, H_h \rangle \models_{\mathcal{P}_{\Sigma,n}} \psi[i/m]$.*

If $\langle H_l, H_h \rangle \models_{\mathcal{P}_{\Sigma,n}} \phi$ holds, then $\langle H_l, H_h \rangle$ is said to *satisfy* ϕ. If Σ and n are clear from the context, we may simply write $\models$ instead of $\models_{\mathcal{P}_{\Sigma,n}}$.

Definition 7 *Let ϕ be a preference axiom of $\mathcal{P}_{\Sigma,n}$. For histories H_l, H_h over Σ of maximum length n, we define*

$$H_l \preceq_\phi H_h \quad \textit{iff} \quad \langle H_l, H_h \rangle \models_{\mathcal{P}_{\Sigma,n}} \phi.$$

The use of the symbol '$\preceq_\phi$' is purely suggestive at this point, since $\preceq_\phi$ may have none of the properties of an ordering.

We give some illustrations next.

Example 1 *The formula*

$$\phi = \big(\boldsymbol{h} : (\exists i f_1(i) \wedge \forall i \neg f_2(i))\big) \wedge \big(\boldsymbol{l} : (\exists i f_2(i) \wedge \forall i \neg f_1(i))\big)$$

expresses a preference of f_1 over f_2 in the sense that, for all histories H_l, H_h, we prefer H_h over H_l whenever it holds that H_h satisfies f_1 but not f_2, while H_l satisfies f_2 but not f_1.

This example can be more perspicuously given using our previous abbreviations (perhaps as part of a higher-level preference language) by:

$$\phi = \big(\boldsymbol{h} : (\Diamond f_1 \wedge \neg \Diamond f_2)\big) \wedge \big(\boldsymbol{l} : (\Diamond f_2 \wedge \neg \Diamond f_1)\big)$$

Example 2 *For fluent or action name b over a propositional action signature, and a variable i, let $min[b, i]$ abbreviate the formula*

$$\big(b(i) \wedge \forall k((k < i) \supset \neg b(k))\big).$$

For fluent or action name f, define

$$\phi = \exists i, j \big(\boldsymbol{l} : min[f, i] \wedge \boldsymbol{h} : min[f, j] \wedge (j < i)\big). \tag{3}$$

Then, $H_l \preceq_\phi H_h$ holds iff f is true in both H_l and H_h, but f is established earlier in H_h than in H_l.

Analogously, for fluent or action names f_1, f_2, define

$$\begin{aligned} \phi' \;=\; & \big(\boldsymbol{l} : \exists i, j((i \geq j) \wedge min[f_1, i] \wedge min[f_2, j])\big) \\ & \wedge \big(\boldsymbol{h} : \exists i, j((i < j) \wedge min[f_1, i] \wedge min[f_2, j])\big). \end{aligned} \tag{4}$$

Then, $H_l \preceq_{\phi'} H_h$ holds iff both f_1, f_2 are true in H_l and H_h, but f_1 is established before f_2 in H_h while this is not the case in H_l.

In an easy extension of the preceding example, we can express that we prefer first that f_1 and f_2 occur together, and then that f_1 occurs before f_2. As well, conditional preferences are trivially expressible in our approach. Indeed, a preference of β over α whenever γ holds can be realized in terms of the formula

$$(\boldsymbol{l} : \gamma \wedge \alpha) \wedge (\boldsymbol{h} : \gamma \wedge \beta).$$

So, for instance,

$$\begin{aligned} & \big(\boldsymbol{l} : \exists i, j((i < j) \wedge min[a, i] \wedge min[b, j])) \wedge f_1\big) \\ & \quad \wedge \big(\boldsymbol{h} : \exists i, j((i < j) \wedge min[a, i] \wedge min[b, j])) \wedge f_2\big). \end{aligned}$$

expresses a preference of f_2 over f_1 given that a occurs before b.

Having the preference-specification language at hand, we formally define a *preference frame* as follows:

Definition 8 *Let Σ be an action signature and $n \geq 0$.*
A preference frame over Σ with horizon n *is a pair $\langle \boldsymbol{H}, \boldsymbol{P} \rangle$, where*

- *$\boldsymbol{H}$ is a set of histories over Σ having maximum length n, and*
- *$\boldsymbol{P}$ is a set of preference axioms over $\mathcal{P}_{\Sigma,n}$.*

The question then is how to select maximally preferred histories, given a preference frame $\langle \boldsymbol{H}, \boldsymbol{P} \rangle$. If $\boldsymbol{P}$ contains more than one axiom, this question involves the general problem of combining different relations and is actually independent from the concrete form of our preference language. We say more on this next.

4.4 From Preferences to Ordering on Histories

We consider here how, given a preference frame, one may determine those histories that are maximally preferred. In a preference frame, each preference formula defines a binary relation whose instances are pairs of relatively less- and more-preferred histories. Thus, one can express various independent preference relations that must in some sense be combined in order to come up with maximally preferred histories. However this problem, of combining differing preference orderings, is a general and difficult problem in and of itself, and is the object of ongoing research involving areas such as multiple criteria decision making and social choice theory [27, 31]. Nonetheless, it is instructive to consider ways in which one may determine an overall preference ordering on histories, given a preference frame.

To begin, we can identify two base or generic approaches for determining (maximally) preferred histories. Recall that each $\phi \in \boldsymbol{P}$ induces a binary relation over $\boldsymbol{H}$ by setting $H_l \preceq_\phi H_h$ iff $\langle H_l, H_h \rangle \models_{\mathcal{P}_{\Sigma,n}} \phi$. In what follows, for any binary relation R we denote the reflexive and transitive closure of R by R^+.

As a base approach we can define:

Definition 9 *Let $\langle \boldsymbol{H}, \boldsymbol{P} \rangle$ be a preference frame over action signature Σ with horizon $n \geq 0$.*
Then, $H \in \boldsymbol{H}$ is a (general) maximally preferred history *iff H is a maximal element of*

$$\left(\bigcup_{\phi \in \boldsymbol{P}} \preceq_\phi \right)^+ .$$

Similarly, we can define a base approach founded on cardinality:

Definition 10 *Let $(\boldsymbol{H}, \boldsymbol{P})$ be a preference frame over action signature Σ with horizon $n \geq 0$. Furthermore, for $H \in \boldsymbol{H}$, let*

$$c(H) = |\{H' \in \boldsymbol{H} \mid \phi \in \boldsymbol{P} \text{ and } H' \preceq_\phi H\}|.$$

Then, $H \in \boldsymbol{H}$ is a (general) cardinality-based maximally preferred history *iff there is no $H' \in \boldsymbol{H}$ such that $c(H) < c(H')$.*

We do not give a full discussion of the above approaches here—rather, for illustrative purposes, we elaborate on the cardinality-based approach for propositional action signatures.

Consider where one is given a set of desirable outcomes, and the goal is to determine the history which satisfies the maximum number. Examples include fluents which are simply preferred to be true somewhere in a history, and temporal preferences in which one prefers that (pairs of) fluents become true in a specific order. In such cases, one wants to maximise the set of these desiderata. Assume then that we are given a set of (for simplicity) fluents $D = \{f, g, h, \dots\}$, where we wish to prefer a history in which as many of these fluents are true as possible.

Given a set of histories $\boldsymbol{H}$ and preferences D, we define a suitable preference framework $\langle \boldsymbol{H}, \boldsymbol{P} \rangle$, where

$$\boldsymbol{P} = \{(\boldsymbol{l} : \Box \neg d) \wedge (\boldsymbol{h} : \Diamond d) \mid d \in D\}. \tag{5}$$

Definition 10 yields a total preorder on histories, the maximal elements of which constitute the set of preferred histories. A refinement of this approach is to use set containment on satisfied preferences, rather than cardinality:

Definition 11 *Let $\langle \boldsymbol{H}, \boldsymbol{P} \rangle$ be a preference frame over action signature Σ with horizon $n \geq 0$, and assume that $\boldsymbol{P}$ is given by (5). For $H \in \boldsymbol{H}$, let*

$$s(H) = \{H' \in \boldsymbol{H} \mid \phi \in \boldsymbol{P} \textit{ and } H' \preceq_\phi H\}.$$

Then, $H \in \boldsymbol{H}$ is a (general) set containment-based maximally preferred history *iff there is no $H' \in \boldsymbol{H}$ such that $s(H) \subset s(H')$.*

Example 3 *Consider a preference frame as in (5), where we have simple preferences given by the set $D = \{f, g, h\}$. Assume that we have histories H_1, H_2, and H_3, such that the following relations hold:*

$$H_1 \models \Diamond f \wedge \Diamond h, \quad H_2 \models \Diamond g, \quad \textit{and} \quad H_3 \models \Diamond h.$$

According to Definition 10, H_1 is preferred; according to Definition 11, H_1 and H_2 are preferred.

Besides the base approaches for determining maximally preferred histories, as discussed above, another possibility for generating a global ordering on histories would be to employ methods based on the approach by Brewka [6] for building complex combinations of different preference strategies. An elaboration of such techniques is an issue for future work.

5 Modelling Choice and Temporal Preferences

Having introduced our general framework for specifying preference relations between histories, we now turn to the issue of modelling concrete types of preference relations in the context of reasoning about actions, thereby illustrating the usability and generality of our approach. Specifically, in this section, we discuss strategies for dealing with choice and temporal preferences and show how these can be captured in terms of preference frames.

As argued in Section 2, choice and temporal preferences deal with different ways in which subgoals may be attained. For a choice preference, the idea is that some condition is preferred *over* another (like preferring soup to grubs as an appetizer); whereas for a temporal preference the idea is that some condition should hold *before* another (like preferring having an appetizer before a main course).

Choice and temporal preferences were introduced in [7], in terms of *prioritised transition systems*. Intuitively, a prioritised transition system is a triple of form $\langle T, \leq_c, \leq_t \rangle$, where

- T is a transition system [19], i.e., a specification consisting of states, assignments of values at states, and a transition relation between states, and
- $\leq_c$ and $\leq_t$ are partial preorders, defined between Boolean combinations of action and fluent names, representing choice and temporal preferences, respectively.[3]

A transition system induces a set of possible histories, and the task of the choice and temporal preferences is to select preferred histories respecting these preferences. Here, our focus is on ways in which to discriminate among histories, rather than on ways that a transition system, together with choice and temporal preferences, induce preferred histories. As a consequence, we will base our discussion not on prioritised transition systems but rather on *prioritised history sets*, which are triples of the form $\langle \boldsymbol{H}, \leq_c, \leq_t \rangle$, where $\boldsymbol{H}$ is a set of histories and $\leq_c$ and $\leq_t$ are as above.

Intuitively, given a prioritised history set, a choice or temporal preference $\leq$ induces a relation $\preceq$ on *histories* such that the maximal elements of the latter are viewed as the preferred histories. Depending on whether $\leq$ models a choice or temporal preference, the way in which $\preceq$ is obtained from $\leq$ differs. In fact, there is not a unique method in which a choice or temporal preference $\leq$ determines a preference relation $\preceq$ on histories; rather, different desiderata require different strategies. The goal of this section is to consider specific strategies and to express them in terms of preference frames. In particular, we use strategies introduced in earlier work [7] as well a novel strategy for choice preferences.

We start with the general setting in which choice and temporal preferences are formulated.

5.1 Prioritised History Sets

In what follows, we assume a propositional action signature $\Sigma = \langle V, F, A \rangle$. We define the language $\mathcal{B}_\Sigma$ consisting of formulas formed from elements from $F \cup A$ as atomic formulas, and using the standard Boolean connectives in the usual recursive fashion. Note that formulas of $\mathcal{B}_\Sigma$ do not employ explicit time points, but just fluent and action names.

