

Planning with graded nondeterministic actions: a possibilistic approach

Célia Da Costa Pereira*, Frédérick Garcia*
Jérôme Lang**, Roger Martin-Clouaire*

*INRA/SBIA, Auzeville, BP 27
31326 Castanet Tolosan cedex FRANCE

** IRIT, Université Paul Sabatier, 118 Route de Narbonne
31062 Toulouse cedex FRANCE

Abstract

This paper proposes a framework for planning under uncertainty given a partially known initial state and a set of actions having nondeterministic (disjunctive) effects, some being more possible (normal) than the others. The problem, henceforth called possibilistic planning problem, is represented in an extension of the STRIPS formalism in which the initial state of the world and the graded nondeterministic effects of actions are described by possibility distributions. Two notions of solution plans are introduced: γ -*acceptable plans* that lead to a goal state with a certainty greater than a given threshold γ , and *optimally safe plans* that lead to a goal state with maximal certainty. It is shown that the search of a γ -acceptable plan amounts to solve a derived planning problem that has only pure (non graded) nondeterministic actions. A sound and complete partial order planning algorithm, called NDP, has been developed for such classical nondeterministic planning problems. The generation of γ -acceptable and optimally safe plans is achieved by two sound and complete planning algorithms: POSPLAN that relies on NDP, and POSPLAN* that can be seen as a hierarchical version of POSPLAN. The possibilistic planning framework is illustrated throughout the paper by an example in the agronomic domain.

1 Introduction

In a “classical” planning problem, it is assumed that actions are deterministic, the initial state is known and the goal is defined by a set of final states; a solution plan is then a sequence of actions that leads from the initial state to a goal state. However, most practical problems do not satisfy these conditions of complete and deterministic information. In order to relax them, some authors have proposed approaches of planning under uncertainty in which the effects of actions are described by probabilities over possible resulting states. More particularly, Kushmerick *et al.* [1] have developed an extension of the STRIPS [2] formalism that enables to represent actions having conditional and probabilistic effects and to cope with partial observability of the initial state. They have constructed a planning algorithm that

generates a partially ordered plan (if any exists) such that the probability to reach a goal state is greater than a user-supplied threshold. The implementation of this algorithm based on nonlinear planning primitives drawn from the SNLP [3] approach gave rise to the BURIDAN [1] system.

This article considers the class of planning problems in which, firstly, the environment is *static* (which means that all changes that take place result from actions specified in the plan given by the agent), and, secondly, the environment is assumed to be *unobservable* during plan execution, thus requiring the search of non-conditional plans that must be robust to uncertainty. It proposes a possibilistic counterpart of the Kushmerick et al.[1]’s approach in which possibility distributions are used to represent the uncertainty both on the initial and subsequent states and on the outcomes of the execution of an action in a particular context. Two notions of solution plans for such a possibilistic planning problem are introduced: γ -*acceptable plans* that lead to a goal state with a certainty greater than a given threshold γ , and *optimally safe plans* that lead to a goal state with maximal certainty.

The main benefit expected from the possibility theory [4] framework concerns the ability to represent more qualitatively and, thus more faithfully, what is known about the initial state and the possible effects of actions; the possibilistic approach is likely to be less sensitive to a lack of precision in the assessment of uncertainty. Using a model in which actions have possibilistic effects is particularly well-suited for cases in which the probabilities of the resulting effects of actions are not available, not very reliable, or hard to obtain, that is, in situation of partial or total ignorance about the immediate consequence of applying an action. Moreover, the notion of action with possibilistic effects properly generalizes the notion of nondeterministic actions by enabling the representation of ordinal grading in the uncertainty that characterize the uncontrollable choice process through which the real effect of an action will be determined. What is represented is simply that one or several effects are normal in essence (nothing prevents them from occurring) and that some are more normal (less exceptional) than others, that is, some may be considered more plausible than others in the absence of any further information.

One of the most important results of this paper states that the search of a γ -acceptable plan amounts to solve a planning problem strait-forwardly derived from the original possibilistic one and constituted only of pure (non graded) nondeterministic actions. A partial order planning algorithm, called NDP, has been developed for such classical nondeterministic planning problems. The generation of γ -acceptable and optimally safe plans is accomplished by the algorithms POSPLAN that relies on NDP, and POSPLAN* that can be seen as a hierarchical version of POSPLAN applied iteratively to increasing levels of acceptability until the highest one is reached. Thus our approach has interesting computational properties and naturally embeds a form of anytime planning method.

Section 2 gives some recalls on possibility theory and further develops arguments in favor of using possibility theory in planning. Section 3 describes a representation

framework for possibilistic planning problems and formally defines the two notions of solution plans. Section 4 presents the partial order planning algorithms NDP, POSPLAN and POSPLAN* and proves their soundness and completeness. Lastly, we conclude by summarizing our contribution and by pointing out some limitations and future works concerning several possible extensions of our framework. The main concepts are illustrated throughout the paper by a simple example in the agronomic domain.

2 Possibility theory and planning: motivations

2.1 Issues in flexible planning

The classical planning paradigm usually refers to “searching for an unconditional sequence of actions leading from the initial state to a goal state”. There are several directions in which these strict assumptions can be relaxed so as to enable more flexibility in the representation of a planning problem. Some of the works done in this purpose make a significant use of decision theory, and are thus referred as as “decision-theoretic planning”. We give here a (non-exhaustive) list of issues in flexible planning and locate our work with respect to these issues.

- Enabling the representation of states, action effects and goals in a *concise, structured language*, thus avoiding a heavy, explicit representation of all states. The first seminal work in this direction was STRIPS [2] where states and goals are represented by conjunctions of literals and where action effects only mention the atoms that are modified by the action. This kind of representation was extended to deductive models [5, 6] [7] [8], or revision models [9] in order to deal with context-dependent effects, i.e., effects that depend on the state of the world at execution time. Some other approaches make use of logical languages for designing planning problems ([10] [11] [12]).
- Enabling the representation of *uncertainty about the initial state and/or the possible effects of actions*. This issue has lead to several approaches to *probabilistic planning* ([13], [14], [1], [15] ...), of which one of the most important is Kushmerick et al.’s BURIDAN planner [1] where the effects of actions as well as the initial state description are represented by means of probabilistic state operators that are a probabilistic extension of STRIPS’ operators. A solution plan is there an unconditional sequence of actions leading to a goal state with a probability not less than a given threshold. See also Section 3.
- Enabling the *flexible representation of goals*, replacing the set of goal states of classical planning by a preference ordering on goal states ([16] [17]) or a utility function ([18] ...) thus embedding the representation of the planning problem into decision theory.
- Taking account of different assumptions concerning *observability*. In a *fully observable* decision process, the current state is always known before the agent

has to act so that an adequate solution consists in a conditional plan mapping each possible state at each time point to an actual decision. Such an assumption leads naturally to conditional planning, and in a decision-theoretic perspective to fully observable Markov decision processes (FOMDP) [19] [20]. A qualitative, possibilistic variant of FOMDP was proposed in [17]. The other extreme case is *non-observability*, where the agent never gets any feedback from the process, which entails that one looks for an unconditional plan. Between these two extreme cases, *partially observable* decision processes enable the agent to gather some further information about the current state by performing tests (see for instance [21] [22]).

- Enabling the agent to interrupt the planner at any time after it has been launched (“anytime planning”), assuming thus that at any step of the planning process, the planner maintains a solution whose quality increases with execution time, eventually leading to an optimal solution.

In this article we propose an approach to planning under uncertainty which has the following main features:

- actions are expressed in an extended version of the STRIPS representation that supports context-dependent as well as graded nondeterministic effects;
- uncertainty about the initial state and effects of actions is modeled by *possibility distributions* and any solution is qualified by a necessity (or certainty) degree;
- goals are assumed to be *non-flexible*, i.e., a set of goal states is specified;
- the decision process is *non-observable*;
- the framework lends itself to the implementation of an anytime planning method (see Section 4) that yields an optimally safe solution plan if run until natural quiescence or a less safe plan if interrupted sooner (provided one exists and enough time has been allocated to its generation).

2.2 Possibility theory

The basic representation tool of possibility theory [23] is the *possibility distribution*:

Definition 1 *Let S a finite set of states. A (normalized) possibility distribution on U is a mapping from S to $[0, 1]$ such that $\text{Sup}_{u \in U} \pi(u) = 1$ ¹.*

In the rest of the paper, possibility distributions will be used both for :

¹The normalization condition $\text{Sup}_{u \in U} \pi(u) = 1$ is not always required in possibility theory. The reason why we require it in this paper will be clear when defining possibilistic actions (see Section 3.2).

- representing uncertainty about the initial state: $\pi_{init}(s)$ represents the possibility that the initial state may actually be s . In particular, $\pi_{init}(s) = 0$ means that it is fully impossible that s be the actual initial state, while $\pi_{init}(s) = 1$ means that this is completely possible.
- representing uncertainty about the effects of actions. For this purpose, to each action a and each state s we associate a possibility distribution $\pi[.|s, a]$. For each state s' , $\pi[s'|s, a]$ represents the possibility that executing action a in state s actually leads to state s' .

A possibility distribution π induces a possibility measure Π and a necessity measure N , which are both mappings from 2^S to $[0, 1]$, defined respectively by: $\forall A \subseteq S$,

$$\Pi(A) = \text{Sup}_{s \in A} \pi(s)$$

and

$$N(A) = 1 - \Pi(\bar{A}) = \text{Inf}_{s \in \bar{A}} 1 - \pi(s)$$

$\Pi(A)$ measures to what extent it is consistent with the knowledge expressed by π to assume that the actual state is in A . $N(A)$ measures to what extent the knowledge expressed by π implies that the actual state is in A .

