Unifying Nondeterministic and Probabilistic Planning through Imprecise Markov Decision Processes *

Felipe Werndl Trevizan**[1], Leliane Nunes de Barros[1], and Fabio Gagliardi Cozman[2]

[1] Instituto de Matemática e Estatística, Universidade de São Paulo
Rua do Matão, 1010, Cidade Universitária – 05508-090 São Paulo, SP, Brasil
{trevisan, leliane}@ime.usp.br
[2] Escola Politécnica, Universidade de São Paulo
Av. Prof. Mello Moraes, 2231, Cidade Universitária – 05508-900 São Paulo, SP, Brasil
fgcozman@usp.br

Abstract. This paper proposes an unifying formulation for nondeterministic and probabilistic planning. These two strands of AI planning have followed different strategies: while nondeterministic planning usually looks for min-max (or worst-case) policies, probabilistic planning attempts to maximize expected reward. In this paper we show that both problems are special cases of a more general approach, and we demonstrate that the resulting structures are Markov Decision Processes with Imprecise Probabilities (MDPIPs). We also show how existing algorithms for MDPIPs can be adapted to nondeterministic and probabilistic planning in AI.

1 Introduction

Planning is not only ubiquitous in artificial intelligence; it also appears in many different forms. While "classical" planning focuses on deterministic settings without any uncertainty, several "non-classical" approaches have tried to deal with various forms of uncertainty [1]. Among these approaches, *probabilistic planning* has produced significant results in recent years [2,3,4]. Here one finds that probabilities are used to encode risk and uncertainty, while expected utility is used to rank plans. Another important approach is *nondeterministic planning* [5], where one does not even assign probabilities to the consequences of actions.

A particularly apt perspective from which to read this literature is due to Geffner and Bonet [6]. The idea is to make a clear distinction between planning languages, models and algorithms, and to try to capture what is common across approaches by formulating general languages, models, and algorithms. As discussed in Section 2, this perspective has been quite effective in unifying

* Project funded by Fundação de Amparo a Pesquisa do Estado de São Paulo (FAPESP)
** Funded by Conselho Nacional de Pesquisa (CNPq)

various strands of planning, from classical to probabilistic, including variants of non-deterministic planning. A unified understanding of planning problems is obviously beneficial not only to artificial intelligence but to several other fields such as operations research and management.

The just mentioned general formulation takes probabilistic and nondeterministic approaches as two extreme and unrelated positions concerning planning. And in fact there are apparently great differences between them. They are based on different assumptions concerning beliefs (either always translated into probabilities, or never translated into probabilities), and different prescriptions for action (either focused on average behavior through expected utility, or on worst-case guarantees coming from min-max). Accordingly, communities in probabilistic and nondeterministic planning have had little real interaction. In a sense, this is the general decision-theoretic contrast between Bayesian position that prescribes expected utility, and a min-max position that looks at worst case behavior. But in decision theory there are many other options, and in particular there are interesting options that can handle not only expected and min-max positions, but also other positions in between. Thus one can have a decision problem where some events have probability values attached to them, while other events may be associated with "nondeterministic" phenomena.

In this paper we propose a unifying formulation for planning problems, where we can smoothly transition between probabilistic and nondeterministic planning. These two approaches are viewed as simple special cases, and our analysis reveals a spectrum of new planning problems that has not been considered by the literature in artificial intelligence so far. We demonstrate that the resulting structures are Markov Decision Processes with Imprecise Probabilities (MDPIPs), a model proposed in operations research to solve control problems. We also show how existing algorithms for MDPIPs can be adapted to nondeterministic and probabilistic planning in AI.

The remainder of this paper is organized as follows. In Section 2 we summarize Geffner and Bonet's unifying perspective on planning — thus defining the probabilistic and nondeterministic varieties. Section 3 introduces basic concepts underlying risk and uncertainty. Section 4 defines our proposal model for probabilistic and nondeterministic planning, named **NDP model**. In Section 5 we demonstrate that the NDP model is a variant of Markov Decision Processes with Imprecise Probabilities (called **MDPIPs** in the literature). Section 6 adapts MDPIP algorithms for NDP models. Finally, Section 7 suggests extensions on PPDDL so as to represent NDP models, and in Section 8 we draw some conclusions.