Definition 12 *Let Σ be a propositional action signature and $n \geq 0$.*

A prioritised history set (over Σ with horizon n) *is a triple*

$$\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle,$$

where $\boldsymbol{H}$ is set of histories over Σ having maximum length n and $\leq_c, \leq_t \subseteq \mathcal{B}_\Sigma \times \mathcal{B}_\Sigma$ are partial preorders.

The relation $\leq_c$ is called choice preference *and $\leq_t$ is called* temporal preference.

If $\boldsymbol{H}$, $\leq_c$, and $\leq_t$ are finite, then $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ is called finite.

We assume finite prioritised history sets only in what follows. As usual, for a partial preorder $\leq$ over some set S, its *strict part*, $<$, is defined by the condition that, for all $a, b \in S$, $a < b$ iff $a \leq b$ but $b \not\leq a$. Using preorders for preference relations has the advantage that one may distinguish between *indifference* (where both $a \leq b$ and $b \leq a$ hold) and *incomparability* (where neither $a \leq b$ nor $b \leq a$ holds). Intuitively, if $\alpha_1 \leq_c \alpha_2$, then we prefer histories in which α_2 is true no less than

[3]Recall that a partial preorder is a binary relation which is both reflexive and transitive.

histories in which α_1 is true. A temporal preference $\alpha_1 \leq_t \alpha_2$, on the other hand, specifies that, if possible, α_1 should become true not later than α_2 becoming true in a history.

Our first task is to translate formulas of $\mathcal{B}_\Sigma$ into formulas of $\mathcal{Q}_{\Sigma,n}$. This is done in the following way: given a formula $\alpha \in \mathcal{B}_\Sigma$ and a variable or natural number i, we associate with α the formula $\alpha[i] \in \mathcal{Q}_{\Sigma,n}$ which results from α by replacing each action name a by the atom $a(i)$ and each fluent name f by the atom $f(i) = 1$.[4] Furthermore, in analogy to the operators $\Box$ and $\Diamond$, as introduced in Section 4.2, we can associate to $\alpha, \beta \in \mathcal{B}_\Sigma$ the following formulas from $\mathcal{Q}_{\Sigma,n}$:

- $\Box\alpha = \forall i\, \alpha[i]$; and
- $\Diamond\alpha = \exists i\, \alpha[i]$.

Definition 13 *Let Σ be a propositional action signature and $n \geq 0$.*

Then, for any history H over Σ with $|H| \leq n$ and any $\alpha \in \mathcal{B}_\Sigma$, we define

$$H \models_{\mathcal{B}_\Sigma} \alpha \quad \textit{iff} \quad H \models_{\mathcal{Q}_{\Sigma,n}} \Diamond\alpha.$$

If $H \models_{\mathcal{B}_\Sigma} \alpha$, we say that α is *true* in history H, or that H *satisfies* α. We sometimes write $\models$ for $\models_{\mathcal{B}_\Sigma}$ if Σ is clear from the context. The definition of $\models_{\mathcal{B}_\Sigma}$ allows, for example, having both $H \models_{\mathcal{B}_\Sigma} \alpha$ and $H \models_{\mathcal{B}_\Sigma} \neg\alpha$ for some α, while $H \models_{\mathcal{B}_\Sigma} \alpha \wedge \neg\alpha$ will not hold.

Given a choice or temporal preference $\leq$, the main issue is to determine those histories which are maximally preferred under $\leq$. This is realised in terms of the following concept:

Definition 14 *Let Σ be a propositional action signature and $\boldsymbol{H}$ a set of histories over Σ.*

A preference strategy, *σ, (over Σ and $\boldsymbol{H}$) is a mapping assigning to each partial preorder $\leq$ over $\mathcal{B}_\Sigma$ a partial preorder $\sigma(\leq)$ over $\boldsymbol{H}$.*

We now are in a position to define maximal histories:

Definition 15 *Let Σ be a propositional action signature, $\langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ a prioritised history set over Σ, σ a preference strategy over Σ, and $\leq \in \{\leq_c, \leq_t\}$.*

A history $H \in \boldsymbol{H}$ is preferred under $\leq$ (relative to σ) *iff H is a maximal element of $\sigma(\leq)$.*

Our main goal can now be restated more formally as follows: given a set $\boldsymbol{H}$ of histories over some action signature Σ and a choice or temporal preference $\leq$ over $\mathcal{B}_\Sigma$, along with a preference strategy σ, we seek a preference frame $\langle \boldsymbol{H}, \boldsymbol{P}_{\leq,\sigma} \rangle$ such that the histories preferred under $\leq$ relative to σ are given by the histories maximally preferred under $\langle \boldsymbol{H}, \boldsymbol{P}_{\leq,\sigma} \rangle$.

[4]Since we assume propositional action signatures here, following our convention from Section 4.2, the atom $f(i) = 1$ may be identified with $f(i)$.

5.2 Expressing Choice Preferences

We first introduce the preference strategy σ_c, following the method proposed in [7]; afterwards we discuss how choice preferences under this strategy can be expressed in our framework.

To begin with, for any binary relation R, define

$$dom(R) = \{x, y \mid \langle x, y \rangle \in R\}.$$

Furthermore, for any prioritised history set $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ and histories $H, H' \in \boldsymbol{H}$, define

$$\Delta_{\mathcal{H}}(H, H') \;\; = \;\; \{\alpha \in dom(\leq_c) \mid H \models \alpha \text{ and } H' \not\models \alpha\}.$$

That is, $\Delta_{\mathcal{H}}(H, H')$ consists of all formulas related by $\leq_c$ which are satisfied at some point in H but never in H'.

Next, we define an intermediate relation between histories, on which the construction of σ_c rests.

Definition 16 *Let $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ be a prioritised history set.*

Then, for any two histories $H, H' \in \boldsymbol{H}$, we define $H \trianglelefteq_c^{\mathcal{H}} H'$ iff, for any formula $\alpha \in \Delta_{\mathcal{H}}(H, H')$, there is some $\alpha' \in \Delta_{\mathcal{H}}(H', H)$ such that $\alpha \leq_c \alpha'$.

If $\mathcal{H}$ is unambiguously fixed, we write $H \trianglelefteq_c H'$ instead of $H \trianglelefteq_c^{\mathcal{H}} H'$. The reason for using $\Delta_{\mathcal{H}}(H, H')$ in the above definition is that we are only interested in formulas which are not jointly satisfied by the two histories H and H'. A similar construction using "difference sets", though defined on different objects of discourse, was used by Geffner and Pearl [17].

Clearly, $\trianglelefteq_c$ is reflexive, that is, we have $H \trianglelefteq_c H$ for any history H. However, $\trianglelefteq_c$ is generally not transitive. To see this, consider an action signature Σ involving fluents, f, g, and h, and a prioritised history set $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ over Σ such that

- $f \leq_c g \leq_c h$ and $h \leq_c g$, and
- $\boldsymbol{H}$ consists of three histories, H, H', and H'', satisfying the following relations:

$$\begin{array}{lll} H \models f, & H \not\models g, & H \models h; \\ H' \not\models f, & H' \models g, & H' \not\models h; \\ H'' \not\models f, & H'' \not\models g, & H'' \models h. \end{array}$$

From this, we get the following sets of differing fluents:

$$\begin{array}{rr} \Delta_{\mathcal{H}}(H, H') = \{f, h\}; & \Delta_{\mathcal{H}}(H', H) = \{g\}; \\ \Delta_{\mathcal{H}}(H', H'') = \{g\}; & \Delta_{\mathcal{H}}(H'', H') = \{h\}; \\ \Delta_{\mathcal{H}}(H, H'') = \{f\}; & \Delta_{\mathcal{H}}(H'', H) = \emptyset \; . \end{array}$$

Then, it is easy to check that both $H \trianglelefteq_c H'$ and $H' \trianglelefteq_c H''$ hold, but not $H \trianglelefteq_c H''$.

In view of the non-transitivity of $\trianglelefteq_c$, the relation $\sigma_c(\leq_c)$, for a given choice preference $\leq_c$, is then defined as the transitive closure of $\trianglelefteq_c$. To this end, let R^* to denote the transitive closure of R, for any binary relation R.

Definition 17 *The strategy σ_c assigns to each choice preference $\leq_c$ of a prioritised history set $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ the relation $\preceq_c$, given as follows:*

$$H \preceq_c H' \quad \textit{iff} \quad H \trianglelefteq_c^* H', \qquad \textit{for every } H, H' \in \boldsymbol{H}.$$

Given the properties of $\trianglelefteq_c^*$, the relation $\preceq_c$ is clearly a partial preorder. Note that $\preceq_c$ may posses non-trivial cycles, i.e., there may exist pairwise distinct histories $H_1, \ldots, H_k$, for $k > 1$, such that, for all $i \in \{1, \ldots, k-1\}$, $H_i \trianglelefteq_c H_{i+1}$ and $H_k \trianglelefteq_c H_1$. However, the strict part $\prec_c$ of $\preceq_c$ is always cycle free.[5] If one wants that $\preceq_c$ does not have any non-trivial cycles, one can replace Definition 17 by setting $H \preceq_c H'$ iff $H \triangleleft_c^* H'$ or $H = H'$, where $\triangleleft_c^*$ is the strict part of the transitive closure of $\trianglelefteq_c$.

Any history H satisfying Definition 17 would seem to be undeniably "preferred". However, other definitions are certainly possible. For example, one could rank histories by the number of choice preferences violated. This alternative would make sense where the preferences were *absolute*, i.e., of the form $\neg\alpha <_c \alpha$. Other alternatives can be obtained by using variations of the relation $\trianglelefteq_c$. One way to realise such variants is to change the basic "$\forall$-$\exists$" quantification in Definition 16 for relating the elements in $\Delta_P(H, H')$ and $\Delta_P(H', H)$ by other quantification forms. For example, the following relations may be defined:

- $H \trianglelefteq_{c,\forall\forall} H'$ iff, for all $\alpha \in \Delta_P(H, H')$ and all $\beta \in \Delta_P(H', H)$, $\alpha \leq_c \beta$ holds.
- $H \trianglelefteq_{c,\exists\forall} H'$ iff there is some $\beta \in \Delta_P(H', H)$ such that, for all $\alpha \in \Delta_P(H, H')$, $\alpha \leq_c \beta$ holds.
- $H \trianglelefteq_{c,\exists\exists} H'$ iff there is some $\alpha \in \Delta_P(H, H')$ and some $\beta \in \Delta_P(H', H)$ such that $\alpha \leq_c \beta$ holds.

These relations are in general incompatible with $\trianglelefteq_c$, except for $\trianglelefteq_{c,\exists\forall}$ (which satisfies $\trianglelefteq_{c,\exists\forall} \subseteq \trianglelefteq_c$), and are mostly too weak to yield satisfactory properties. In any case, the basic patterns $\trianglelefteq_c$ and its variants described above mirror the general problem concerning how one can move from an order relation defined between elements of a set S to an order relation defined between elements of the *power set* of S (i.e., between subsets of S), for which there is no unique solution. Such a problem is encountered in various applications and solutions corresponding to the above patterns have been studied in the literature (see, e.g., Halpern [24] for a discussion about this issue and related approaches). Many other alternatives for $\preceq_c$ are obtainable within our framework; later on, we detail further such an alternative.

Before describing the encoding of σ_c in terms of preference frames, let us consider some basic examples illustrating the choice preference order $\preceq_c$. The first example describes the most obvious case in which two histories are distinguished.