Lastly, let us mention that possibility theory can be viewed as a graded generalization of nondeterminism (for more discussion see Section 2.4). Nondeterminism, or qualitative uncertainty, is the restriction of possibility theory obtained by enforcing $\pi(s) \in \{0, 1\}, \forall s$; for instance, a nondeterministic action a whose set of possible resulting states when performed in state s is $Res(s, a) \subseteq S$ is encoded by $\pi[s'|s, a] = 1$ if $s' \in Res(s, a)$ and $\pi[s'|s, a] = 0$ otherwise. Thus, requiring that all possibility values are equal to 0 or 1 comes down exactly to the definition of classical subsets of S . Note that, contrarily to possibility theory, *probability is not a generalization of nondeterminism*: indeed, there is no way of recovering classical subsets from probability distributions². This is due to the fact that possibility theory is well-suited to representing states of partial or total ignorance, and probability to representing randomness, which appears to be very different (see [4]).

2.3 Possibility theory and decision making

Possibility theory and closely related formalisms have been used as a basis for a qualitative version of decision theory. We first recall that a decision-theoretic model needs two scales, one for uncertainty and one for “goalness” (i.e., satisfaction) – in Bayesian decision theory these are $[0, 1]$ (probability scale) and the real line (utility scale), the optimal decisions being those maximizing the expected utility $\sum_s pr(s) \cdot u(s, a)$,

²One may think of doing it by specifying equiprobability, i.e., if $|S| = n$, then $\forall s \in S, pr(s) = \frac{1}{n}$; however, this would be wrong, since equiprobability is much more informative than the specification of a classical subset – or equivalently, equipossibility; for instance, notice that from equipossibility we get $\Pi(\{s_1, s_2\}) = \Pi(\{s_3\}) = 1$, hence the subsets $\{s_1, s_2\}$ and $\{s_3\}$ have the same possibility degree (and more generally all nonempty subsets have the same possibility, i.e., 1) – while from equiprobability we get $Pr(\{s_1, s_2\}) = \frac{2}{n}$ and $Pr(\{s_3\}) = \frac{1}{n}$.

where $u(s, a)$ is the utility resulting in performing action a in state s –. Dubois and Prade [24] have proposed a qualitative counterpart to Bayesian decision theory, where uncertainty is modeled by a possibility distribution $\pi : S \rightarrow [0, 1]$ and goalness by a qualitative utility function $q : S \times A \rightarrow [0, 1]$, the optimal decision(s) being these maximizing $\min_s \max(1 - \pi(s), q(s, a))$ which is a pessimistic qualitative counterpart to expected utility. In [24], normative approaches to qualitative decision theory are proposed. On the other hand, Yager [25] proposed to maximize the optimistic criterion $\max_s \min(\pi(s), q(s, a))$ ³. Boutilier [27] gives a logical approach to qualitative decision theory which can be equivalently expressed in possibility theory (although it differs from Dubois and Prade’s model because his pessimistic utility function takes account only on the most plausible worlds).

Now, some possibilistic approaches to *sequential decision making* have been proposed, assuming *full observability* (the agent always knows the actual state of the world before acting), among them [28] and [17]. These approaches (see [17] for a discussion and further details) propose various possibilistic versions of Markov decision processes and dynamic programming, and can thus be considered as a first step to possibilistic decision-theoretic planning. Yet these approaches assume the set of states is represented explicitly and thus they may be inefficient for real planning applications. Our approach is much more in the spirit of AI planning, with a STRIPS-like representation of actions (and an assumption of non-observability, contrarily to the above-mentioned approaches). The planning algorithm implementing our approach can be seen as genuine extension of classical AI planners such as SNLP[3].

2.4 Possibility theory and planning: meaning and benefits

Apart of the aforementioned uses of possibility theory in decision making it seems (to our knowledge at least) that there are up to now few significant applications of possibility theory and fuzzy sets to planning. One noticeable exception is Saffiotti et al.’s multivalued logic for planning and control, used in the robot Flakey [29]; they use fuzzy rules to represent conditional plans, and standard goal regression to generate them. Note that the use of fuzzy sets in their approach differs from ours mainly in the fact that while we use possibility degrees for representing uncertain effects of actions, they use truth degrees for representing the desirability of a plan – thus their interpretation of fuzzy sets is in terms of graded preference.

Our approach makes the following uses of possibility theory:

- action effects are described by possibility distributions $\pi[.|s, a]$ (where s is a state and a an action);
- the initial state is described by a possibility distribution;
- in the computational process, the current state is modeled by a possibility distribution over possible states.

³Note that Yager’s proposal can be recast in Dubois and Prade’s axiomatic framework by modifying one of their postulates – see [26].

This means that possibility theory is here interpreted in terms of *uncertainty*. Let us now focus on the benefits of using possibility theory in planning, especially in comparison with probabilistic planning:

- possibility theory is an *ordinal* model: only the order induced by possibility distributions is important, not the precise values of the degrees. Indeed, the only operations on $[0, 1]$ needed in our framework are min, max and order reversal ($1 - \cdot$). This ordinal aspect of a possibilistic representation (contrarily to probabilistic representations which are intrinsically quantitative) gives the model more *robustness* to imprecision on the degrees and is thus particularly suited to cases where there is a lack of statistical data.
- actions with possibilistic transitions functions $\pi[\cdot|s, a]$ are a *graded generalization of nondeterministic actions*, so that possibilistic planning encompasses nondeterministic planning as a particular case (obtained by allowing only 0 and 1 as possibility degrees). *This is not the case with probability theory*, at least if we work with a single probability distribution as in Bayesian approaches: there is no mean of encoding graded nondeterminism by a probability distribution (see Section 2.2 and the corresponding footnote⁴). This is a particularly interesting point, more so because planning with nondeterministic actions has only received little attention in the literature.
- a possibilistic representation of uncertain effects of actions and uncertain states is consistent with the usual non-monotonic encoding of action effects and default knowledge. The knowledge representation community often represents effects of actions, or more generally evolution of the world, by complete pre-orders, or equivalently by possibility distributions, on states. As a consequence, these approaches can be used upstream of ours, since they provide us with models for representing and/or generating possibilistic actions which can then be used in our planning framework.

3 Representation of a possibilistic planning problem

In this section we define the basic components of a possibilistic planning problem, and two different notions for a solution plan.

⁴Note, however, that a planning problem that involves only pure nondeterministic effects can be encoded correctly, though without any probabilistic meaning, with a probability distribution since all what matters then is to distinguish feasible effects from impossible ones. Any probability distribution that assigns a non-zero probability to each feasible effect will do the job – but this use of probability here is weird and artificial!. For such a problem, a planner of the BURIDAN type, queried to return a solution plan having probability 1 of success, is able to find the kind of secure plan that lead certainly to a goal state (provided any such plan exists).

3.1 Expressions and states

The facts or properties that need to be talked about in the application domain are represented in a finite propositional language by expressions that are conjunctions of atomic sentences (symbols) in either positive or negative form, i.e. conjunctions of literals. For convenience, we shall also occasionally represent an expression as the set of literals involved in the conjunction, an empty set representing no specification at all. A state is a complete description of the world at a time point, that is, a particular expression in which all atomic sentences of the language appear exactly once in a positive or negative form. A state is said to be satisfied by an expression ϵ (denoted $s \models \epsilon$) if and only if each literal of ϵ is in s . We define the set of states satisfied by an expression ϵ as $\mathcal{S}(\epsilon) = \{s \in S/s \models \epsilon\}$.

Agronomic example: In the simple illustrative planning problem considered throughout the rest of this paper, the state of the world is described via five atomic sentences: *favorable-spring*, *sown*, *good-potential*, *pest*, *good-yield*. They express respectively that the spring climate offers favorable sowing and growing conditions, the field has been sown, the early stages of growth create good potential conditions toward the production goal, a pest problem occurs, and the final crop yield is good. For short, we shall express these atomic sentences by f , s , g , p and y respectively. The state corresponding to a favorable spring, a field sown, a good potential of production with respect to early growth stages, the occurrence of a pest problem and a yield that is not good is represented by $f \wedge s \wedge g \wedge p \wedge \bar{y}$. The \mathcal{S} function maps an expression to the set of compatible states. For instance, $\mathcal{S}(f \wedge s \wedge g \wedge p) = \{f \wedge s \wedge g \wedge p \wedge y, f \wedge s \wedge g \wedge p \wedge \bar{y}\}$.

Definition 2 (uncertain states)

Let S denotes the set of all conceivable states. The uncertainty about the current state of the world (at time t) is represented by a possibility distribution π_t over the set S of states such that $\max_{s \in S} \pi_t(s) = 1$. The initial state is described by the possibility distribution π_{init} .

$\pi_t(s)$ conveys what is known at time t about the actual state of the world. $\pi_t(s)$ expresses to what extent it is possible that the real world state is s ; in particular, $\pi_t(s) = 0$ means that s is surely not the real world state, and $\pi_t(s) = 1$ means that nothing prevents s from being the real state. Note that there may be several states s with $\pi_t(s) = 1$.

3.2 Possibilistic actions

The actions considered here can be executed in any world state and their effect depends both on the execution-time state (context-dependent effect) and on chance (nondeterministic effect). The feasible nondeterministic results of the application of an action can be specified by a possibility distribution that enables a ranking of the possible outcomes on the scale of *normality* (i.e. non-exceptionality). More formally a possibilistic action is defined as follows.