2 Planning Models

We briefly review the mathematical models needed to characterize planning tasks with full observability for different action dynamics (we simplify the presentation by assuming full observability; partial observability can be addressed with minor changes in the framework). Every state model that we consider can be defined in terms of the elements of the following basic state model [7]:

BSM1 a discrete and finite state space $\mathcal{S}$,
BSM2 a non-empty set of initial states $S_0 \subseteq \mathcal{S}$,
BSM3 a goal given by a non-empty set $S_G \subseteq \mathcal{S}$,
BSM4 a non-empty set of actions $\mathcal{A}(s) \subseteq \mathcal{A}$ representing the actions applicable in each state s,
BSM5 a state transition function $F(s,a) \subseteq \mathcal{S}$ mapping states s and actions $a \in \mathcal{A}(s)$ into sets of states, i.e. $\|F(s,a)\| \geq 1$, and
BSM6 a positive action cost $C(a,s)$ for doing $a \in \mathcal{A}(s)$ in s.

Differents models can be defined adding new restrictions or modifing the statements 2, 5 and 6. Those models are:

- **Deterministic Models**, where the dynamics are defined by a deterministic state transition function, i.e., $\|F(s,a)\| = 1$. This is the basis of the *classical* planning scenario, where one has additional constraints of initial state $\|S_0\| = 1$ and goal $\mathcal{C}(a,s) = 1 \quad \forall s \in \mathcal{S}, a \in \mathcal{A}(s)$. The goal is generally to find a plan (sequence of actions) that moves from the initial state to the goal.
- **Nondeterministic Models**, where the actions may result in more than one successor state without preferences among them. So we have the same model as in deterministic planning, but uncertainty in actions. In fact, the term “nondeterminism” should here be understood as “automata-style” nondeterminism. This may be a bit confusing as nondeterminism is usually associated with probabilities in decision theory — using the terminology discussed in Section 3, we actually have *planning under pure Knightian uncertainty*. The goal is generally to find a plan that moves from the initial state to the goal *no matter what* nondeterministic actions do — that is, that works even in the worst-case.
- **Probabilistic Models**, where actions have probabilistic consequences. Not only the function $\|F(s,a)\| \geq 1$ is given, but also the model includes a probability distribution $P(\cdot|s,a)$ over $F(s,a) \quad \forall s \in \mathcal{S}, a \in \mathcal{A}(s)$. The goal is to maximize expected utility; there are in fact several ways to take into account costs and rewards. The important point is that the model prescribes a *unique* probability distribution; thus each policy is associated with a *unique* value, and selecting a policy means finding a policy such that no other policy has a larger value.

There are algorithms that compute policies for each one of these problems. A recent development is the derivation of a single algorithm that can be instantiated for different models, including the ones just described [2]. However it should be emphasized that this generalized formulation does not yield a smooth family of solutions that moves from one case to the other. In particular, there is no algorithm that has the probabilistic and nondeterministic cases as special ones, and also that copes with mixtures of these cases. The main goal of this paper is to start the construction of such a framework.

3 Risk, Knightian Uncertainty and Sets of Probabillies

Instead of moving directly to our general formulation, it is instructive to start with an open-minded review of decision theory. Here a decision maker contemplates a set of options (in our setting, policies); each option yields a utility depending on the state of nature that obtains [8]. We consider a set of states of nature Ω; then each option is a function $f : \Omega \to \Re$. If a decision maker can specify a single probability measure P over a field of events defined on Ω, then this "Bayesian agent" will evaluate each option f by expected utility, $E_P[f]$. Typically one assumes that such an agent can select any option that is be dominated by expected utility — a simple criterion that leads to a rich theory [9].

However, there may be situations where an agent does not have a single probability measure. A common assumption then is that the agent will have *no* probability at all. The usual solution then is to look at worst-case scenarios: select f that displays the highest worst utility — a minimax solution [8]. The difference between these extremes (one/no probability) is well studied in economics and psychology. Usually the presence of probabilities is associated with the expression *risk*, while the absence of probabilities is associated with *uncertainty*, or rather, *Knightian uncertainty* (from the work of Knight [10]). To indicate the pervasiveness of these concepts in economics practice, it suffices to quote from a relevant speech by Alan Greenspan, read in January 3 2004:

> ...uncertainty is not just a pervasive feature of the monetary policy landscape; it is the defining characteristic of that landscape. The term "uncertainty" is meant here to encompass both "Knightian uncertainty," in which the probability distribution of outcomes is unknown, and "risk," in which uncertainty of outcomes is delimited by a known probability distribution...