Example 4 *Consider a prioritised history set, $\mathcal{H}_4$, over fluents f and g such that $f \leq_c g$, and assume histories H and H' satisfying $H \models f$, $H \not\models g$, $H' \not\models f$, and $H' \models g$.*

Then, we have that

$$\Delta_{\mathcal{H}_4}(H, H') = \{f\} \quad \textit{and} \quad \Delta_{\mathcal{H}_4}(H', H) = \{g\}.$$

It follows that $H \trianglelefteq_c H'$, and thus $H \preceq_c H'$. In fact, it holds that $H \prec_c H'$.

[5] Indeed, it is a straightforward matter to verify that, for any preorder $\leq$, its strict part $<$ possesses no cycles.

The next example illustrates the particular interpretation of choice preference under σ_c.

Example 5 *Let $\mathcal{H}_5$ be a prioritised history set, comprised again of fluents f and g, and ordered by $f \leq_c g$, and consider histories H and H', where $H \models f$ and $H \not\models g$ as before, but H' now satisfies $H' \not\models f$ and $H' \not\models g$.*

We obtain

$$\Delta_{\mathcal{H}_5}(H, H') = \{f\} \quad \textit{and} \quad \Delta_{\mathcal{H}_5}(H', H) = \emptyset.$$

Therefore, we get that $H \not\trianglelefteq_c H'$ and $H' \trianglelefteq_c H$.

Informally, this example illustrates the type of choice preference that σ_c implements: A preferred history with respect to $\mathcal{H}$ is one that satisfies the choice preferences as much as possible, but disfavours histories like H' with an empty set $\Delta_{\mathcal{H}}(H', H)$ of distinguishing preference formulas. Observe that no relation among H and H' is obtained in the aforementioned $\trianglelefteq_{c,\exists\exists}$ variant of $\trianglelefteq_c$.

Example 6 *Let $\mathcal{H}_6$ be a prioritised history set defined similarly to those in Examples 4 and 5, and consider H and H' such that $H \models f$, $H \models g$, $H' \not\models f$, and $H' \models g$.*

Then,

$$\Delta_{\mathcal{H}_6}(H, H') = \{f\} \quad \textit{and} \quad \Delta_{\mathcal{H}_6}(H', H) = \emptyset.$$

Both histories agree on (in fact, satisfy) the $\leq_c$-higher fluent, but differ on the $\leq_c$-lesser fluent. The result here is the same as in the previous example.

Just as $\neg\alpha \leq_c \alpha$ expresses an "absolute" preference for α, the reverse condition $\alpha \leq_c \neg\alpha$ expresses a "negative" preference for α in the sense that those histories are preferred in which α is avoided.

Let us consider a more involved example now, based on the well-known monkey-and-bananas scenario.

Example 7 *A monkey wants a bunch of bananas, hanging from the ceiling, or a coconut, found on the floor; as well, the monkey wants a chocolate bar, found in a drawer.*

In order to get the bananas, the monkey must push a box to the empty place under the bananas and then climb on top of the box. In order to get the chocolate, the drawer must be opened. Each object is initially at a different location.

We assume that the monkey wants the chocolates, and either the coconuts or the bananas, and he prefers bananas over coconuts.

Formally, we use a prioritised history set $\langle \boldsymbol{H}, \leq_c, \emptyset \rangle$ over a propositional action signature $\Sigma = \langle \{0,1\}, F, A \rangle$, specified as follows:

$$\begin{aligned}
F \quad &= \quad \big\{ loc(I, l_i) \mid I \in \{Monkey, Box, Ban, Drawer, Coco\}, 1 \leq i \leq 5 \big\} \\
&\quad\ \cup \{onBox, hasBan, hasChoc, hasCoco\}; \\
A \quad &= \quad \{walk(l_i), pushBox(l_i) \mid 1 \leq i \leq 5\} \\
&\quad\ \cup \{climbOn, climbOff, graspBan, graspChoc, graspCoco, openDrawer\}; \\
\leq_c \quad &= \quad \{hasCoco \leq_c hasBan\}.
\end{aligned}$$

The set $\boldsymbol{H}$ is assumed to be given as the collection of all histories determined by a concrete transition system T [19] based on the action language $\mathcal{C}$ [22] satisfying the query

$$Q = \mathit{hasChoc} \wedge (\mathit{hasBan} \vee \mathit{hasCoco})[7].$$

For the sake of brevity, we omit the full (and straightforward) details about the specification of $\boldsymbol{H}$.

Initially, the monkey does not have the chocolates, bananas, or coconuts, and each object is at a different location. There are, among others, two histories, H and H', satisfying Q:[6]

History H	*Action*
STATE 0:	*go to the drawer*
STATE 1:	*open the drawer*
STATE 2:	*grasp the chocolates*
STATE 3:	*walk to the box*
STATE 4:	*push the box to the bananas*
STATE 5:	*climb on the box*
STATE 6:	*grasp the bananas*

History H'	*Action*
STATE 0:	*go to the drawer*
STATE 1:	*open the drawer*
STATE 2:	*grasp the chocolates*
STATE 3:	*walk to the coconuts*
STATE 4:	*grasp the coconuts*

Given the monkey's preference of bananas over coconuts, we expect that H is preferred over H'. This is indeed the case, as it can be shown that H is $\preceq_c$-preferred, but H' is not. For this, observe that $\mathit{hasCoco} \in \Delta_P(H, H')$ as well as $\mathit{hasBan} \in \Delta_P(H, H')$.

We note that there are of course more histories satisfying the intended goal if we consider histories of length greater than 7. In particular, there are histories satisfying

$$Q' = \mathit{hasChoc} \wedge (\mathit{hasBan} \vee \mathit{hasCoco})[8]$$

in which the subgoals are achieved in the reverse order as given by H and H' (cf. Example 9 below).

We now reconstruct σ_c by means of preference frames. The following definition is central:

Definition 18 *Let Σ be a propositional action signature and $n \geq 0$.*

For every prioritised history set $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$, let $\Phi_c^{\mathcal{H}}$ be the following formula of $\mathcal{P}_{\Sigma,n}$:

$$\Phi_c^{\mathcal{H}} = \bigwedge_{\alpha \in dom(\leq_c)} \Big((\boldsymbol{h} : \neg\Diamond\alpha \wedge \boldsymbol{l} : \Diamond\alpha) \supset \bigvee_{\beta \in \{\gamma \mid \alpha \leq_c \gamma\}} (\boldsymbol{h} : \Diamond\beta \wedge \boldsymbol{l} : \neg\Diamond\beta) \Big).$$

Note that $\Phi_c^{\mathcal{H}}$ is well defined, since both $dom(\leq_c)$ and $\{\gamma \mid \alpha \leq_c \gamma\}$, for any formula α, are finite given our assumption of prioritised history sets being finite. This formula captures the condition expressed in Definition 16: the subformulas $(\boldsymbol{h} : \neg\Diamond\alpha \wedge \boldsymbol{l} : \Diamond\alpha)$ and $(\boldsymbol{h} : \Diamond\beta \wedge \boldsymbol{l} : \neg\Diamond\beta)$ refer to formulas α and β belonging to $\Delta_{\mathcal{H}}(H, H')$ and $\Delta_{\mathcal{H}}(H', H)$, respectively, while the consequents of the implications, present for each formula α in the domain of $\leq_c$, select at least one $\beta \in \Delta_{\mathcal{H}}(H', H)$ with $\alpha \leq_c \beta$. This is made precise by the following result:

[6] Observe that, strictly speaking, we have to further assume the tick-of-the-clock action $do\_nothing$ and apply it twice after State 4 in H' in order that evaluating Q at H' is defined.

Lemma 1 *Let $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ be a prioritised history set over a propositional action signature Σ with horizon n.*

Then, for any $H, H' \in \boldsymbol{H}$, we have that

$$H \trianglelefteq_c^{\mathcal{H}} H' \quad \textit{iff} \quad \langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \Phi_c^{\mathcal{H}}.$$

Proof. Suppose that $H \trianglelefteq_c^{\mathcal{H}} H'$. We show that $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \Phi_c^{\mathcal{H}}$. Consider some $\alpha \in dom(\leq_c)$ and assume that $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} (\boldsymbol{h} : \neg\Diamond\alpha \wedge \boldsymbol{l} : \Diamond\alpha)$. Then, $H \models_{\mathcal{Q}_{\Sigma,n}} \Diamond\alpha$ and $H' \models_{\mathcal{Q}_{\Sigma,n}} \neg\Diamond\alpha$, which implies that $H \models_{\mathcal{B}_\Sigma} \alpha$ and $H' \models_{\mathcal{B}_\Sigma} \neg\alpha$, and thus $\alpha \in \Delta_{\mathcal{H}}(H, H')$. Since $H \trianglelefteq_c^{\mathcal{H}} H'$, it follows that there is some $\alpha' \in \Delta_{\mathcal{H}}(H', H)$ such that $\alpha \leq_c \alpha'$. By the former condition, we have that $H' \models_{\mathcal{B}_\Sigma} \alpha'$ and $H \models_{\mathcal{B}_\Sigma} \neg\alpha'$, from which we obtain that $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} (\boldsymbol{h} : \Diamond\alpha' \wedge \boldsymbol{l} : \neg\Diamond\alpha')$. But $\alpha \leq_c \alpha'$ means that $\alpha' \in \{\gamma \mid \alpha <_c \gamma\}$ and so $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \bigvee_{\beta \in \{\gamma \mid \alpha \leq_c \gamma\}} (\boldsymbol{h} : \Diamond\beta \wedge \boldsymbol{l} : \neg\Diamond\beta)$ holds as well. Consequently,

$$\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} (\boldsymbol{h} : \neg\Diamond\alpha \wedge \boldsymbol{l} : \Diamond\alpha) \supset \bigvee_{\beta \in \{\gamma \mid \alpha \leq_c \gamma\}} (\boldsymbol{h} : \Diamond\beta \wedge \boldsymbol{l} : \neg\Diamond\beta).$$

Since α was chosen as an arbitrary element of $dom(\leq_c)$, $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \Phi_c^{\mathcal{H}}$ follows. This proves that $H \trianglelefteq_c^{\mathcal{H}} H'$ only if $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \Phi_c^{\mathcal{H}}$.