Definition 3 (possibilistic actions)

A *possibilistic action*, denoted a , is a set of possibilistic effects, i. e. $a = \{ep_i, i = 1, \dots, m\}$, in which ep_i is the i -th possible effect defined by:

$$ep_i = \langle t_i, (\pi_{i1}, e_{i1}), \dots, (\pi_{in_i}, e_{in_i}) \rangle$$

where $\forall i, j, t_i$ and e_{ij} are expressions, $\pi_{ij} \in]0, 1]$, such that

- for all state s , there is a single i such that $s \models t_i$
- for all i , $\max_{1 \leq j \leq n_i} \pi_{ij} = 1$

The ep_i 's, the t_i 's and the e_{ij} 's are called *possibilistic effects*, *discriminants* and *elementary consequences* respectively. The e_{ij} 's play the role of Add/Delete lists of the STRIPS action model.

If the context defined by t_i is verified before the execution of a , then it is possible at degree π_{ij} that effect e_{ij} is verified after the execution. If π_{ij} is equal to 1, then e_{ij} is a normal effect (i.e. nothing prevents it from occurring), else the smaller π_{ij} the more exceptional e_{ij} . For a given discriminant, the elementary consequences together with their associated degrees of possibility constitute a possibility distribution over the changes to the world.

The first condition in the definition means that the discriminants are exhaustive⁵ : $\cup_i \mathcal{S}(t_i) = S$ and mutually exclusive: $\forall i, j, (i \neq j) \mathcal{S}(t_i) \cap \mathcal{S}(t_j) = \emptyset$. The second one means that any action must have at least one normal elementary consequence.

An empty elementary consequence expresses the case where the execution of the action results in no change of the state it was applied in.

Agronomic example: We consider a simple crop management problem that consists in planning a coherent combination of sowing, pest treating and harvesting operations in order to obtain a satisfactory crop production. As presented in Section 3.1 the state of the world and the actions are described via five atomic sentences: *favorable-spring*, *sown*, *good-potential*, *pest*, *good-yield* represented by f , s , g , p and y respectively. Four actions are available: *sow-normal*, *sow-better*, *treat* and *harvest*.

$$\begin{array}{ll} \textit{sow-normal} = \{ & \textit{sow-better} = \{ \\ \langle f \wedge \bar{s}, (1, s \wedge g), (0.2, s \wedge \bar{g}) \rangle, & \langle f \wedge \bar{s}, (1, s \wedge g), (0.3, s \wedge \bar{g}) \rangle, \\ \langle \bar{f} \wedge \bar{s}, (1, s \wedge g), (0.7, s \wedge \bar{g}) \rangle, & \langle \bar{f} \wedge \bar{s}, (1, s \wedge g \wedge p), (0.4, s \wedge \bar{g}) \rangle, \\ \langle s, (1, \{\}) \rangle \} & \langle s, (1, \{\}) \rangle \} \end{array}$$

⁵If in practice the original formulation of an action does not naturally verify the exhaustiveness property, one can always modify that formulation by adding complementary discriminants associated to null effects. Actually this property is only useful for a clear presentation of our formalism and is not required by the planning algorithms presented in the next section.

$$\begin{array}{l}
\text{treat} = \{ \\
\langle p, (1, \bar{p}), (0.1, \{\}) \rangle, \\
\langle \bar{p}, (1, \{\}) \rangle \} \\
\text{harvest} = \{ \\
\langle s \wedge p, (1, \bar{y}) \rangle, \\
\langle s \wedge g \wedge \bar{p}, (1, y), (0.2, \bar{y}) \rangle, \\
\langle s \wedge \bar{g} \wedge \bar{p}, (1, y), (0.8, \bar{y}) \rangle, \\
\langle \bar{s}, (1, \{\}) \rangle \}
\end{array}$$

The first two are sowing actions that correspond to two different choices of seeds. The first one tells that: i) in the context of a field not sown and a favorable spring, the normal effect is to have a good potential and a field sown and very exceptionally (possibility= 0.2) to have a not so good potential and the field sown, ii) in the context of a field not sown and an unfavorable spring, the normal effect is to have a good potential and a field sown but it is much less exceptional (possibility=0.7) to have a not so good potential and the field sown, iii) in the context of a field already sown, this action does not make any change. The *sow-better* action has the same discriminants but differs from *sow-normal* on the first possibilistic effects since the exceptional outcome (i.e. $s \wedge \bar{g}$) is slightly less exceptional (possibility=0.3 instead of 0.2) and on the second possibilistic effect since the normal outcome is to have p in addition to s and g and the exceptional outcome is significantly more exceptional (possibility= 0.4 instead of 0.7). The *treat* action has the normal effect of eradicating the pest problem if any and very exceptionally (possibility=0.1) fails to do so. It does not change the world in case no pest problem occurs. The *harvest* action is effective when the field has been sown. If there is a pest problem, then the yield is definitely not good. Otherwise, if the potential is good, then normally the yield should be good but may be bad very exceptionally and if the potential is not good, the normal outcome is still a good yield but it is not very exceptional (possibility=0.8) to end up with the opposite effect.

Definition 4 (effect of a possibilistic action)

Let s a state, a an action and e an elementary consequence of a , the state resulting from the change on s caused by e is defined by:

$$Res(e, s) = e \cup \{l \in s \mid \bar{l} \notin e\}$$

The result of executing action a on s is given by a possibility distribution on S defined by:

$$\pi[s'|s, a] = \begin{cases} \max_k \pi_{ik} & \text{if } s \in S(t_i) \text{ and } s' = Res(e_{ik}, s) \\ 0 & \text{otherwise} \end{cases}$$

If the initial state is described by a possibility distribution π_{init} over S , then the effect of executing a is defined by the following possibility distribution:

$$\pi[s'|\pi_{init}, a] = \max_{s_0} \min(\pi[s'|s_0, a], \pi_{init}(s_0))$$

For instance, considering the *treat* action, we have:

$$\begin{aligned}\pi[f \wedge s \wedge g \wedge p \wedge y \mid f \wedge s \wedge g \wedge p \wedge y, \textit{treat}] &= 0.1 \\ \pi[f \wedge s \wedge g \wedge p \wedge y \mid f \wedge s \wedge g \wedge \bar{p} \wedge y, \textit{treat}] &= 0\end{aligned}$$

Some interesting families of possibilistic actions (or, equivalently, graded nondeterministic actions) are worth identifying; let a an action and $\pi[\cdot|s, a]$ the possibility distribution induced by the action a applied on state s :

- action a is *nondeterministic* if and only if $\forall s, s' \in S, \pi[s'|s, a] \in \{0, 1\}$;
- action a is *pseudo-deterministic* if and only if $\forall s$ there is a single s' such that $\pi[s'|s, a] = 1$;
- action a is *deterministic* if and only if a is both *nondeterministic* and *pseudo-deterministic*, i.e $\forall s, \exists s'$ such that $\pi[s'|s, a] = 1$ and $\forall s'' \neq s', \pi[s''|s, a] = 0$.

The set relations among these families are summarized in Figure (1). In our example, all actions are pseudo-deterministic.

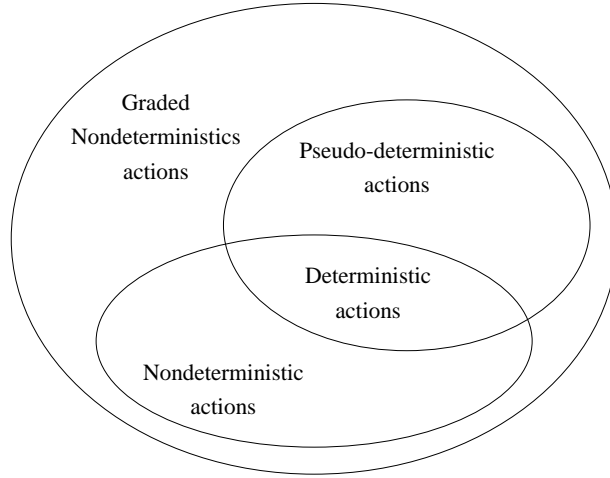


Figure 1: *Different families of possibilistic actions.*

The set \mathcal{A} of actions available in the planning problem under consideration will be said *nondeterministic* (*pseudo-deterministic* or *deterministic* respectively) if all the actions that it contains are *nondeterministic* (*pseudo-deterministic* or *deterministic* respectively). Finally, we will say that \mathcal{A} is a k -level action set if the set of distinct possibility degrees $\{\pi[s'|s, a] \mid s, s' \in S; a \in \mathcal{A}\}$ has k elements⁶. In our example, $\mathcal{A} = \{\textit{sow-normal}, \textit{sow-better}, \textit{treat}, \textit{harvest}\}$ is an 8-level action set.

Thus, given a pseudo-deterministic 3-level action set $\{0, \alpha, 1\}$ and a state s , every action a yields a *normal resulting state* (s' such that $\pi[s'|s, a] = 1$) and at most one *exceptional resulting state*: (s'' such that $\pi[s''|s, a] = \alpha$). A set \mathcal{A} with more than 3 levels allows to take account of different degrees of exception.

⁶We impose that 0 is always in this set; thus, \mathcal{A} has 2 levels if and only if \mathcal{A} is nondeterministic.

3.3 Plans and possibilistic planning

We classically define a plan of actions as a set of ordered actions.

Definition 5 (plan of actions)

A *sequential plan* is a totally ordered set of actions $\langle a_i \rangle_{i=0}^{N-1}$.

A *partially ordered plan* is a pair $\mathcal{P} = (A, O)$ where A is a set of actions and O is a set of ordering constraints between these actions. A *completion* of \mathcal{P} is a sequential plan $\mathcal{CP} = \langle a_i \rangle_{i=0}^{N-1}$ such that $A = \{a_0, \dots, a_{N-1}\}$ and the total ordering $a_0 < \dots < a_{N-1}$ is consistent with O .