Now it is clear that *sequential decision making under risk* is *probabilistic planning*, while *sequential decision making under Knightian uncertainy* is *nondeterministic planning*. In fact, we would like to suggest that the term "nondeterministic" is an unfortunate one in the present setting, as nondeterminism usually suggests some form of probabilistic model. It seems that Knightian uncertainty, although longer, is a less overloaded term.

Once it is recognized that risk and Knightian uncertainty are two challenges a decision maker may face, one is naturally lead to ask about situations of both risk *and* Knightian uncertainty. That is, we may consider the possibility that an agent displays imprecision in probability values or even that the agent considers a set of probability values. There are many reasons where such a general situation may arise. First, it may happen that existing beliefs are incomplete or vague [11,12,13], either because there is no time/resources to spend in their elicitation, or because experts are psychologically unable to specify precise probability values. Second, it may be the case that a group of experts disagrees on probability values, and no compromise can be reached other than the collection of their opinions [14,15]. Another reason to abandon a single probability measure is when one is interested in the robustness of inferences — that is, in evaluating how much inferences can change when probability values are allowed to vary [9,16,17].

Our strategy in this paper is, at a fundamental level, simple: we intend to bring the decision theory of risk *and* Knightian uncertainty to the realm of artificial intelligence planning. In this setting, uncertainty will be represented by *sets of probability measures.* At one extreme, we obtain probabilistic planning (all sets are singletons); at the other extreme, we obtain nondeterministic planning (all sets are as large as possible). Moreover, we obtain a continuum of models as we allow sets of probability measures to transit from vacuously large ones to singletons.

Artificial intelligence has witnessed steady interest in sets of probability measures, for example, in the theory of probabilistic logic [18,19,20], in Dempster-Shafer theory [21,22], in theories of argumentation [23,24], and in techniques that generalize graph-theoretic models such as Bayesian networks [25,26,27].[3] Our contribution here is to identify the probabilistic/nondeterministic planning spectrum with the theory of sets of probability measures.

4 Planning under Risk and Knightian Uncertainty: the Probabilistic/Nondeterministic Spectrum

The *Nondeterministic-Probabilistic model*, named **NDP model**, can be seen as a more general model since it gives a precise semantics to a new planning task, i.e., a task envolving nondeterministic and probabilistic effects of actions. Therefore, a planning problem can be solved considering, simultaneously, these two types of action's effects.

Example 1. Given the following situation, ... In this case the action ? can be envisaged as a nondeterministic-probabilistic action, where ....

Once the NDP model has to represent non-deterministic effects, the transition function $F(s, a)$, from the basic state model described in Section 2, instead of being defined as $F(s, a) \subseteq \mathcal{S}$, will map states and actions to sets of sets of the state space. That is, for all k in $F(s, a)$, k is a subset of or equal to $\mathcal{S}$.

Definition 1. *In the NDP model, the transition function $F(s, a)$ maps states s and actions $a \in \mathcal{A}(s)$ into non-empty sets of the parts of the state space, i.e. $F(s, a) \subseteq 2^{\mathcal{S}}$.*

Definition 2. *A* ***possible-state set*** *k is a set composed of possible resulting states achieved with the execution of an action a, i.e. $k \in F(s, a)$ with $F(s, a)$ being the state transition of Definition 1.*

With the above definions the probabilistic function $P(k|s, a), \quad k \in F(a, s)$ has a diffrent interpretation: in the NDP model, $P(k|s, a)$ represents the probability of the next state be one of the states in k. A complete and formal description of the NDP model is given by:

[3] There is now significant literature on the theory and applications of sets of probability measures [28,29,30,31].