For the converse direction, assume that $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \Phi_c^{\mathcal{H}}$. We show that for any $\alpha \in \Delta_{\mathcal{H}}(H, H')$ there is some $\alpha' \in \Delta_{\mathcal{H}}(H', H)$ such that $\alpha \leq_c \alpha'$. So, fix an $\alpha \in \Delta_{\mathcal{H}}(H, H')$. Then, we have that $\alpha \in dom(\leq_c)$, as well as $H \models_{\mathcal{B}_\Sigma} \alpha$ and $H' \models_{\mathcal{B}_\Sigma} \neg\alpha$. The latter two conditions imply that $H \models_{\mathcal{Q}_{\Sigma,n}} \Diamond\alpha$ and $H' \models_{\mathcal{Q}_{\Sigma,n}} \neg\Diamond\alpha$, from which $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} (\boldsymbol{h} : \neg\Diamond\alpha \wedge \boldsymbol{l} : \Diamond\alpha)$ follows. Given the hypothesis that $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \Phi_c^{\mathcal{H}}$, we in turn obtain that

$$\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \bigvee_{\beta \in \{\gamma \mid \alpha \leq_c \gamma\}} (\boldsymbol{h} : \Diamond\beta \wedge \boldsymbol{l} : \neg\Diamond\beta)$$

must hold. Hence, by the semantics of disjunction, there is some $\alpha' \in \mathcal{B}_\Sigma$ such that $\alpha \leq_c \alpha'$ and $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} (\boldsymbol{h} : \Diamond\alpha' \wedge \boldsymbol{l} : \neg\Diamond\alpha')$. The latter yields that $H' \models_{\mathcal{B}_\Sigma} \alpha'$ and $H \models_{\mathcal{B}_\Sigma} \neg\alpha'$, which means that $\alpha' \in \Delta_{\mathcal{H}}(H', H)$. Consequently, from our assumption that $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \Phi_c^{\mathcal{H}}$, we derived that for any $\alpha \in \Delta_{\mathcal{H}}(H, H')$ there is some $\alpha' \in \Delta_{\mathcal{H}}(H', H)$ such that $\alpha \leq_c \alpha'$. In other words, $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \Phi_c^{\mathcal{H}}$ entails $H \trianglelefteq_c^{\mathcal{H}} H'$. ∎

Now, according to Definition 15, given $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$, a history $H \in \boldsymbol{H}$ is preferred under $\leq_c$ relative to σ_c iff H is a maximal element of $\sigma_c(\leq_c)$. The latter relation, in turn, is defined as the transitive closure of $\trianglelefteq_c^{\mathcal{H}}$. By Definition 7 and Lemma 1, and since $\trianglelefteq_c^{\mathcal{H}}$ is reflexive, we therefore obtain that the transitive closure of $\trianglelefteq_c^{\mathcal{H}}$ and the reflexive and transitive closure of $\preceq_{\Phi_c^{\mathcal{H}}}$ coincide. Hence, in view of Definition 9, it follows that H is a maximal element of $\sigma_c(\leq_c)$ iff H is a maximally preferred history of the preference frame $\langle \boldsymbol{H}, \{\Phi_c^{\mathcal{H}}\} \rangle$. This proves the following result:

Theorem 2 *Let $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ be a prioritised history set over a propositional action signature Σ.*

Then, for any $H \in \boldsymbol{H}$, H is preferred under $\leq_c$ relative to σ_c iff H is a maximally preferred history of the preference frame $\langle \boldsymbol{H}, \{\Phi_c^{\mathcal{H}}\} \rangle$.

5.3 Expressing Temporal Preferences

With choice preference, the order $\leq_c$ specifies the relative desirability that a formula be true in a history; thus, $\alpha_2 \leq_c \alpha_1$ implicitly expresses a preference that holds *between* histories. For temporal preferences, the order $\leq_t$ specifies the desired order in which formulas become true *within* a history. Thus, $\alpha_2 \leq_t \alpha_1$ implicitly expresses a preference that the establishment of α_2 is not later than that of α_1. Hence, a history in which α_2 becomes true before α_1 is preferred to one where this is not the case. To this end, for $\leq_t$, it is convenient to be able to refer to the ordering on formulas given by a history. The following definition is taken from [7]; clearly other definitions may serve equally well, depending on the intended application.

Definition 19 *For a propositional action signature Σ and a history H over Σ of length n, define $\leq_H \subseteq \mathcal{B}_\Sigma \times \mathcal{B}_\Sigma$ by $\alpha_1 \leq_H \alpha_2$ iff there are i, j, where $i < j$ and*

1. *$H \models_{\mathcal{Q}_{\Sigma,n}} \alpha_1[i]$ and $H \models_{\mathcal{Q}_{\Sigma,n}} \alpha_2[j]$, and*
2. *for any $i' < i$ and any $j' < j$, we have $H \models_{\mathcal{Q}_{\Sigma,n}} \neg\alpha_1[i']$ and $H \models_{\mathcal{Q}_{\Sigma,n}} \neg\alpha_2[j']$.*

We want to compare histories, say H and H', based on "how well" $\leq_H$ and $\leq_{H'}$ agree with $\leq_t$. Below we introduce the strategy σ_t implementing a way to achieve this.

First, a history H will be temporally preferred if $\leq_H$ does not disagree with $\leq_t$; that is, if $<_t \cap \leq_H^{-1} = \emptyset$.[7] We extend this to relative preference among histories as follows:

Definition 20 *The strategy σ_t is defined as the mapping assigning each temporal preference $\leq_t$ of a prioritised history set $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ the relation $\preceq_t$, given as follows:*

$$H \preceq_t H' \quad \textit{iff} \quad \left(<_t \cap \leq_{H'}^{-1}\right) \subseteq \left(<_t \cap \leq_H^{-1}\right) , \qquad \textit{for every } H, H' \in \boldsymbol{H}.$$

That is, H' violates fewer preferences in $<_t$ than H does. Obviously, $\preceq_t$ is a partial preorder on $\mathcal{H}$.

Example 8 *Consider a prioritised history set in which we are given $\alpha_1 \leq_t \alpha_2$ only, where α_1 and α_2 are two distinct formulas, and in which histories H_1, H_2, and H_3 are such that:*

- *H_1 satisfies the given preference, in that α_1 becomes true prior to α_2;*
- *in H_2, α_2 becomes true prior to α_1; and*
- *α_1 does not become true in H_3.*

According to Definition 19, we have

$$\begin{aligned} <_t \cap \leq_{H_1}^{-1} &= \emptyset, \\ <_t \cap \leq_{H_2}^{-1} &= \{\langle \alpha_1, \alpha_2 \rangle\}, \quad \textit{and} \\ <_t \cap \leq_{H_3}^{-1} &= \emptyset. \end{aligned}$$

H_1 and H_3 are temporally preferred histories, since neither violates the preference in $<_t$.

[7] For binary relation R, R^{-1} is the relation satisfying $\langle x, y \rangle \in R^{-1}$ iff $\langle y, x \rangle \in R$.

Example 9 *Consider again the monkey-and-bananas scenario from Example 7. Assume now a prioritised history set $\langle \boldsymbol{H}, \emptyset, \leq_t \rangle$, where $\leq_t$ consists of just the pair*

$$hasBan \leq_t hasChoc,$$

and $\boldsymbol{H}$ is the collection of all histories induced by the transition system T as in Example 7 satisfying the query

$$Q' = hasChoc \wedge (hasBan \vee hasCoco)[8].$$

The temporally preferred histories are those in which the bananas are obtained and then chocolate, and those in which coconuts and bananas are obtained. Another way of saying this is that the histories that are not temporally preferred are those violating $hasBan \leq_t hasChoc$.

If we add the preference $hasCoco \leq_t hasChoc$, then the preferred histories are those where one of bananas or coconuts are obtained, and then chocolate. If we combine this preference with the choice preference in Example 7, we obtain a history in which both preferences can be satisfied. If the preferences were to conflict, then obviously only one can be satisfied; nonetheless, clearly such a conflict does not prevent us from finding a successful plan. As before, there are four histories satisfying Q'. Of these three are temporally preferred:

- *one in which bananas are obtained and then chocolate, and*
- *two others in which coconuts and bananas are obtained.*

Another way of saying this is that there is one history that is not temporally preferred, and that is the history that violates $hasBan \leq_t hasChoc$. If we add preference $hasCoco \leq_t hasChoc$, then there are two preferred histories, where first bananas or coconuts are obtained, and then chocolate.

We capture temporal preferences under σ_t in our framework as follows. In analogy to the formula $min[b, i]$ from Example 2, defined for an action or fluent name b and a variable i, we introduce, for every $\alpha \in \mathcal{B}_\Sigma$, $min[\alpha, i]$ as an abbreviation for

$$\big(\alpha[i] \wedge \forall k((k < i) \supset \neg\alpha[k])\big).$$

Definition 21 *Let Σ be a propositional action signature and $n \geq 0$.*

1. *For every $\alpha, \beta \in \mathcal{B}_\Sigma$, let $\mathcal{T}[\alpha, \beta]$ be the following formula of $\mathcal{Q}_{\Sigma,n}$:*

$$\mathcal{T}[\alpha, \beta] = \exists i, j\big((i < j) \,\wedge\, min[\alpha, i] \,\wedge\, min[\beta, j]\big).$$

2. *For every prioritised history set $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$, let $\Phi_t^{\mathcal{H}} \in \mathcal{P}_{\Sigma,n}$ be given as follows:*

$$\Phi_t^{\mathcal{H}} = \bigwedge_{\alpha \in dom(<_t)} \Big(\bigwedge_{\beta \in \{\gamma | \alpha \leq_t \gamma\}} \big(\boldsymbol{h} : \mathcal{T}[\beta, \alpha] \supset \boldsymbol{l} : \mathcal{T}[\beta, \alpha]\big)\Big).$$

The central properties of these two formulas are as follows:

Lemma 2 *Let Σ be a propositional action signature and $n \geq 0$.*

1. *For every history H over Σ with $|H| \leq n$ and every $\alpha, \beta \in \mathcal{B}_\Sigma$, we have that*

$$\alpha \leq_H \beta \quad \textit{iff} \quad H \models_{\mathcal{Q}_{\Sigma,n}} \mathcal{T}[\alpha, \beta] \,.$$

2. *For every prioritised history set $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ over Σ with horizon n and every $H, H' \in \boldsymbol{H}$, we have that*

$$H \preceq_t H' \quad \textit{iff} \quad \langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \Phi_t^{\mathcal{H}} \,,$$

for $\sigma_t(\leq_t) = \preceq_t$.

Proof. Part 1 follows straightforwardly from the definition of $\leq_H$ and the properties of the module $min[\cdot, \cdot]$.

As regards Part 2, by Definition 20, we have that $H \preceq_t H'$ iff $<_t \cap \leq_{H'}^{-1} \subseteq <_t \cap \leq_{H'}^{-1}$, for every $H, H' \in \boldsymbol{H}$. But Part 1 implies that $\alpha \leq_H^{-1} \beta$ iff $H \models_{\mathcal{Q}_{\Sigma,n}} \mathcal{T}[\beta, \alpha]$, for every $\alpha, \beta \in \mathcal{B}_\Sigma$. So, for every $H, H' \in \boldsymbol{H}$, it holds that $\leq_{H'}^{-1} \subseteq \leq_H^{-1}$ is equivalent to the condition that $H \models_{\mathcal{Q}_{\Sigma,n}} \mathcal{T}[\beta, \alpha]$ whenever $H' \models_{\mathcal{Q}_{\Sigma,n}} \mathcal{T}[\beta, \alpha]$, for every $\alpha, \beta \in \mathcal{B}_\Sigma$, which in turn is equivalent to

$$\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \boldsymbol{h} : \mathcal{T}[\beta, \alpha] \supset \boldsymbol{l} : \mathcal{T}[\beta, \alpha],$$

for every $\alpha, \beta \in \mathcal{B}_\Sigma$. Now, observing further that the condition $\alpha <_t \beta$ holds precisely in case $\alpha \in dom(<_t)$ and $\beta \in \{\gamma \mid \alpha \leq_t \gamma\}$, it follows that $H \preceq_t H'$ iff $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \Phi_t^{\mathcal{H}}$. ∎

From this, we obtain the characterisation of σ_t in terms of preference frames as follows:

Theorem 3 *Let $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ be a prioritised history set over a propositional action signature Σ.*

Then, for every $H \in \boldsymbol{H}$, H is preferred under $\leq_t$ relative to σ_t iff H is a maximally preferred history of the preference frame $\langle \boldsymbol{H}, \{\Phi_t^{\mathcal{H}}\} \rangle$.