A *consistent partially ordered plan* is a plan $\mathcal{P} = (A, O)$ with a consistent set O of ordering constraints.

To execute a plan \mathcal{P} is to execute in sequence a_0, a_1, \dots, a_{N-1} where $\langle a_i \rangle_{i=0}^{N-1}$ is a completion of \mathcal{P} . The strong assumption of non-observability underlies the way the execution is monitored. Indeed we suppose that despite the uncertainty concerning the effects of the actions, a plan is executed *blindly*, without any information gathering between steps. The following proposition relies on this assumption.

Proposition 1 (effect of a sequential plan on a state)

Let s_0 be a state of S . The possibility to reach a given state s_N by executing a sequential plan of possibilistic actions $\langle a_i \rangle_{i=0}^{N-1}$ starting in s_0 , is given by:

$$\begin{aligned} \pi[s_N | s_0, \langle a_i \rangle_{i=0}^{N-1}] &= \max_{s_1} \min(\pi[s_1 | s_0, a_0], \pi[s_N | s_1, \langle a_i \rangle_{i=1}^{N-1}]) \\ &= \max_{\langle s_1 \dots s_{N-1} \rangle} \min_{i=0 \dots N-1} \pi[s_{i+1} | s_i, a_i] \end{aligned}$$

where $\langle s_1 \dots s_{N-1} \rangle$ represents a sequence of states visited from s_1 to s_{N-1} .

Proof: Executing $\langle a_i \rangle_{i=0}^{N-1}$ from s_0 is equivalent to execute $\langle a_i \rangle_{i=1}^{N-1}$ from the possibilistic distribution $\pi[\cdot | s_0, a_0]$. Thus:

$$\begin{aligned} \pi[s_N | s_0, \langle a_i \rangle_{i=0}^{N-1}] &= \pi[s_N | \pi[\cdot | s_0, a_0], \langle a_i \rangle_{i=1}^{N-1}] \\ &= \max_{s_1} \min(\pi[s_N | s_1, \langle a_i \rangle_{i=1}^{N-1}], \pi[s_1 | s_0, a_0]) \end{aligned}$$

Then,

$$\begin{aligned} &\max_{s_1} \min(\pi[s_1 | s_0, a_0], \pi[s_N | s_1, \langle a_i \rangle_{i=1}^{N-1}]) \\ &= \max_{s_1} \min(\pi[s_1 | s_0, a_0], \max_{s_2} \min(\pi[s_2 | s_1, a_1], \pi[s_N | s_2, \langle a_i \rangle_{i=2}^{N-1}])) \\ &= \max_{s_1} \max_{s_2} \min(\pi[s_1 | s_0, a_0], \min(\pi[s_2 | s_1, a_1], \pi[s_N | s_2, \langle a_i \rangle_{i=2}^{N-1}])) \\ &= \max_{s_1} \max_{s_2} \min(\pi[s_1 | s_0, a_0], \pi[s_2 | s_1, a_1], \pi[s_N | s_2, \langle a_i \rangle_{i=2}^{N-1}]) \\ \dots &= \max_{s_1} \max_{s_2} \dots \max_{s_{N-1}} \min(\pi[s_1 | s_0, a_0], \dots, \pi[s_N | s_{N-1}, a_{N-1}]) \\ &= \max_{\langle s_1 \dots s_{N-1} \rangle} \min_{i=0, \dots, N-1} \pi[s_{i+1} | s_i, a_i] \end{aligned}$$

□

Let $Goals \subseteq S$ the set of the goal states, and π_{init} a possibility distribution over S that describes the initial state s_0 . The possibility and necessity measures to reach a goal state after the execution from s_0 of the sequential plan $\langle a_i \rangle_{i=0}^{N-1}$ are given by:

$$\begin{aligned} \Pi[Goals | \pi_{init}, \langle a_i \rangle_{i=0}^{N-1}] &= \max_{s_0 \in S} \min(\Pi[Goals | s_0, \langle a_i \rangle_{i=0}^{N-1}], \pi_{init}(s_0)) \\ &= \max_{s_0 \in S, s_N \in Goals} \min(\pi[s_N | s_0, \langle a_i \rangle_{i=0}^{N-1}], \pi_{init}(s_0)) \\ N[Goals | \pi_{init}, \langle a_i \rangle_{i=0}^{N-1}] &= 1 - \Pi[\overline{Goals} | \pi_{init}, \langle a_i \rangle_{i=0}^{N-1}] \\ &= \min_{s_0 \in S, s_N \in \overline{Goals}} \max(1 - \pi_{init}(s_0), 1 - \pi[s_N | s_0, \langle a_i \rangle_{i=0}^{N-1}]) \end{aligned}$$

In our example, if $\pi_{init} = \{(1, \bar{f} \wedge \bar{g} \wedge \bar{s} \wedge \bar{p} \wedge \bar{y})\}$ and $Goals = \mathcal{S}(y)$ then:

$$\begin{aligned} \Pi[Goals | \pi_{init}, \langle \text{sow-better}, \text{treat}, \text{harvest} \rangle] &= 1 \\ N[Goals | \pi_{init}, \langle \text{sow-better}, \text{treat}, \text{harvest} \rangle] &= 0.6 \end{aligned}$$

Definition 6 (possibilistic planning problem)

A *possibilistic planning problem* Δ is a triplet $\langle \pi_{init}, \epsilon_{Goals}, \mathcal{A} \rangle$ where π_{init} is the possibility distribution associated to the initial state, ϵ_{Goals} is an expression defining the set of goal states $Goals$ and \mathcal{A} is the set of available possibilistic actions.

Given a possibilistic planning problem, two criteria may be considered to define a *solution plan*:

Definition 7 (solution plans)

Let Δ be a possibilistic planning problem and \mathcal{P} be a sequential plan:

- \mathcal{P} is a **γ -acceptable plan** if $N[Goals | s_0, \mathcal{P}] \geq \gamma$;
- \mathcal{P} is an **optimally safe plan**, or simply, **optimal plan** if $N[Goals | s_0, \mathcal{P}]$ is maximal among all possible sequential plans.

This definition can be extended to partially ordered sets of actions. Let \mathcal{P} be a consistent partially ordered plan :

- \mathcal{P} is a **γ -acceptable plan** if $N[Goals | s_0, \mathcal{CP}] \geq \gamma$ for all totally ordered completion \mathcal{CP} of \mathcal{P} ;
- \mathcal{P} is an **optimal plan** if $N[Goals | s_0, \mathcal{CP}]$ is maximal among all possible sequential plans for all totally ordered completion \mathcal{CP} of \mathcal{P} .

4 Generation of solution plans

The two planning algorithms that we present in this section for solving a possibilistic planning problem are based on the equivalence between the search of γ -acceptable plans and the resolution of a derived planning problem that has only pure nondeterministic actions. Before presenting these general planning algorithms we first recall the characteristics of nondeterministic planning.

4.1 Nondeterministic planning and γ -acceptability

We have seen that a nondeterministic action a is a special case of a possibilistic action where $\forall s, s' \in S, \pi[s'|s, a] \in \{0, 1\}$. It is equivalent to impose on this action the condition $\forall i, j, \pi_{ij} = 1$. Thus we can simplify its definition as follows.

Definition 8 (nondeterministic actions)

A nondeterministic action is a set of nondeterministic effects $a = \{ef_i, i = 1, \dots, m\}$, in which ef_i is the i -th nondeterministic effect defined by $ef_i = \langle t_i, e_{i1}, \dots, e_{in_i} \rangle$ where $\forall i, j, t_i$ and e_{ij} are expressions, and the discriminants t_i are exhaustive and mutually exclusive.

Since possibility distributions are no longer necessary to describe nondeterministic actions, we can also simplify the description of the effect of an action on a state: let s be a state and a a nondeterministic action; the result of executing action a on s is given by a subset of S defined by:

$$Res(a, s) = \{s' \mid s \in S(t_i) \text{ and } s' = Res(e_{ik}, s)\}$$

Let $I \subset S$ be a set of states. The result of executing nondeterministic action a on I is given by:

$$Res(a, I) = \bigcup_{s \in I} Res(a, s)$$

We can also define the effect of a sequential plan on a set of states: let $\langle a_i \rangle_{i=0}^{N-1}$ be a sequence of nondeterministic actions; the result of executing $\langle a_i \rangle_{i=0}^{N-1}$ on I is recursively defined by:

$$Res(\langle a_i \rangle_{i=0}^{N-1}, I) = Res(\langle a_i \rangle_{i=1}^{N-1}, Res(a_0, I))$$

We can now define a planning problem with nondeterministic actions, and a solution of that problem:

Definition 9 (planning with nondeterministic actions)

A nondeterministic planning problem Δ is a triplet $\langle I_{init}, \epsilon_{Goals}, \mathcal{A} \rangle$ where I_{init} is the set of possible initial states, ϵ_{Goals} is an expression defining the set of goal states $Goals$ and \mathcal{A} is the set of available nondeterministic actions.

*A sequential plan $\langle a_i \rangle_{i=0}^{N-1}$ is a **safe plan** for Δ if*

$$Res(\langle a_i \rangle_{i=0}^{N-1}, I_{init}) \subset Goals$$

A partially ordered plan \mathcal{P} is a safe plan for Δ if each totally ordered completion \mathcal{CP} of \mathcal{P} is a safe plan for Δ .

We finally present how, given a possibilistic planning problem, the search of a plan having a necessity greater than γ can be transformed equivalently into the problem of finding a safe plan for a particular nondeterministic planning problem derived from the original possibilistic one.