Proof. For every $H \in \boldsymbol{H}$, we have, on the one hand, that H is preferred under $\leq_t$ relative to σ_t iff H is a maximal element of $\sigma_t(\leq_t)$, and, on the other hand, that H is a maximally preferred history of $\langle \boldsymbol{H}, \{\Phi_t^{\mathcal{H}}\} \rangle$ iff H is a maximal element of the reflexive and transitive closure $\preceq_{\Phi_t^{\mathcal{H}}}^{+}$ of $\preceq_{\Phi_t^{\mathcal{H}}}$ (cf. Definitions 15 and 9, respectively). Now, since, by definition, $H \preceq_{\Phi_t^{\mathcal{H}}} H'$ iff $\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \Phi_t^{\mathcal{H}}$, for every $H, H' \in \boldsymbol{H}$, Part 2 of Lemma 2 tells us that $\sigma_t(\leq_t) = \preceq_{\Phi_t^{\mathcal{H}}}$. In view of the fact that $\sigma_t(\leq_t)$ is already reflexive and transitive, we therefore actually have that $\sigma_t(\leq_t) = \preceq_{\Phi_t^{\mathcal{H}}}^{+}$, and so the maximal elements of $\sigma_t(\leq_t)$ coincide with the maximal elements of $\preceq_{\Phi_t^{\mathcal{H}}}^{+}$. From this, the statement of the theorem follows. ∎

Clearly, various alternatives for temporal preference are obtainable by varying the underlying order $\leq_H$ in Definition 19. For instance, instead of using the minimal time point for selecting the earliest state of satisfaction, one may choose the maximal time point for focusing on the latest such states. As well, one may simply require there to be two (arbitrary) time points, so that one formula becomes true before the other. More elaborated orderings could even take into account the number of times a formula is satisfied before another, etc. All this is possible within our framework.

5.4 A Variant Choice Preference Strategy

The strategy σ_c discussed in Section 5.2 can be viewed as implementing a certain "absolute" notion of choice preference in the following sense: Suppose we only have a choice preference $f \leq_c g$ and two histories, H and H', such that $H \not\models f$, $H \not\models g$, $H' \models f$, and $H' \not\models g$. Then, $H \preceq_c H'$ holds. That is, even though the preferred fluent or action g is false in both H and H', the strategy σ_c assigns a preference to H', in effect implicitly assuming an "absolute" preference $\neg f \leq_c f$. In what follows, we describe a strategy, σ_c', realising a more agnostic kind of preference.

Let $\langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ be a prioritised history set over $\Sigma = \langle V, F, A \rangle$, and let H and H' be two histories of $\boldsymbol{H}$. We say that $\alpha \leq_c \beta$ *prefers* H' *to* H just if

$$H \models \alpha \text{ and } H \not\models \beta, \text{ but } H' \models \beta.$$

Thus, if the above condition holds, then $\alpha \leq_c \beta$ provides a reason for preferring one history over another. For example, if we have that $f \leq_c g$ and $g \leq_c h$, along with $H \not\models f$, $H \models g$, $H \not\models h$, $H' \models f$, $H' \not\models g$, and $H' \models h$, then $f \leq_c g$ prefers H to H' and $g \leq_c h$ prefers H' to H. Clearly though, if this is all the information that we have, then we would want to say that H' is preferred to H (and not vice versa) since H' alone satisfies the most preferred fluent h. The definition below generalises this in the obvious fashion:

Definition 22 *Let $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ be a prioritised history set.*
Then, for any histories $H, H' \in \boldsymbol{H}$, we define $H \ll_c^{\mathcal{H}} H'$ iff

1. *there are formulas $\alpha, \beta \in \mathcal{B}_\Sigma$ such that $\alpha \leq_c \beta$ prefers H' to H, and*
2. *for each $\phi, \psi \in \mathcal{B}_\Sigma$ such that $\phi \leq_c \psi$ prefers H to H', there is some $\mu \in \mathcal{B}_\Sigma$ such that $\psi \leq_c \mu$ prefers H' to H.*

If $\mathcal{H}$ is unambiguously fixed, we write $H \ll_c H'$ instead of $H \ll_c^{\mathcal{H}} H'$.

In contrast to $\trianglelefteq_c$, $\ll_c$ is clearly irreflexive; that is, we have not $H \ll_c H$, for any history H. However, like $\trianglelefteq_c$, $\ll_c$ is not transitive and may possess non-trivial cycles. This is witnessed by the following example: consider $\langle \boldsymbol{H}, \leq_c, \leq_t \rangle$, where $\boldsymbol{H} = \{H, H'\}$ and $\leq_c$ consists of the two items $f \leq_c g$ and $g \leq_c f$. Furthermore, the following relations are assumed:

- $H \models f$ and $H \not\models g$,
- $H' \not\models f$ and $H' \models g$.

Then, it holds that

- $f \leq_c g$ prefers H' to H, and
- $g \leq_c f$ prefers H to H'.

It follows that $H \ll_c H'$ and $H' \ll_c H$. Indeed, the former is the case since $f \leq_c g$ prefers H' to H, and for the only pair $g \leq_c f$ which prefers H to H', there is the pair $f \leq_c g$ preferring H' to H. Likewise, $H' \ll_c H$ holds by symmetry. So, there is a non-trivial cycle $H \ll_c H'$ and $H' \ll_c H$, and since $\ll_c$ is irreflexive, $H \ll_c H$ does not hold and hence $\ll_c$ is not transitive. In view of the non-reflexivity and non-transitivity of $\ll_c$, we consider its reflexive and transitive closure, $\ll_c^+$, towards defining σ_c', as follows.

Definition 23 *The strategy σ'_c assigns to each choice preference $\leq_c$ of a prioritised history set $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ the relation $\preceq'_c$, given as follows:*

$$H \preceq'_c H' \quad \text{iff} \quad H \lll^+_c H', \qquad \text{for every } H, H' \in \boldsymbol{H}\ .$$

Let us consider some basic example illustrating this form of choice preference.

Example 10 *Consider a prioritised history set $\langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ over fluents f and g such that $\boldsymbol{H} = \{H_1, H_2, H_3, H_4\}$ and $\leq_c$ consists of the single pair $f \leq_c g$. Assume that the following conditions hold:*

$$\begin{array}{ll} H_1 \models f, H_1 \models g; & H_2 \models f, H_2 \not\models g; \\ H_3 \not\models f, H_3 \models g; & H_4 \not\models f, H_4 \not\models g. \end{array}$$

Then, we have that $\preceq'_c$ consists precisely of the two pairs $H_2 \preceq'_c H_1$ and $H_2 \preceq'_c H_3$. Thus, H_1, H_3, and H_4 are the preferred histories under σ'_c. Note that for σ_c, besides $H_2 \preceq_c H_1$ and $H_2 \preceq_c H_3$, we would also have $H_4 \preceq_c H_2$.

This example illustrates that it is only the highest-ranked formula(s) that determine whether a history is preferred or not. Thus, since H_1 and H_3 both satisfy g, they are equally preferred, despite the fact that they differ on f. If one wishes to have a lexicographic notion of preference where preference is determined by the highest-ranked formula on which they differ, then this can be captured by a suitable modification to Definition 22. Alternately, this can be captured in a quite ad-hoc fashion in our preceding example by adding the preference $g \leq_c f \wedge g$. This would yield the additional preference $H_3 \preceq'_c H_1$.

Last, assume that a choice preference represents some absolute notion of desirability, where $f \leq_c g$ indicates that f is desirable, just not more so than g. We can capture this by adding the preference $\neg f \leq_c f$; adding this preference in Example 10 gives the additional preference on histories that $H_4 \preceq'_c H_2$.

We now turn to the issue of expressing σ'_c in terms of preference frames.

Definition 24 *Let Σ be a propositional action signature and $n \geq 0$.*

1. *For every $\alpha, \beta \in \mathcal{B}_\Sigma$, let $\mathcal{C}[\alpha, \beta]$ and $\mathcal{D}[\alpha, \beta]$ be the following formulas of $\mathcal{P}_{\Sigma,n}$:*

$$\begin{array}{rcl} \mathcal{C}[\alpha, \beta] & = & (\boldsymbol{l} : \Diamond\alpha) \wedge (\boldsymbol{l} : \neg\Diamond\beta) \wedge (\boldsymbol{h} : \Diamond\beta); \\ \mathcal{D}[\alpha, \beta] & = & (\boldsymbol{h} : \Diamond\alpha) \wedge (\boldsymbol{h} : \neg\Diamond\beta) \wedge (\boldsymbol{l} : \Diamond\beta). \end{array}$$

2. *For every prioritised history set $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$, the following formulas of $\mathcal{P}_{\Sigma,n}$ are introduced:*

$$\begin{array}{rcl} \mathcal{E}_\mathcal{H} & = & \bigwedge_{\phi \in dom(\leq_c)} \Big(\bigwedge_{\psi \in \{\gamma \mid \phi \leq_c \gamma\}} \big(\mathcal{D}[\phi, \psi] \supset \bigvee_{\mu \in \{\gamma \mid \psi \leq_c \gamma\}} \mathcal{C}[\psi, \mu]\big)\Big) \text{ and} \\ \Psi_c^{\mathcal{H},\alpha,\beta} & = & \mathcal{C}[\alpha, \beta] \wedge \mathcal{E}_\mathcal{H}, \end{array}$$

for every $\alpha, \beta \in \mathcal{B}_\Sigma$.

The following properties are straightforward to establish:

Lemma 3 *Let $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ be a prioritised history set over a propositional action signature Σ with horizon n.*

1. *For any $H, H' \in \boldsymbol{H}$ and any $\alpha, \beta \in \mathcal{B}_\Sigma$, the following three conditions are equivalent:*
   (a) *$\alpha \leq_c \beta$ prefers H' to H;*
   (b) *$\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \mathcal{C}[\alpha, \beta]$;*
   (c) *$\langle H', H \rangle \models_{\mathcal{P}_{\Sigma,n}} \mathcal{D}[\alpha, \beta]$.*

2. *For any $H, H' \in \boldsymbol{H}$, the following two conditions are equivalent:*
   (a) *for each $\phi, \psi \in \mathcal{B}_\Sigma$, if $\phi \leq_c \psi$ prefers H to H', then there is some $\mu \in \mathcal{B}_\Sigma$ such that $\psi \leq_c \mu$ prefers H' to H;*
   (b) *$\langle H, H' \rangle \models_{\mathcal{P}_{\Sigma,n}} \mathcal{E}_{\mathcal{H}}$.*

Let us now define

$$\boldsymbol{P}^{\mathcal{H}}_{\sigma'_c} = \{\Psi^{\mathcal{H},\alpha,\beta}_c \mid \alpha, \beta \in \mathcal{B}_\Sigma \text{ and } \alpha \leq_c \beta\}, \tag{6}$$

for every prioritised history set $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$. Then, $\ll^{\mathcal{H}}_c$ coincides with $\cup\{\preceq_\phi \mid \phi \in \boldsymbol{P}^{\mathcal{H}}_{\sigma'_c}\}$, which in turn implies that $\preceq'_c$ coincides with the reflexive and transitive closure of $\cup\{\preceq_\phi \mid \phi \in \boldsymbol{P}^{\mathcal{H}}_{c'}\}$. We therefore have the following characterisation:

Theorem 4 *Let $\mathcal{H} = \langle \boldsymbol{H}, \leq_c, \leq_t \rangle$ be a prioritised history set over a propositional action signature Σ.*

Then, for any $H \in \boldsymbol{H}$, H is preferred under $\leq_c$ relative to σ'_c iff H is a maximally preferred history of the preference frame $\langle \boldsymbol{H}, \boldsymbol{P}^{\mathcal{H}}_{\sigma'_c} \rangle$, where $\boldsymbol{P}^{\mathcal{H}}_{\sigma'_c}$ is as in (6).

6 Multi-valued Fluents and Aggregates

So far, we have mostly been dealing with propositional fluents, whose values are restricted to $\{0, 1\}$. In realistic applications, however, we are also faced with non-binary quantities specifying things like temperature, amount of rain, etc. Moreover, actions often have associated costs and other measures, such as duration or distance. These types of quantities permit preference statements expressing certain optimisations or desiderata. In this section, we show how, with the addition of a simple but very powerful construct, we can express aggregates of fluent values. Such aggregates could express, for example, the total amount of rainfall received in a history, or the total cost of a plan. One can then use the preference language $\mathcal{P}_{\Sigma,n}$ as before to express that one prefers histories with least rainfall or, more pertinently, those histories with lowest overall cost or duration. We present an abbreviated account here; more details can be found in [9].

For simplicity, we restrict ourselves to multi-valued fluents taking non-negative integer values from a finite set $\{0, \ldots, n\}$. In analogy to time terms, we introduce *value terms* as follows.

We extent the alphabet of $\mathcal{Q}_{\Sigma,n}$ by

1. a set $\mathcal{U}$ of *value variables*, and

2. the symbols ‘?’ and ‘:’.

Syntactically, a value term is one of:

1. a variable in $\mathcal{U}$,
2. a (multi-valued) fluent in F in an action signature $\langle V, F, A\rangle$,
3. a number in $\{0, \ldots, n\}$,
4. an arithmetic term recursively built from value terms, employing $+$ and $\cdot$ (as well as parentheses) in the usual manner, and
5. a *fluent conditional*, which is an expression of form $(\phi \mathrel{?} a : b)$, where ϕ is a query and a, b are value terms.

Semantically, the value, $val(t)$, of a ground value term t is defined as follows; let $\Sigma = \langle V, F, A\rangle$ be an action signature and $H = (s_0, a_1, s_1, , \ldots, s_n)$ a history over Σ.

1. If t is a fluent of form $f(t')$, then $val(t) = s_j(f)$, where $j = \min(val(t'), n)$.
2. If t is term of form $(\phi \mathrel{?} a : b)$, then $val(t) = val(a)$ if $H \models_{\mathcal{Q}_{\Sigma,n}} \phi$, and $val(t) = val(b)$ otherwise.
3. Otherwise, t is evaluated according to the rules of integer arithmetic.

A major advantage of fluent conditionals is that they allow us to express aggregating expressions over histories—notably without any further extension of the language. We illustrate this by showing how aggregates can be defined in terms of the previous extension to language $\mathcal{Q}_{\Sigma,n}$. Moreover, we show below how this approach can also be used for a direct semantic account of aggregates.

For example, we may define an aggregate $\max[f]$ corresponding to the maximum value of fluent f in a history of length n by means of the following macros representing value terms:

$$\begin{aligned} \max[f] &= \max[f, n]; \\ \max[f, 0] &= f(0); \\ \max[f, i] &= \big((f(i) > \max[f, i-1]) \mathrel{?} f(i) : \max[f, i-1]\big). \end{aligned}$$

An analogous definition can be given for the minimum value of f by taking $\min$ instead of $\max$.

So, for example, in a history of length 2 the macro $\max[f]$ simply stands for the following expression (highlighting the structure by underlining):

$$\underline{(f(2) > \underline{\underline{(f(1) > \underline{f(0)} \mathrel{?} f(1) : \underline{f(0)})}} \mathrel{?} f(2) : \underline{\underline{(f(1) > \underline{f(0)} \mathrel{?} f(1) : \underline{f(0)})}})}. \tag{7}$$

That is, neither $\max[f]$ nor $\max[f, n]$ appear in the actual expression. Thus, given the associated mappings $s_0 : f \mapsto 2$, $s_1 : f \mapsto 3$, and $s_2 : f \mapsto 1$, the expression in (7) evaluates to 3.

It is important to note that aggregates like $\max[f, i]$ are merely macros representing nested value terms. Thus, in particular, they are not fluents.

As a second example, for summing the values of f in a history, we define:

$$\begin{aligned}\mathrm{sum}[f] &= \mathrm{sum}[f,n];\\ \mathrm{sum}[f,0] &= f(0);\\ \mathrm{sum}[f,i] &= (f(i)+\mathrm{sum}[f,i-1]).\end{aligned}$$

A similar definition can be given for counting all occurrences of f being true:

$$\begin{aligned}\mathrm{count}[f] &= \mathrm{count}[f,n];\\ \mathrm{count}[f,0] &= (f(0)\ ?\ 1:0);\\ \mathrm{count}[f,i] &= \mathrm{count}[f,i-1]+(f(i)\ ?\ 1:0).\end{aligned}$$

When expressing preferences among histories H_l and H_h, the language $\mathcal{P}_{\Sigma,n}$ is augmented by labelled value terms of form $\ell : t$, where $\ell \in \{\boldsymbol{l},\boldsymbol{h}\}$ is a label and t is a value term. As done in Section 4, we then understand by a formula (value term) $\ell : \alpha$ the formula (value term) resulting from α by replacing each atom (value term) β in α by its labelled counterpart $\ell : \beta$. The value of a labelled value term $\ell : t$, then, is defined with respect to history H_ℓ. Given this, we can, for example, state a preference for histories with the globally least amount of rain (r) as follows:

$$\boldsymbol{h} : \mathrm{sum}[r] \leq \boldsymbol{l} : \mathrm{sum}[r].$$

Further refinements are easily specified. For example, preferring histories with the globally fewest days (states) on which it rained more than t litres can be modelled by an extension of the count macro:

$$\begin{aligned}\mathrm{count}[r,t] &= \mathrm{count}[r,n,t];\\ \mathrm{count}[r,0,t] &= ((r(0)\leq t)\ ?\ 0:1);\\ \mathrm{count}[r,i,t] &= \mathrm{count}[r,i-1,t]+((r(i)\leq t)\ ?\ 0:1).\end{aligned}$$

For modelling action costs, we associate with each action a fluent yielding the cost of the corresponding action. This is simple, and moreover allows us to associate with an action further measures, such as duration.

Consider the action *drive* and the fluent *driven*, whose value represents the corresponding distance. Summing up the driving costs within the first five states can then be expressed as follows: Let

$$\Phi(i) = (1 \leq i) \wedge (i \leq 5) \wedge \mathit{drive}(i)$$

and

$$\begin{aligned}\mathrm{sum5}[\mathit{drive}] &= \mathrm{sum5}[\mathit{drive},n],\\ \mathrm{sum5}[\mathit{drive},0] &= 0,\\ \mathrm{sum5}[\mathit{drive},i] &= \big(\Phi(i)\ ?\ \mathrm{sum5}[\mathit{drive},i-1]+\mathit{driven}(i+1) : \mathrm{sum5}[\mathit{drive},i-1]\ \big).\end{aligned}$$

Now, let us suppose we prefer histories in which “at least half of the driving is done in the first five hops”. This could be expressed by the preference formula

$$\boldsymbol{l} : \mathrm{sum}[\mathit{drive}] \leq \boldsymbol{h} : (2\cdot \mathrm{sum5}[\mathit{drive}]),$$

where sum is the appropriate global sum value macro.

A similar straightforward semantical description of aggregates can be obtained by appeal to the language extension of $\mathcal{Q}_{\Sigma,n}$ by $(\phi\ ?\ a : b)$. To this end, let us combine the interpretation of terms in the following way.

Definition 25 *Let $H = (s_0, a_1, s_1, \ldots, a_n, s_n)$ be a history over $\Sigma = \langle V, F, A\rangle$ of length n.*

We define $\mathcal{I}$ as follows:

1. *If t is a ground time term, then $\mathcal{I}(t) = \min(val(t), n)$.*
2. *If t is a ground value term, then $\mathcal{I}(t)$ is given as follows:*
   (a) *If $t = f(t')$, for $f \in F$, then $\mathcal{I}(t) = s_i(f)$, where $i = \min(val(t'), n)$.*
   (b) *If $t = (\phi\ ?\ e_1 : e_2)$, then $\mathcal{I}(t) = \mathcal{I}(e_1)$ if $H \models_{\mathcal{Q}_{\Sigma,n}} \phi$, and $\mathcal{I}(t) = \mathcal{I}(e_2)$ otherwise.*
   (c) *Otherwise, $\mathcal{I}(t) = val(t)$.*

The semantics of aggregates can now be defined by specific extensions of the definition of $\mathcal{I}$.

To begin with, consider where we wish to define a fluent that will correspond to the maximum value of some other fluent obtained so far in a history. Specifically, for fluent f we want to define a fluent $maxf$, where

$$maxf(i) = \max_{0\le j\le i} f(j).$$

We do this by extending Definition 25 as follows:

$$\mathcal{I}(maxf(i)) = \begin{cases} \mathcal{I}(f(0)), & \text{if } i = 0;\\ \mathcal{I}(f(i) > maxf(i-1)\ ?\ f(i) : maxf(i-1)), & \text{otherwise.}\end{cases}$$

Similarly we can define a fluent $sumf$ that will correspond to the sum of the values of fluent f obtained so far in a history; i.e., we wish to define

$$sumf(i) = \sum_{0\le j\le i} f(j).$$

We do this by extending Definition 25 as follows:

$$\mathcal{I}(sumf(i)) = \begin{cases} \mathcal{I}(f(0)), & \text{if } i = 0;\\ \mathcal{I}(f(i) + sumf(i-1)), & \text{otherwise.}\end{cases}$$

In both cases we recursively define an aggregate function in terms of a fluent along with that (aggregate) function at an earlier time point. In each case, the value of a defined fluent (such as $maxf$) can be determined from the underlying fluent (viz., f) by a process akin to macro expansion.

The overall scheme can be obviously extended to more than one fluent, and more than a single time point for each step. For example, we can define a fluent ex that counts the number of times the value of fluent f exceeds that of g two time points ago, as follows:

$$\mathcal{I}(ex(i)) = \begin{cases} 0, & \text{if } i = 0 \text{ or } i = 1;\\ \mathcal{I}(f(i) > g(i-2)\ ?\ ex(1-i) + 1 : ex(1-i)), & \text{otherwise.}\end{cases}$$

Practically then, it is a simple matter to determine, for example,

- the number of days (states) with more than 1 cm of rain,
- the number of hikes that one took of distance over 20 km, or
- the maximum rainfall that occurred on a hike of over 20 km.

From this, it is an easy matter to extend the language so that we can assert that we prefer seasons (histories) in which the maximum rainfall that occurred on a hike of over 20 km is as low as possible, or, more prosaically, that we prefer a history in which total action costs are as low as possible.

7 Expressing Other Approaches

We conclude our technical exposition by showing how other approaches to preference handling and temporal reasoning can be captured within our framework. We start our treatment with a well-known preference method due to Son and Pontelli [33]. Afterwards, we deal with PDDL3 [20], the extension of the *Planning Domain Definition Language* that includes constraints and preferences. Lastly, we briefly consider Allen's interval algebra [1].

7.1 Son and Pontelli's Language $\mathcal{PP}$

Son and Pontelli [33] present a language $\mathcal{PP}$ for specifying preferences between possible solutions of planning problems. We show below how this language can be encoded in $\mathcal{Q}_{\Sigma,n}$ and $\mathcal{P}_{\Sigma,n}$. Interestingly, $\mathcal{P}_{\Sigma,n}$ comes mainly into play for formalising meta-language definitions of $\mathcal{PP}$. For brevity, we proceed by directly giving the encoding of $\mathcal{PP}$ rather than first giving its semantics. To this end, we provide a translation $\tau(\cdot)$ mapping elements of $\mathcal{PP}$ into $\mathcal{Q}_{\Sigma,n}$ and $\mathcal{P}_{\Sigma,n}$, respectively, such that the corresponding reasoning task is equivalent. We begin with the most basic language constructs.