Definition 10 (from possibilistic to nondeterministic planning)

Let $\Delta = \langle \pi_{init}, \epsilon_{Goals}, \mathcal{A} \rangle$ a possibilistic planning problem and $\gamma \in]0, 1]$. The nondeterministic planning problem $\Delta_{1-\gamma}$ constructed from Δ is defined by:

$$\Delta_{1-\gamma} = \langle I_{init1-\gamma}, \epsilon_{Goals}, \mathcal{A}_{1-\gamma} \rangle \quad \text{where}$$

- $I_{init1-\gamma} = \{s \in S \mid \pi_{init}(s) > 1 - \gamma\}$
- $\mathcal{A}_{1-\gamma} = \{a_{1-\gamma} \mid a \in \mathcal{A}\}$ such that
if $a = \{\langle t_i, \dots (\pi_{ij}, e_{ij}) \dots \rangle\}$ then $a_{1-\gamma} = \{\langle t_i, \dots e_{ij} \dots \rangle \mid \pi_{ij} > 1 - \gamma\}$.

$I_{init1-\gamma}$ is the set of initial states that have a possibility greater than $1 - \gamma$. The nondeterministic action $a_{1-\gamma} \in \mathcal{A}_{1-\gamma}$ is the result of transforming the action a in \mathcal{A} by retaining only the effects having a possibility greater than $1 - \gamma$.

In our example $\mathcal{A}_{0.5} = \{sow-normal_{0.5}, sow-better_{0.5}, treat_{0.5}, harvest_{0.5}\}$, where

$$\begin{array}{llll} sow-normal_{0.5} = \{ & sow-better_{0.5} = \{ & treat_{0.5} = \{ & harvest_{0.5} = \{ \\ \langle f \wedge \bar{s}, s \wedge g \rangle, & \langle f \wedge \bar{s}, s \wedge g \rangle, & \langle p, \bar{p} \rangle, & \langle s \wedge p, \bar{y} \rangle, \\ \langle \bar{f} \wedge \bar{s}, s \wedge g, s \wedge \bar{g} \rangle, & \langle \bar{f} \wedge \bar{s}, s \wedge g \wedge p \rangle, & \langle \bar{p}, \{\} \rangle \} & \langle s \wedge g \wedge \bar{p}, y \rangle, \\ \langle s, \{\} \rangle \} & \langle s, \{\} \rangle \} & & \langle s \wedge \bar{g} \wedge \bar{p}, y, \bar{y} \rangle, \\ & & & \langle \bar{s}, \{\} \rangle \} \end{array}$$

The following key result shows that it is equivalent to search the safe plans of $\Delta_{1-\gamma}$ and to search the γ -acceptable plans that solve Δ .

Proposition 2 (equivalence)

A partial plan \mathcal{P} is γ -acceptable for Δ if and only if it is safe for $\Delta_{1-\gamma}$.

Proof: we prove it for a sequential plan $\langle a_i \rangle_{i=0}^{N-1}$.

$$\begin{aligned} N[Goals \mid \pi_{init}, \langle a_i \rangle_{i=0}^{N-1}] \geq \gamma &\Leftrightarrow \min_{s_N \in \overline{Goals}} \min_{s_0} \max(1 - \pi_{init}(s_0), 1 - \pi[s_N \mid s_0, \langle a_i \rangle_{i=0}^{N-1}]) \geq \gamma \\ &\Leftrightarrow \forall s_N \in \overline{Goals} \max_{s_0} \min(\pi_{init}(s_0), \pi[s_N \mid s_0, \langle a_i \rangle_{i=0}^{N-1}]) \leq 1 - \gamma \\ &\Leftrightarrow \forall s_N \in \overline{Goals}, \forall s_0 \mid \pi_{init}(s_0) > 1 - \gamma, \\ &\quad \pi[s_N \mid s_0, \langle a_i \rangle_{i=0}^{N-1}] \leq 1 - \gamma \\ &\Leftrightarrow \forall s_N \in \overline{Goals}, \forall s_0 \mid \pi_{init}(s_0) > 1 - \gamma, \end{aligned}$$

$$\begin{aligned}
& \max_{\langle s_1 \dots s_{N-1} \rangle} \min_{i=0 \dots N-1} \pi[s_{i+1}|s_i, a_i] \leq 1 - \gamma \\
\Leftrightarrow & \forall s_N \in \overline{Goals}, \forall s_0 \mid \pi_{init}(s_0) > 1 - \gamma, \\
& \forall \langle s_1 \dots s_{N-1} \rangle, \exists i \mid \pi[s_{i+1}|s_i, a_i] \leq 1 - \gamma \\
\Leftrightarrow & \forall s_0 \mid \pi_{init}(s_0) > 1 - \gamma, \forall \langle s_1 \dots s_N \rangle, \\
& \forall i \pi[s_{i+1}|s_i, a_i] > 1 - \gamma \Rightarrow s_N \in Goals \\
\Leftrightarrow & Res(\langle a_i \rangle_{i=0}^{N-1}, I_{init}) \subset Goals
\end{aligned}$$

The proposition extends to partial plans by the way they are defined:

$$\begin{aligned}
\mathcal{P} \text{ is } \gamma\text{-acceptable for } \Delta & \Leftrightarrow \text{each completion } \mathcal{CP} \text{ of } \mathcal{P} \text{ is } \gamma\text{-acceptable for } \Delta \\
& \Leftrightarrow \text{each completion } \mathcal{CP} \text{ of } \mathcal{P} \text{ is safe for } \Delta_{1-\gamma} \\
& \Leftrightarrow \mathcal{P} \text{ is safe for } \Delta_{1-\gamma}
\end{aligned}$$

□

In particular, for pseudo-deterministic actions with 3 levels $\{0, \alpha, 1\}$ we have:

- $\gamma = 1$: $\langle a_i \rangle_{i=1}^{N-1}$ is 1-acceptable if and only if for every effect of action a_i (normal or exceptional), the plan $\langle a_i \rangle_{i=1}^{N-1}$ leads to a goal state.
- $\gamma = 1 - \alpha$: $\langle a_i \rangle_{i=1}^{N-1}$ is $(1 - \alpha)$ -acceptable if and only if it leads to a goal state whenever the effects of all actions are normal. In addition, the problem Δ_α contains only deterministic actions. So, a $(1 - \alpha)$ -acceptable plan can be computed with a classical planning method.

4.2 A planning algorithm for γ -acceptability

Given a possibilistic planning problem Δ we have seen how to transform it into a nondeterministic planning problem $\Delta_{1-\gamma}$. This transformation is the core of our possibilistic planning system entitled POSPLAN (Figure 2). The nondeterministic planning algorithm NDP, described in the rest of this section, takes a nondeterministic planning problem $\Delta_{1-\gamma} = \langle I_{init}, \epsilon_{Goals}, \mathcal{A} \rangle$ as input, and generates a partially ordered safe plan, which is γ -acceptable for Δ .

<p>Algorithm POSPLAN(Δ, γ)</p> <ol style="list-style-type: none"> 1. transform Δ into a nondeterministic planning problem $\Delta_{1-\gamma}$; 2. return $\mathcal{P} \leftarrow NDP(\Delta_{1-\gamma})$
--

Figure 2: POSPLAN: general algorithm

4.2.1 NDP: a planning algorithm for nondeterministic actions

Like most of the classical planning algorithms, NDP first transforms a nondeterministic planning problem Δ into an null plan containing the two pseudo-operators $a_{begin} = \{\langle \emptyset, s_0^1, \dots, s_0^n \rangle\}$, with $I_{init} = \{s_0^1, \dots, s_0^n\}$, and $a_{end} = \{\langle \epsilon_{Goals}, \emptyset \rangle\}$ with the ordering constraint $a_{begin} < a_{end}$. Then NDP planning algorithm explores a search tree of partial plans whose root is the null plan and the branches represent refinements of the current plan in order to establish subgoals or to confront threats on some already established causal links (Figure 3).




Figure 3: *NDP planning algorithm*

4.2.2 Refinement of a partial plan

The main function of a planning algorithm like NDP is the refinement which transforms a partial plan \mathcal{P} into a new partial plan \mathcal{P}' , by eliminating *flaws*. A flaw in \mathcal{P} can be a subgoal not yet established, or a threat on a subgoal already established, that prevents \mathcal{P} from being a solution plan (Figure 4).

Like in SNLP [3], UCPOP [8] or BURIDAN [1] NDP establishes a subgoal proposition p by adding a *causal link* between the effect of an operator that add p and the action of which p is a precondition. More formally, if $a_j : t_k : p$ is a subgoal of \mathcal{P} where $p \in t_k$ and t_k a discriminant of a_j in \mathcal{P} , and if $p \in e_{lm}$ where e_{lm} is an elementary consequence of a nondeterministic effect of the action a_i , then NDP can add the causal link $a_i : e_{lm} \xrightarrow{p} a_j : t_k$. Consequently, each proposition $q \in t_l$ becomes a new subgoal to establish. To each partial plan in the search graph is associated a set of subgoals \mathcal{SG} and a set of causal links \mathcal{CL} . They are initialized with $\mathcal{SG} = \{end : \epsilon_{Goals} : q \forall q \in \epsilon_{Goals}\}$ and $\mathcal{CL} = \emptyset$.

A threat in a partial plan $\mathcal{P} = (A, O)$ is a pair $\langle a; l \rangle$ where l is a causal link $a' : e \xrightarrow{p} a'' : t$ in \mathcal{CL} establishing a proposition p , and a is an action in A that potentially can delete p , i.e. one of its elementary consequences contains $\neg p$. In a classical propositional planner like SNLP, there are two ways to confront such a threat: you can add the constraint $a'' < a$ (*promotion*) or add the constraint $a < a'$ (*demotion*). NDP naturally keeps these two techniques.

Due to the specificity of the problem of planning actions with context-dependent effects, the refinement function of NDP must use another trick to protect a causal link as it is done in the two planners UCPOP and BURIDAN. Consider the example

Algorithm *refine*(\mathcal{P})

1. choose *flaw* from \mathcal{P}
2. if *flaw* = $a : t : p$ is a subgoal of \mathcal{SG} then:
 - (a) nondeterministically choose an existing action a_i in \mathcal{P} or add $a_i \in \mathcal{A}$ to \mathcal{P} such that a_i has an elementary consequence e_{kl} of a nondeterministic effect e_k with $p \in e_{kl}$;
 - (b) add the causal link $a_i : e_{kl} \xrightarrow{p} a : t$ to \mathcal{CL} and $a_i < a$ to \mathcal{P} ;
 - (c) add the subgoals $a_i : t_k : q$ to \mathcal{SG} for all $q \in t_k$ discriminant of e_k ;
3. if *flaw* = $\langle a, l \rangle$ where a is an action of \mathcal{P} that threatens the causal link $l = a_j : e \xrightarrow{p} a_j : t$ in \mathcal{CL} , then nondeterministically choose:
 - (a) (*demotion*) add $a < a_i$ to \mathcal{P} ;
 - (b) (*promotion*) add $a_j < a$ to \mathcal{P} ;
 - (c) (*confrontation*) add a *safety proposition* sp unique to this threat to each non threatening elementary consequence $a : e'$ and to the discriminant $a_j : t$. Add the subgoal $a_j : t : sp$ to \mathcal{SG} ;
4. *return* the new plan \mathcal{P} .

Figure 4: *Refinement function*

in Figure (5) that contains two deterministic actions A and B . The plan $\langle A \rangle$ is not safe since the effect associated to the third discriminant of A threatens the goal proposition p . If we only use demotion and promotion, $\langle A \rangle$ leads to failure. However, there exists a solution that consists in adding an action B to prevent the occurrence of effect $\neg p$, even though B does not directly establish a subgoal of \mathcal{P} .

In UCPOP, the following particular form of *separation* is used: if $a : e$ is a threatening effect for a causal link l of \mathcal{CL} , and if t is the trigger of this effect $a : e$, then a subgoal $a : \neg q$ is added to \mathcal{SG} where q is any proposition of t , to prevent the trigger t to be active and the effect e to threaten the causal link l . BURIDAN uses a similar approach, called *confrontation*. If $a : e$ is a threatening effect for a causal link l of \mathcal{CL} , BURIDAN will try to make another non threatening effect $a : e'$ occur. Thus, instead of planning for $\neg t$ where t is the trigger of $a : e$, BURIDAN will try to reinforce the possibility to realize $a : e'$. This is done by adding to the discriminant of the causal link consumer a *safety proposition* sp unique to this threat, which becomes a new subgoal to achieve, and by adding to each non threatening elementary consequence $a : e'$ the proposition sp .

These confrontation techniques are necessary in UCPOP and BURIDAN planning algorithms to prove their completeness. Actually they are very similar since the discriminants t_i 's induce a partition of S and thus to make false one discriminant




Figure 5: *Demotion and promotion are not enough.*

t_i is equivalent to make true one of the discriminant t_j , with $j \neq i$. In the NDP planner, we retain the confrontation form of BURIDAN.

4.2.3 Assessment of a partial plan

In classical planning with STRIPS like operators, a polynomial truth criterion can be used to check whether or not a partial plan is a solution plan[30, 31]. In that case the causal link structure is a way to simplify the management of this truth criterion. The problem of planning actions with context-dependent effects like in UCPOP is not so simple. It has been shown that the assessment problem for actions with context-dependent effect was NP-hard [30, 32]. This result is still valid in our non-deterministic action representation since we are not in the favorable situation in which a subgoal proposition in \mathcal{P} can be established by the same action whatever the completions \mathcal{CP} and the uncertainty outcomes might be. Hence, like in BURIDAN, a subgoal is not necessary achieved by a unique operator in NDP, and it is possible to add several different causal links to establish it⁷. A consequence is that we cannot produce a general assessment algorithm that runs efficiently on every problem. Like in BURIDAN, we propose to check if a partial plan \mathcal{P} is a safe plan by directly executing each of its totally ordered completions from each possible initial state belonging to I_{init} (Figure 6). In that case, the causal link structure attached to a partial plan is not used to establish whether it is a safe plan or not, but instead to implement efficiently the *threat* management.




Figure 6: *Assessment algorithm*

⁷See also [33] for an analysis of the multi-contributors causal link structure in planning.

4.2.4 Soundness and completeness of POSPLAN

Given the assessment procedure we have retained, the soundness property of NDP is evident. From the proposition (2) we directly show the following result:

Proposition 3 (Soundness of POSPLAN)

Let $\Delta = \langle \pi_{init}, \epsilon_{Goals}, \mathcal{A} \rangle$ be a possibilistic planning problem and $\gamma \in]0, 1]$.

If $POSPLAN(\Delta, \gamma)$ returns a partial solution plan \mathcal{P} , then \mathcal{P} is a γ -acceptable plan for Δ : each totally ordered completion $\langle a_i \rangle_{i=0}^{N-1}$ of \mathcal{P} has a necessity measure greater than γ .

The completeness of the nondeterministic planning algorithm NDP can also be established. The demonstration of this property is very similar to the ones for SNLP, UCPOP or BURIDAN. In the following, an *essential solution plan* is a totally ordered solution plan $\langle a_i \rangle_{i=0}^{N-1}$ such that no subplan $\langle a_i \rangle_{i=0}^{p-1}$, $p < N$, can be a solution plan.

Proposition 4 (Completeness of NDP)

Let $\Delta = \langle I_{init}, \epsilon_{Goals}, \mathcal{A} \rangle$ be a nondeterministic planning problem. If there exists an essential solution plan $\langle a_i \rangle_{i=0}^{N-1}$ that is safe for Δ , then $NDP(\Delta)$ will generate a partial solution plan \mathcal{P} such that $\langle a_i \rangle_{i=0}^{N-1}$ is one of its completions.

Proof: Based on the following lemma.

Lemma 1

Let \mathcal{P} be a plan generated by $NDP(\Delta)$, and $\langle a_i \rangle_{i=0}^{N-1}$ a completion of \mathcal{P} . If $\langle a_{i_j} \rangle_{j=0}^{M-1}$, $M \geq N$, is an essential safe plan for Δ such that $\langle a_i \rangle_{i=0}^{N-1}$ is a subplan of $\langle a_{i_j} \rangle_{j=0}^{M-1}$, then $NDP(\Delta)$ will generate a path between \mathcal{P} and an essential safe \mathcal{P}' , such that $\langle a_{i_j} \rangle_{j=0}^{M-1}$ is a completion of \mathcal{P}' .

proof (sketch) : let $B = \{a_{i_j}, j = 0, \dots, M-1\} - \{a_i, i = 0, \dots, N-1\} = \{b_i, i = 0, \dots, l-1\}$. By induction on $l = M - N$: If $l = 0$, we just have $\mathcal{P}' = \mathcal{P}$. If $l > 0$, we show that it is possible to add the last action b_{l-1} of B . Indeed, since $\langle a_{i_j} \rangle_{j=0}^{M-1}$ is essential, either b_{l-1} adds a new trajectory to the goal, or b_{l-1} eliminates an existing trajectory that prevents the achievement of the goal. In the first case, b_{l-1} necessarily establishes a discriminant proposition of an operator directly or indirectly implied in the establishment of ϵ_{Goals} , and then can be added by NDP with the line 2(a) in Figure 4; In the second case, the effect of b_{l-1} is to establish a non-interfering discriminant proposition of an operator that could otherwise make ϵ_{Goals} false. Thus b_{l-1} can be added by NDP, with the lines 3(c) and 2(a) in Figure 4. Then it is also possible to show that the necessary constraints can be added with b_{l-1} , to define a partial plan \mathcal{P}^1 such that $\langle a_0 \dots b_{l-1} \dots a_{N-1} \rangle$ is a completion of \mathcal{P}^1 . Then the inductive assumption insures that from \mathcal{P}^1 NDP can generate a path leading to \mathcal{P}' , with $\langle a_{i_j} \rangle_{j=0}^{M-1}$ a completion of \mathcal{P}' . \square

From this lemma, we set $\mathcal{P} = \mathcal{P}_0$, and $N = 0$. We deduce the proposition (4). \square

Note that as any planning algorithm NDP is only semi-decidable.

Then from the completeness of NDP and the equivalence proposition (2), we easily derive the completeness of POSPLAN:

Proposition 5 (Completeness of POSPLAN)

Let $\Delta = \langle \pi_{init}, \epsilon_{Goals}, \mathcal{A} \rangle$ be a possibilistic planning problem and $\gamma \in]0, 1]$. If there exists an essential solution plan $\langle a_i \rangle_{i=0}^{N-1}$ such that $N[Goals | \pi_{init}, \langle a_i \rangle_{i=0}^{N-1}] \geq \gamma$ then $POSPLAN(\Delta, \gamma)$ will generate a partial solution plan \mathcal{P} such that $\langle a_i \rangle_{i=0}^{N-1}$ is one of its completions.