For histories of length n over a propositional action signature $\Sigma = \langle V, F, A \rangle$, we map so-called *basic desire formulas* in $\mathcal{PP}$ onto formulas in $\mathcal{Q}_{\Sigma,n}$ in the following way:

1. $\tau(\textbf{goal}(\psi)) = \psi(n)$;
2. $\tau(f) = \psi(0)$, for $f \in F$;
3. $\tau(\textbf{occ}(a)) = a(1)$, for $a \in A$;
4. $\tau(f(t)) = f(t)$, for a fluent atom $f(t) \in \mathcal{Q}_{\Sigma,n}$;
5. $\tau(a(t)) = a(t)$, for an action atom $a(t) \in \mathcal{Q}_{\Sigma,n}$;
6. $\tau(\psi_1 \land \psi_2) = \tau(\psi_1) \land \tau(\psi_2)$;
7. $\tau(\psi_1 \lor \psi_2) = \tau(\psi_1) \lor \tau(\psi_2)$;
8. $\tau(\neg\psi) = \neg\tau(\psi)$;

9. $\tau(\textbf{next}(\psi)) = \tau(\psi)[i_1/i_1{+}1, \ldots, i_j/i_j{+}1, m_1/m_1{+}1, \ldots, m_k/m_k{+}1]$, where $\{i_1, \ldots, i_j\}$ is the set of all time-stamp variables in $\tau(\psi)$ and $\{m_1, \ldots, m_k\}$ is the set of all natural numbers in $\tau(\psi)$;

10. $\tau(\textbf{always}(\psi)) = \forall i \; \tau(\psi[i])$, where i is a newly-introduced variable.

11. $\tau(\textbf{eventually}(\psi)) = \exists i \; \tau(\psi[i])$, where i is a newly-introduced variable.

12. $\tau(\textbf{until}(\psi_1, \psi_2)) = \exists i \; \tau(\psi_2[i]) \wedge \forall j \; (j < i \supset \tau(\psi_1[j]))$, where i, j are newly-introduced variables.

Here, similar to the notation introduced in Section 5.1, $\psi[i]$ stands for the formula in $\mathcal{Q}_{\Sigma,n}$ obtained from $\psi \in \mathcal{PP}$ by replacing in ψ each fluent $f \in F$ by fluent atom $f(i)$ and by replacing in ψ each expression $\textbf{occ}(a) \in \mathcal{PP}$ for $a \in A$ by action atom $a(i)$. The language $\mathcal{PP}$ itself is given by the constructs mapped in 1-3 and 6-12 above. Items 4 and 5 are necessary to map time-stamp variables introduced by $\psi[\cdot]$ in 9-12. The proviso in Items 10-12 is due to the fact that modalities may be iterated; the proviso prevents variable clashes.

As an example, consider the basic desire formula

$$\mathit{rain} \wedge \textbf{occ}(\mathit{take\_umbrella}) \wedge \textbf{always}(\textbf{occ}(\mathit{take\_umbrella}) \supset \textbf{next}(\neg \mathit{rain})).$$

This is turned by $\tau(\cdot)$ into the following formula:

$$\mathit{rain}(0) \wedge \mathit{take\_umbrella}(1) \wedge \forall i(\mathit{take\_umbrella}(i) \supset \neg \mathit{rain}(i+1)).$$

The following characterisation can be shown via an inductive argument, which is omitted here.

Theorem 5 *Let $\models_{\mathcal{PP}}$ be the satisfaction relation of $\mathcal{PP}$ between a history and a basic desire formula.*

Then, for any history H of length n and any basic desire formula φ,

$$H \models_{\mathcal{PP}} \varphi \text{ iff } H \models_{\mathcal{Q}_{\Sigma,n}} \tau(\varphi).$$

Given a basic desire formula $\varphi \in \mathcal{PP}$, Son and Pontelli [33] define a preference relation among histories H_h and H_l as follows:

$$H_l \prec_\varphi H_h \text{ iff } H_h \text{ satisfies } \varphi \text{ but } H_l \text{ does not satisfy } \varphi.$$

This meta-level definition can now be easily encoded in $\mathcal{P}_{\Sigma,n}$ in terms of the formula

$$\Theta_{\prec_\varphi} = \boldsymbol{h} : \tau(\varphi) \wedge \boldsymbol{l} : \neg\tau(\varphi),$$

satisfying

$$H_l \prec_\varphi H_h \quad \text{iff} \quad \langle H_l, H_h \rangle \models_{\mathcal{P}_{\Sigma,n}} \Theta_{\prec_\varphi}.$$

Further, the condition that two histories H_h and H_l are indistinguishable with respect to φ, written $H_l \approx_\varphi H_h$, is expressed in $\mathcal{P}_{\Sigma,n}$ by the formula

$$\Theta_{\approx_\varphi} = (\boldsymbol{h} : \tau(\varphi) \wedge \boldsymbol{l} : \tau(\varphi)) \vee (\boldsymbol{h} : \neg\tau(\varphi) \wedge \boldsymbol{l} : \neg\tau(\varphi)),$$

likewise obeying

$$H_l \approx_\varphi H_h \quad \text{iff} \quad \langle H_l, H_h \rangle \models_{\mathcal{P}_{\Sigma,n}} \Theta_{\approx_\varphi}.$$

Next, so-called *atomic preferences*, $\varphi_1 \lhd \varphi_2 \lhd \cdots \lhd \varphi_m$, were added to $\mathcal{PP}$ for lexically ordering different basic desires φ_i ($1 \le i \le m$). Although they were again characterised at the meta level, we can encode them thus:

$$\begin{aligned} \Theta_{\approx_{\varphi_1 \lhd \varphi_2 \lhd \cdots \lhd \varphi_m}} &= \bigwedge_{i \in \{0,\ldots,m\}} \Theta_{\approx_{\varphi_i}}; \\ \Theta_{\prec_{\varphi_1 \lhd \varphi_2 \lhd \cdots \lhd \varphi_m}} &= \bigvee_{i \in \{0,\ldots,m\}} \bigwedge_{j \in \{0,\ldots,i-1\}} \Theta_{\approx_{\varphi_1 \lhd \varphi_2 \lhd \cdots \lhd \varphi_j}} \wedge \Theta_{\prec_{\varphi_i}}. \end{aligned}$$

In $\mathcal{PP}$, atomic preferences are then further extended to so-called *general preferences*, which are mixed Boolean and lexicographic combinations of atomic and general preferences. This type of preferences can be encoded in a similar way.

7.2 PDDL3

PDDL [21], the Planning Domain Definition Language, is intended as a common problem specification language for (among other things) planning competitions. PDDL3 [20] extends PDDL by allowing the expression of state trajectory and preference constraints. Such constraints may be *strong* or *soft*. In our approach, we expect strong constraints to be satisfied by all histories under consideration, while soft constraints are handled via our preference framework.

The constraint specification part of PDDL3 can be expressed in our basic language $\mathcal{Q}_{\Sigma,n}$ once the objects of discourse are "harmonized". In fact, the semantics of PDDL3 relies on sequences of states, which may occur in non-uniform time steps. More formally, it uses sequences of the form

$$((s_0, 0), (s_1, t_1), (s_2, t_2), \ldots, (s_{n-1}, t_{n-1}), (s_n, t_n)), \tag{8}$$

where each s_i is a state and t_i is a point in time, for $0 \le i \le n$ and $0 \le t_1 \le \cdots \le t_n$. Such a sequence corresponds to a history of length n that proceeds in non-uniform time steps t_1, $t_2 - t_1$,...,$t_n - t_{n-1}$. Given that in the context of PDDL plans are to be constructed, no actions are taken into account.

For comparing our approach with PDDL3 on an equal basis, we associate with a sequence as in (8) a history $(s'_0, a_1, s'_1, a_2, s'_2, \ldots, s'_{n-1}, a_n, s'_n)$, where

1. $s'_i = s_i \cup \{t = t_i\}$, for $0 \le i \le n$, and
2. $a_i = \emptyset$, for $1 \le i \le n$.

Here, t is a new fluent indicating the time point associated with state s_i. Following this recipe, the PDDL3 sequence $((s_0, 0), (s_1, 3), (s_2, 5))$, for example, is associated with the history $(s_0 \cup \{t = 0\}, \emptyset, s_1 \cup \{t = 3\}, \emptyset, s_2 \cup \{t = 5\})$.[8]

Given this, we are able to provide a translation, $\rho(\cdot)$, whose image is defined on histories, as given in Definition 2: For histories of length n over an action signature Σ, some (duration)

[8]If time points were integer-valued, then a more natural mapping is to introduce *null* actions for time points 1, 2, and 4. However, time points may be real-valued, so we adopt the convention of having a new "time point" fluent t.

d such $0 \leq d \leq n$, and atomic formulas ψ, ψ_1, ψ_2,[9] we map an expression in PDDL3 (cf. [20, Definition 3]) onto a formula in $\mathcal{Q}_{\Sigma,n}$ in the following way:

1. $\rho((\texttt{at end}\,\psi)) = \psi(n)$;
2. $\rho(\psi) = \psi(n)$;
3. $\rho((\texttt{always}\,\psi)) = \forall i\; \psi(i)$;
4. $\rho((\texttt{sometime}\,\psi)) = \exists i\; \psi(i)$;
5. $\rho((\texttt{within}\,d\,\psi)) = \exists i\; (\psi(i) \wedge (t(i) \leq d))$;
6. $\rho((\texttt{at-most-once}\,\psi)) =$
$$\forall i\; \Big(\psi(i) \supset \exists j\; (j \geq i) \wedge \big([\forall k\; (i \leq k) \wedge (k \leq j) \supset \psi(k)] \wedge [\forall l\; (l > j) \supset \neg\psi(l)]\big)\Big);$$
7. $\rho((\texttt{sometime-after}\,\psi_1\;\psi_2)) = \forall i\; \Big(\psi_1(i) \supset \big(\exists j\; (i \geq j) \wedge (j \leq n) \wedge \psi_2(j)\big)\Big)$;
8. $\rho((\texttt{sometime-before}\,\psi_1\;\psi_2)) = \forall i\; \Big(\psi_1(i) \supset \big(\exists j\; (0 \geq j) \wedge (j \leq i) \wedge \psi_2(j)\big)\Big)$;
9. $\rho((\texttt{always-within}\,d\,\psi_1\;\psi_2)) =$
$$\forall i\; \Big(\psi_1(i) \supset \big(\exists j\; (i \geq j) \wedge (j \leq n) \wedge \psi_2(j) \wedge (t(j) - t(i) \leq d)\big)\Big);$$
10. $\rho((\texttt{and}\,\psi_1\;\psi_2)) = \rho(\psi_1) \wedge \rho(\psi_2)$.

In addition, the language of PDDL has a predicate language, encompassing logical connectives, like `implies`, as well as object variables along with quantifiers, like `forall`.

The above translation is of great importance from an applications point of view, since it demonstrates that our basic language is powerful enough to express practical examples. For illustration, let us just consider two examples from the long list given in [20].

Example 11 *The statement*

"a fragile block should never have something above it"

is expressed in PDDL3 thus:

```
(always  (forall (?b - block)
         (implies (fragile ?b)
                  (clear ?b)))).
```

The translation of this expression into $\mathcal{Q}_{\Sigma,n}$ *is as follows:*

$$\forall i\; \forall b\; \big(\mathit{block}(b) \wedge \mathit{fragile}(b) \supset \mathit{clear}(b, i)\big).$$

This assumes that there are only finitely many blocks as well as that clear *is the only fluent.*

[9] Note that this precludes nesting.