Example: for $\gamma = 0.5$, and $\Delta = \langle \pi_{init} = \{(1, \bar{f} \wedge \bar{g} \wedge \bar{s} \wedge \bar{p} \wedge \bar{y})\}, \epsilon_{Goals} = y, \mathcal{A} \rangle$, $POSPLAN(\Delta, \gamma)$ considers the nondeterministic planning problem $\Delta_{0.5} = \langle \{\bar{f} \wedge \bar{g} \wedge \bar{s} \wedge \bar{p} \wedge \bar{y}\}, y, \mathcal{A}_{0.5} \rangle$. Then NDP generates the safe plan $\langle sow\text{-}better_{0.5}, treat_{0.5}, harvest_{0.5} \rangle$ and thus $\langle sow\text{-}better, treat, harvest \rangle$ is a 0.5-acceptable solution plan. Note, however, that $N[Goals | \pi_{init}, \langle sow\text{-}better, treat, harvest \rangle] = 0.6$.

4.3 A planning algorithm for generating optimal plans

4.3.1 Optimizing acceptability

The link between γ -acceptable plans and optimal plans is clear: $\langle a_i \rangle_{i=1}^{N-1}$ is optimal if and only if it is γ -acceptable and $\forall \gamma' > \gamma$, there is no γ' -acceptable plan. Therefore a meta-algorithm for computing an optimal plan might consist in searching γ -acceptable plans with well-chosen successive values of γ . Several strategies can be thought of; they all require that the state transition degrees (i.e. the possibility degrees involved in the actions and the initial state) be ordered beforehand. Assume that $\{\alpha_i, i = 0, n\}$ is the set of possibility degrees named such that $0 = \alpha_0 < \alpha_1 < \alpha_2 < \dots < \alpha_n < 1$ or, equivalently, $1 = \gamma_0 > \gamma_1 > \dots > \gamma_{n-1} > \gamma_n > 0$ with $\gamma_i = 1 - \alpha_i$.

One can then search iteratively a γ -acceptable plan by letting γ decrease from 1 till a plan is found or, conversely, by letting γ increase from 0 till no plan can be found. Alternatively, another strategy, called dichotomic search, would be to start with γ somewhere in the middle of the scale and iterate the process with a scale truncated from the above in case of absence of solution and truncated from the below in the opposite case.

The search by increasing acceptability has some worth-mentioning computational properties:

- the algorithm is “anytime” in the sense that it can supply a solution at any time (provided at least one exists) and the supplied solution is all the better as the algorithm runs longer;
- the planning problems generated at each iteration are more and more difficult; consequently, the first iteration should be faster and, in particular, if the actions are pseudo-deterministic, the planning problem solved at the first iteration is a classical (deterministic) one;
- at each γ -iteration, it is possible to reuse the search tree developed so far and only realize the updating of the current plan that is required to cope with the newly added elementary effects. Thus it is not necessary to replan from scratch at each iteration.

4.3.2 POSPLAN*: a hierarchical planning algorithm

The planning algorithm POSPLAN* is an implementation of the increasing acceptability method. It can be considered as a hierarchical version of POSPLAN, where the different hierarchy levels correspond successively to each γ -acceptability threshold from $\gamma = \gamma_n$ to $\gamma = \gamma_i = \gamma_{max}$.

Since the planning algorithm NDP is only semi-decidable, the POSPLAN* algorithm does not necessarily stop. In that case we retain the current γ_{max} at the interruption time. In practice, we classically use a depth limit of the search, within the *planning* function.

In Figure (7)⁸ we denote by \mathcal{P}_{γ_i} the plan obtained from the plan \mathcal{P} by transforming each of its actions so as to incorporate the elementary consequences that have a possibility $\pi > 1 - \gamma_i$.

Algorithm POSPLAN* (Δ)

1. $\mathcal{P} = \mathcal{P}_{null}$; $\gamma_{max} = 0$; $\gamma_i = \gamma_n$;
2. while $\gamma_i \leq 1$ do:
 - $\mathcal{P} \leftarrow \mathcal{P}_{\gamma_i}$
 - $\Delta = \Delta_{1-\gamma_i}$
 - $\mathcal{P} \leftarrow \text{planning}(\mathcal{P})$
 - if $\gamma_{max} < \gamma_i$ then $\gamma_{max} = \gamma_i$
 - $\gamma_i \leftarrow \gamma_{i-1}$

Figure 7: POSPLAN*: general algorithm

It is interesting to compare this acceptability hierarchy with the principles of abstraction hierarchies that have been developed recently [34, 35, 36]. The main results are summarized in Figure (8).

4.3.3 Completeness of POSPLAN*

In abstraction hierarchies, the completeness of the method is trivial since the abstraction levels are only used to express preferences concerning the subgoal order. In POSPLAN*, the completeness result is more subtle, as we have to show that not considering potential establishers at a given node will not prevent a solution plan from being present in the search tree. The following result establishes the completeness of POSPLAN*.

Proposition 6 (Completeness of POSPLAN*)

Let $\Delta = \langle \pi_{init}, \epsilon_{Goals}, \mathcal{A} \rangle$ be a possibilistic planning problem, and γ_{max} be the greater γ_i such that there exists an essential solution plan $\langle a_i \rangle_{i=0}^{N-1}$ with $N[Goals | \pi_{init}, \langle a_i \rangle_{i=0}^{N-1}] =$

⁸In the version presented here, γ_i is locally changed into γ_{i-1} when the current plan is a solution for Δ_i . Another alternative would consist in changing globally from γ_i to γ_{i-1} all partial plans of the search tree whenever a solution plan for Δ_i has been found.

abstraction hierarchy	acceptability hierarchy
Precondition propositions are partitioned into abstraction levels. Within a given abstraction level only a subset of the subgoals are considered.	All precondition propositions are considered at each iteration.
For a given subgoal, the set of potential establishers is independent of the current abstraction level.	Elementary consequences are partitioned into acceptability levels. Within a given acceptability level, only a subset of the potential establishers of a given subgoal are considered.
Abstraction hierarchies implement a specific search control heuristic concerning the choice of the flaw to consider.	Acceptability hierarchies implement a specific search control heuristic concerning both the choice of the flaw and the choice of the way this flaw will be removed (lines 1 and 2(a) in the refinement function).

Figure 8: *Comparison of hierarchy principles*

γ_i . Then POSPLAN* will generate a partial solution plan \mathcal{P} with $N[\text{Goals}|\pi_{init}, \mathcal{P}] = \gamma_{max}$.

Proof: The following lemma insures that for each i , if there exists a safe plan for $\Delta_{1-\gamma_i}$, then POSPLAN* will find a solution at the γ_i level. Then from the equivalence proposition (2), the proposition (6) is proved.

Lemma 2

$\forall i$, if $\langle a_i \rangle_{i=0}^{N-1}$ is an essential safe plan for $\Delta_{1-\gamma_i}$, then POSPLAN* will generate a partial plan \mathcal{P} essential and safe for $\Delta_{1-\gamma_i}$, such that $\langle a_i \rangle_{i=0}^{N-1}$ is a completion of \mathcal{P} .

proof: By induction on i . For $i = n$ and $\gamma_i = \gamma_n$, this is a consequence of POSPLAN algorithm's completeness (proposition (5)). Assume the proposition true till $i + 1$. If $\langle a_i \rangle_{i=0}^{N-1}$ is an essential safe plan for $\Delta_{1-\gamma_i}$, it is also a safe plan for $\Delta_{1-\gamma_{i+1}}$, but not necessarily essential. Let $\langle a_j \rangle_{j=0}^{p-1}$, $p \leq N$ be a subplan of $\langle a_i \rangle_{i=0}^{N-1}$, essential and safe for $\Delta_{1-\gamma_{i+1}}$. By induction we know that there exists a partial solution plan \mathcal{P} essential and safe for $\Delta_{1-\gamma_{i+1}}$, such that $\langle a_j \rangle_{j=0}^{p-1}$ is a completion of \mathcal{P} . Then from the lemma 1 we deduce that POSPLAN* can generate a partial safe plan \mathcal{P}' for $\Delta_{1-\gamma_i}$, such that $\langle a_i \rangle_{i=0}^{N-1}$ is a completion of \mathcal{P} . \square

Example: We still consider the possibilistic planning problem $\Delta = \langle \pi_{init} = \{(1, \bar{f} \wedge \bar{g} \wedge \bar{s} \wedge \bar{p} \wedge \bar{y})\}, \epsilon_{Goals} = y, \mathcal{A} \rangle$. $\mathcal{A} = \{sow-normal, sow-better, treat, harvest\}$ is an 8-level action set, with $\gamma_i \in \{0.2, 0.3, 0.6, 0.7, 0.8, 0.9, 1.0\}$. At $\gamma_6 = 0.2$, $\mathcal{A}_{0.8}$ is deterministic and POSPLAN*(Δ) generates the safe plan $\langle sow-normal, harvest \rangle$. At

$\gamma_5 = 0.3$, $\langle \textit{sow-normal}, \textit{harvest} \rangle$ is still a safe plan but at $\gamma_4 = 0.6$, due to the fact that $\textit{sow-normal}_{0.4}$ has two possible effects for its second discriminant POSPLAN* must generate a new safe plan which is $\langle \textit{sow-better}, \textit{treat}, \textit{harvest} \rangle$. This plan remains safe for γ_6 and γ_5 but does not for $\gamma_3 = 0.7$ because the second discriminant of $\textit{sow-better}_{0.3}$ is then associated with two possible effects. Since no other plan can be generated, $\gamma_{max} = 0.6$ and the returned plan is $\langle \textit{sow-better}, \textit{treat}, \textit{harvest} \rangle$.

5 Conclusions

5.1 Contributions

The main goal of this paper was to formalize a possibilistic approach of planning under uncertainty in domain models characterized by incomplete knowledge of the initial state and by actions having both context-dependent and graded nondeterministic effects. It was inspired by the work done on the BURIDAN[1] planner that relies on a probabilistic representation of uncertainty. In practice, it seems more natural and easier to see actions in terms of normal and more or less exceptional effects rather than probable ones. Thus, the ordinal nature of a possibilistic representation of uncertainty is quite appealing for this purpose. Besides, its representational adequacy, the possibilistic approach has interesting computational properties since the search for γ -acceptable or optimally safe plans amounts to solve induced planning problems that have only crisp nondeterministic actions (i.e. each action having then only normal effects). Moreover, in the case of optimal plan generation, the proposed sound and complete POSPLAN* algorithm is an anytime least-commitment planner; it possesses the additional feature of iteratively solving derived planning problems that are progressively more complex and exploits at each iteration the partial plans developed in the previous ones. The POSPLAN* planner and, consequently, the NDP and POSPLAN algorithms have been implemented in Common Lisp reusing part of BURIDAN's code (in particular its SNLP basis).

The generation of γ -acceptable and optimally safe plans relies on the NDP algorithm which is a planner operating on pure nondeterministic actions. Thus, our approach gives practical usefulness to such a kind of planner that so far were only of theoretical interest [5] [6].

Several recent planning approaches are based on Bayesian decision theory (probabilistic uncertainty and additive utility functions)[20]. Now, possibility theory offers a well-suited base for a more qualitative version of decision theory [24]; thus, our approach can be seen as a preliminary step toward a more qualitative approach to decision-theoretic planning.

5.2 Limitations and future work

Although our approach enables an extension of the STRIPS representation by supporting graded nondeterministic actions and partially known initial state, it is still far from being able to cope with the complexity of most practical problems. One of

the main limitations comes from the syntactic restrictions that the propositional nature of the action representation language imposes. The other limitations concern the capabilities of hierarchical planning (the possibility to specify and search plans at various levels of detail), of dealing with time constraints (duration) and of handling constraints on resources. Important progresses have been accomplished on these last two issues over the past ten years. It would be interesting to reconsider them in the context of graded nondeterministic actions examined in this paper. Especially appealing is the development within the framework of possibility theory of an homogeneous treatment of uncertainty on initial state and effects of actions together with imprecision and uncertainty on durations and resource consumptions and productions.

As pointed out in Section 2.4, the possibilistic transition functions representing the uncertain effects of actions may be generated automatically from an initial set of hard and default rules. This is left for further research.

Our approach of planning under uncertainty assumes a complete non-observability of states during execution. In the opposite case of complete observability, one may switch to dynamic programming approaches that have been extended to possibilistic transition functions (see, for instance, [17]). For the intermediate case of partial observability, we hope to extend our system so that information-gathering actions can also be incorporated in the planning process as has been done with C-Buridan[37] for probabilistic planning.

In this paper, a goal is simply a conjunction of literals that defines a set of equally good goal states (i.e. reaching any of them would be fine). We are thinking of extending our framework so as to take account of preferences over the goal states, using a qualitative utility function or equivalently a fuzzy set of goal states⁹. More generally, the approach could be reconsidered to cope with a more elaborate notion of goals in the spirit of utility functions in decision theoretic planning approaches such that a preference would be associated to each state that might be gone through in the execution of a sequence of actions. The objective would then be to find the best sequence with respect to a criteria embedding a compromise between uncertainty and preference. This notion of compromise has been formalized in the setting of a qualitative possibilistic decision theory by Dubois and Prade[24] and requires a commensurateness assumption. Another extension would consist in using additive utility instead of the above mentioned preferences that correspond to a notion of qualitative utility. Planning would then requires to compare fuzzy numbers representing fuzzy utilities.

It is our belief that, although much work remains to be done towards practical systems of planning under uncertainty, our approach provides an original and profitable basis for future developments.

⁹Possibility distributions would then be used to represent both *uncertainty* as explained in this paper and *preference*.

Acknowledgements

We are indebted to Henri Prade and Alessandro Saffiotti for helpful comments on a previous version of this paper.

References

- [1] Kushmerick N., Hanks S. and Weld D.W. An algorithm for probabilistic planning. *Artificial Intelligence* 76: 239-286 (1995).
- [2] Fikes R.E. and Nilsson N.J. STRIPS: a new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 2: 189-208 (1971).
- [3] McAllester D. and Rosenblitt D. Systematic nonlinear planning. *Proc. of AAAI'91*, 634-639 (1991).
- [4] Dubois D. and Prade H. Possibility theory - an approach to computerized processing of uncertainty, Plenum Press, 1988.
- [5] Pednault E. Synthesizing plans that contain actions with context-dependent effects. *Computational Intelligence*, 4(4):356-372, (1988).
- [6] Pednault E. Generalizing nonlinear planning to handle complex goals and actions with context-dependent effects. *Proc of IJCAI'91*, 240-245, (1991).
- [7] Garcia F. Planning actions with context-dependent effects: a state-space based approach. *Current Trends in AI Planning*, C. Backstrom and E. Sandewall (Eds.), IOS Press, 1994.
- [8] Penberthy J.S. and Weld D. UCPOP: A sound, complete, partial-order planner for ADL. *Proc. of KR'92*, 103-114 (1992).
- [9] Brewka G. and Hertzberg J. How to do things with worlds: on formalizing actions and plans. *J. Logic and Computation*, 3(5):517-532 (1993).
- [10] C. Boutilier, R. Dearden, M. Goldszmidt. Exploiting structure in policy construction, *Proc. of IJCAI'95*.
- [11] Poole D. A framework for decision-theoretic planning I: combining the situation calculus, conditional plans, probability and utility, *Proc. of UAI'96* 436-445.
- [12] Geffner H. and Wainer J. A model for actions and contingent plans. *Proc. of AAAI'97*.
- [13] Goldman R. and Boddy M. Epsilon-safe Planning. *Proc. of UAI'94*, 253-261 (1994).

- [14] De Glas M. and Jacopin E. An Algebraic Framework for Uncertain STRIPS Planning. Proc. of AIPS'94, 62-68 (1994).
- [15] Thiébaux S., Hertzberg J., Shoaff W., and Schneider M. A stochastic model of actions and plans for anytime planning under uncertainty. *Int. J. Intelligent Systems*, 10(2):155-183, 1995.
- [16] Bonet B., Geffner H. Arguing for decisions: a qualitative model for decision making, Proc. of UAI'96. (1996).
- [17] Fargier H., Lang J. and Sabbadin R. 1996. Towards Qualitative Approaches to Multistage Decision Making. Proc. of IPMU'96, Granada (Spain), (1996).
- [18] Haddawy P., Hanks P. Issues in decision-theoretic planning: utility functions for deadline goals. Proc. of KR'92, 71-82 (1992).
- [19] Dean T. Decision-Theoretic Planning and Markov Decision Processes. Technical report, Brown University, 1994.
- [20] Boutilier C., Dean T. and Hanks S. Planning under uncertainty: structural assumptions and computational leverage, *New Directions in AI Planning*, M. Ghallab and A. Milani (Eds.), IOS Press, 157-171, 1996
- [21] Cassandra A.R., Kaelbling L.P. and Littman M.L. Acting optimally in partially observable stochastic domains. *Proc. of AAAI'94*, 1023-1028.
- [22] C. Boutilier, D. Poole. Computing optimal policies for partially observable decision processes using compact representations, *Proc. of AAAI'96*, 1168-1175.
- [23] Zadeh L.A. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and Systems*, 1: 3-28 (1978).
- [24] Dubois D. and Prade H. 1995. Possibility theory as a basis for qualitative decision theory. Proc. of IJCAI'95, Montreal, 19-25 (1995).
- [25] Yager R. An approach to ordinal decision making, *International Journal of Approximate Reasoning* 12, 237-261 (1995).
- [26] Dubois D. and Prade H. Possibilistic mixtures and their applications to qualitative utility theory – Part II: Decision under incomplete knowledge. *Proc. of the Conf. on Foundations and Applications of Possibility Theory (FAPT'95)* (G. de Cooman et al. eds.), Ghent, 1995, World Scientific, 256-266.
- [27] Boutilier C. Towards a logic for qualitative decision theory, Proc. of KR'94. (1994)

- [28] Kacprzyk J. Multi-Stage decision making under fuzziness, Verlag TÜV Rheinland, 1983.
- [29] Saffiotti A., Konolige K. and Ruspini E.H. A multivalued logic approach to integrating planning and control. *Artificial Intelligence* 76 (1995), 481-526.
- [30] Chapman D. Planning for conjunctive goals. *Artificial Intelligence*, 32(3):333-377 (1987).
- [31] Kambhampati S. and Nau D.S., On the nature and Role of Modal Truth Criteria in Planning. *Artificial Intelligence*, (1996).
- [32] Dean T. and Boddy M. Reasoning about partially ordered events. *Artificial Intelligence*, 36(3):375-399 (1988).
- [33] Kambhampati S. Multi-contributor causal structures for planning: a formalization and evaluation. *Artificial Intelligence* 69:235-278 (1994).
- [34] Yang Q. and Tenenber J.D. ABTWEAK: Abstracting a nonlinear least commitment planner. *Proc. AAAI'90*, 204-209 (1990).
- [35] Knoblock C., Automatically generating abstractions for planning. *Artificial Intelligence*, 68:243-302 (1994).
- [36] Garcia F. and Laborie P. Hierarchisation of the Search Space in Temporal Planning. *New Directions in AI Planning*, M. Ghallab and A. Milani (Eds.), IOS Press, 1996.
- [37] Draper D., Hanks S. and Weld D. Probabilistic planning with information gathering and contingent execution. *Proc. of AIPS'94*, 31-36 (1994).