Example 12 *The statement*

> *"whenever the energy of a rover is below 5, it should be at the recharging location within 10 time units"*

is formalised in PDDL3 in the following way:

```
(forall (?r - rover)
        (always-within 10
                       (< (energy ?r) 5)
                       (at ?r recharging-point))).
```

Translating this expression into $\mathcal{Q}_{\Sigma,n}$ yields

$$\forall r\ rover(r) \supset \\ \forall i\ \Big([energy(r,i) < 5] \supset \big(\exists j\ (i \geq j) \wedge (j \leq n) \\ \wedge\ at(r, recharging\text{-}point, j) \wedge (t(j) - t(i) \leq 10)\big)\Big).$$

Similar to the preceding example, this assumes a finite number of rovers as well as that $energy$ *and* at *are fluents, while* $rover$ *is not.*

Gerevini and Long [20, Section 4] provide a substantial list of other real-world preferences, all of which would be expressible in our framework, analogously to the above.

7.3 Allen's Interval Algebra

A major approach in temporal reasoning is the well-known *interval algebra* [1], in which time intervals are the primitive objects. There are 13 basic relations between intervals, including relations such as *before, meets, overlaps*, etc. It may be that a temporal preference language based on the interval algebra would provide a useful high-level language for expressing general preferences. Thus for example one could assert that the interval *during* which coffee is drunk preferentially *overlap*s with or starts *before* the interval in which a seminar takes place. To this end, one might define that a fluent f constitutes an interval just if it is true only for a contiguous set of time points:

$$interval(f) = \exists i, j\big((i \leq j) \wedge \forall k(f(k) \equiv (i \leq k) \wedge (k \leq j))\big).$$

Then, the relation that an interval *meets* another can be defined by:

$$meets(f, g) = interval(f) \wedge interval(g) \wedge \exists i(f(i) \wedge \neg f(i+1) \wedge \neg g(i) \wedge g(i+1)).$$

Other relations follow analogously. The preference language $\mathcal{P}_{\Sigma,n}$ can then be used to encode preferences among interval expressions. Hence, one could in this fashion define a high-level preference language based on the interval algebra which could in turn be translated into our framework.

8 Conclusion

We have addressed the problem of expressing general preferences over histories, inter alia addressing preferences in planning systems. Preferences are "soft" constraints, or desirable (but not required) outcomes that one would like to achieve. We first defined a query language $\mathcal{Q}_{\Sigma,n}$ for specifying arbitrary conditions that may be satisfied by a history. Given this, we defined a second language $\mathcal{P}_{\Sigma,n}$ for specifying preferences. The framework allows the expression of conditional preferences, or preferences holding in a given context, as well as absolute preferences, expressing a general desirability that a formula hold in a history. An interesting extension to our approach would be to consider the question of weights on preferences, thereby addressing the case in which some preferences are considered more important than others.

A preference induces a binary relation on histories, so that in an ordered pair of histories the second history is preferred to the first. From this, one can define a global ordering on the set of histories, the maximal elements of which are the preferred histories. The overall approach is very general and flexible; specifically we argue that previous approaches to preferences in planning are expressible in our formalism. Thus, too, our approach constitutes a "base" language, in terms of which higher-level operators may be defined.

We gave fundamental classifications of domain-specific preference types, constituting action vs. fluent preferences and choice vs. temporal preferences. We encoded choice and temporal preferences in our framework, and showed how preferred histories can be determined for each preference type. Further, we showed how aggregates can be defined and used in preferences, hence for example allowing a preference for histories with lowest action costs.

In a planning context, our approach would amount to generating plans and selecting the most preferred plan based on the preferences. As such, the approach is readily adaptable to an anytime algorithm, in which one may select the currently-best plan(s), but with the hope of a more-preferred plan being generated. An obvious topic for future work is to directly generate a preferred plan (rather than selecting from candidate plans); however, this appears to be a difficult problem in general. Last, while the approach is formulated within the framework of action languages, our results are applicable to general planning formalisms.

Acknowledgements

The first author was partially supported by a Canadian NSERC Discovery Grant. The second author was partially supported by the German Science Foundation (DFG) under grant SCHA 550/6, TP C. The collaboration of the second and third author was partially supported by the Information Society Technologies programme of the European Commission, Future and Emerging Technologies under the IST-2001-37004 WASP project.

We would like to express our thanks to Jérôme Lang for many helpful suggestions and fruitful discussions on earlier drafts of this paper, and to Yannis Dimopoulos for getting us interested in preferences on aggregates. We also thank the referees for their interesting, pertinent, and insightful comments.

References

[1] J. Allen. Maintaining knowledge about temporal intervals. *Communications of the ACM*, 26(1):832–843, 1983.

[2] F. Bacchus and F. Kabanza. Using temporal logics to express search control knowledge for planning. *Artificial Intelligence*, 116(1-2):123–191, 2000.

[3] M. Bienvenu, C. Fritz, and S. McIlraith. Planning with qualitative temporal preferences. In P. Doherty, J. Mylopoulos, and C. A. Welty, editors, *Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning (KR 2006)*, pages 134–144. AAAI Press, 2006.

[4] C. Boutilier, R. Brafman, C. Domshlak, H. H. Hoos, and D. Poole. CP-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.

[5] R. I. Brafman and Y. Chernyavsky. Planning with goal preferences and constraints. In S. Biundo, K. L. Myers, and K. Rajan, editors, *Proceedings of the Fifteenth International Conference on Automated Planning and Scheduling (ICAPS 2005)*, pages 182–191. AAAI Press, 2005.

[6] G. Brewka. A rank based description language for qualitative preferences. In *Proceedings of the Sixteenth European Conference on Artificial Intelligence (ECAI 2004)*, pages 303–307. IOS Press, 2004.

[7] J. Delgrande, T. Schaub, and H. Tompits. Domain-specific preferences for causal reasoning and planning. In D. Dubois, C. Welty, and M. Williams, editors, *Proceedings of the Ninth International Conference on the Principles of Knowledge Representation and Reasoning (KR 2004)*, pages 673–682. AAAI Press, 2004.

[8] J. Delgrande, T. Schaub, and H. Tompits. A general framework for expressing preferences in causal reasoning and planning. In S. McIlraith, P. Peppas, and M. Thielscher, editors, *Proceedings of the Seventh International Symposium on Logical Formalizations of Commonsense Reasoning*, pages 46–54, 2005.

[9] J. Delgrande, T. Schaub, and H. Tompits. An extended query language for action languages (and its application to aggregates and preferences). In A. Hunter and A. Kakas, editors, *Proceedings of the Eleventh International Workshop on Non-Monotonic Reasoning (NMR 2006)*, 2006.

[10] P. Doherty and J. Kvarnstrom. TALplanner: An empirical investigation of a temporal logic-based forward chaining planner. In *TIME*, pages 47–54, 1999.

[11] T. Eiter, W. Faber, N. Leone, G. Pfeifer, and A. Polleres. Answer set planning under action costs. *Journal of Artificial Intelligence Research*, 19:25–71, 2003.

[12] R. Feldmann, G. Brewka, and S. Wenzel. Planning with prioritized goals. In P. Doherty, J. Mylopoulos, and C. A. Welty, editors, *Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning (KR 2006)*, pages 503–514. AAAI Press, 2006.

[13] J. Figueira, S. Greco, and M. Ehrgott, editors. *Multiple Criteria Decision Analysis: State of the Art Surveys*. Springer Verlag, 2005.

[14] M. Fitting. *First-Order Logic and Automated Theorem Proving*. Springer Verlag, New York, 1990.

[15] C. Fritz and S. McIlraith. Decision-theoretic golog with qualitative preferences. In P. Doherty, J. Mylopoulos, and C. A. Welty, editors, *Proceedings of the Tenth International Conference on Principles of Knowledge Representation and Reasoning (KR 2006)*, pages 153–163. AAAI Press, 2006.

[16] D. M. Gabbay. *Labelled Deductive Systems: Volume 1*, volume 33 of *Oxford Logic Guides*. Oxford University Press, 1996.

[17] H. Geffner and J. Pearl. Conditional entailment: Bridging two approaches to default reasoning. *Artificial Intelligence*, 53(2-3):209–244, 1992.

[18] M. Gelfond and V. Lifschitz. Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9:365–385, 1990.

[19] M. Gelfond and V. Lifschitz. Action languages. *Electronic Transactions on AI*, 3, 1998. Available at `http://www.ep.liu.se/ej/etai/`.

[20] A. Gerevini and D. Long. Preferences and soft constraints in PDDL3. In *Proceedings of the ICAPS 2006 Workshop on Preferences and Soft Constraints in Planning*, pages 46–53, 2006.

[21] M. Ghallab, A. Howe, C. Knoblock, D. McDermott, A. Ram, M. Veloso, D. Weld, and D. Wilkins. PDDL – the planning domain definition language. Technical Report CVC TR-98-003/DCS TR-1165, Yale Center for Computational Vision and Control, 1998.

[22] E. Giunchiglia and V. Lifschitz. An action language based on causal explanation: Preliminary report. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI '98)*, pages 623–630. AAAI Press/The MIT Press, 1998.

[23] P. Haddawy and S. Hanks. Representations for decision-theoretic planning: Utility functions for deadline goals. In B. Nebel, C. Rich, and W. R. Swartout, editors, *Proceedings of the Third International Conference on the Principles of Knowledge Representation and Reasoning (KR '92)*, pages 71–82. Morgan Kaufmann Publishers, 1992.

[24] J. Y. Halpern. Defining relative likelihood in partially-ordered preferential structures. *Journal of Artificial Intelligence Research*, 7:1–24, 1997.

[25] U. Junker, J. Delgrande, J. Doyle, F. Rossi, and T. Schaub. Special issue on preferences in CP and AI. *Computational Intelligence*, 2(20), 2004.

[26] H. Levesque, F. Pirri, and R. Reiter. Foundations for the situation calculus. *Linköping Electronic Articles in Computer and Information Science*, 3(18), 1998. Available at `http://www.ep.liu.se/ea/cis/`.

[27] R. Luca and H. Raiffa. *Games and Decisions*. John Wiley & sons, New York, 1957.

[28] K. Myers and T. Lee. Generating qualitatively different plans through metatheoretic biases. In *Proceedings of the Sixteenth National Conference on Artificial Intelligence (AAAI '99)*, pages 570–576. AAAI Press/The MIT Press, July 1999.

[29] D. Nau, Y. Cao, A. Lotem, and H. Muoz-Avila. SHOP: Simple hierarchical ordered planner. In T. Dean, editor, *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI '99)*, pages 968–973. Morgan Kaufmann Publishers, 1999.

[30] M. Oztürk, A. Tsoukiàs, and P. Vincke. Preference modelling. In Figueira et al. [13], pages 27–72.

[31] M. Resnik. *Choices: An Introduction to Decision Theory*. University of Minnesota Press, 1987.

[32] T. Son, C. Baral, T. Nam, and S. McIlraith. Domain-dependent knowledge in answer set planning. *ACM Transactions on Computational Logic*, 2006. To appear.

[33] T. Son and E. Pontelli. Planning with preferences in logic programming. In V. Lifschitz and I. Niemelä, editors, *Proceedings of the Seventh International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR '04)*, volume 2923 of *Lecture Notes in Artificial Intelligence*, pages 247–260. Springer Verlag, 2004.

[34] E. Triantaphyllou. *Multi-Criteria Decision Making Methods: A Comparative Study*. Springer Verlag, 2004.

[35] M. Wellman and J. Doyle. Preferential semantics for goals. In *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI '91)*, pages 698–703. AAAI Press/The MIT Press, 1991.