NDP1 a discrete and finite state space $\mathcal{S}$,
NDP2 a non-empty set of initial states $S_0 \subseteq \mathcal{S}$,
NDP3 a goal situtation given by a non-empty set $S_G \subseteq \mathcal{S}$,
NDP4 a non-empty set of actions $\mathcal{A}(s) \subseteq \mathcal{A}$ representing the actions applicable in each state s,
NDP5 a state transition function $F(s,a) \subseteq 2^{\mathcal{S}}$ mapping states s and actions $a \in \mathcal{A}(s)$ into non-empty sets of the parts of the state space,
NDP6 a probability distribution $P(\cdot|s,a)$ over $F(s,a) \quad \forall s \in \mathcal{S}, a \in \mathcal{A}(s)$ where $P(k|s,a)$ represents the probability of choosing the possible-state set $k \subseteq \mathcal{S}$ when action a is applied in state s, and
NDP7 a positive action cost $C(a,s)$ for doing $a \in \mathcal{A}(s)$ in s.

Notice that there is two types of choices in the NDP model: a probabilistic choice of a possible-state set and a nondeterministic choice of a sucessor state from the possible-state set. As it was shown in the Example 1, the nondeterministic-probabilistic planning task can be characterized by domains for which the action dynamics satisfies the following restrictions: (1) $\|F(s,a)\| > 1$ and (2) $\exists k \in F(s,a)\text{s.t.}\|k\| > 1$, for $s \in \mathcal{S}, a \in \mathcal{A}(s)$. If none of these requirements is true, then NDP model is reduced to one of the models described in Section 2.

If the first requirement is false, i.e. $\|F(s,a)\| = 1$, and the second is true, the NDP model is equivalent to the Non-deterministic model (Section 2). This is because: $\forall s \in \mathcal{S}$, $a \in \mathcal{A}(s)$, if $\quad \|F(s,a)\| = 1$ then $P(k \in F(s,a)|s,a) = 1$, which means that the choice of a possible state set will be deterministic while the occurrence of a single state $s' \in k$ will be nondeterministic.

For the planning set where the first requirement is true and the second is false, then the model corresponds to the Probabilistic Model from Section 2. This is due to the fact that $\forall s \in \mathcal{S}, a \in \mathcal{A}(s), k \in F(s,a) \quad \|k\| = 1$, implying that there will be only one candidate to the nondeterministic choice, with probability $P(k|s,a)$ after executing a in the state s. Under this assumptions the probability distribuition over $2^{\mathcal{S}}$ is equivalent to a probability distribuition over $\mathcal{S}$.

Finally, when both requirements are false, the model is equivalent to the Deterministic Model once there is not any point of choice: neither in the probabilistic choice of a possible state set nor in the nondeterministic choice of a sucessor state.

Furthermore, the complete NDP model is equivalent to a Markov Decision Process having imprecisely known transition probabilities. This equivalence, proved in the next section, gives a formal semantics for the NDP model.

5 The relation between NDP and MDPIP model

Markov Decision Processes with Imprecise Probabilities (MDPIPs) [32,33] are an extension of Markov Decision Processes (MDPs) [34] where the probabilities describing the transition between states are not defined as a number, but as a finite set of linear inequalities. Consequently, the possible effects of an action are modelled by a credal set $\mathcal{K}$ (cite ?) over the state space instead of a probability distribution over the same space. A precise definition of an MDPIP is:

MIP1 a discrete and finite state space $\mathcal{S}$,
MIP2 a goal situtation given by a non-empty set $S_G \subseteq \mathcal{S}$,
MIP3 a non-empty set of actions $\mathcal{A}(s) \subseteq \mathcal{A}$ representing the actions applicable in each state s,
MIP4 a non-empty credal set $\mathcal{K}_s(a)$ representing the possibles probability distributions $P(\cdot|a, s)$ over $\mathcal{S}$, and
MIP5 a positive action cost $C(a, s)$ for doing $a \in \mathcal{A}(s)$ in s.

The formulation above is based on Game Theory and considers the existence of a mechanism that selects the exact probability distribution after an action has been selected. This mechanism is usually called nature and an MDPIP can be solved only if an assumption is made about its behavior. In this paper, we assume that nature is intent on maximizing the expected total discount cost for each state that the planner wishes to minimize (1). Therefore, a min-max criterion is adopted to find a policy.

$$V(s) = \min_{a \in \mathcal{A}(s)} \max_{P(\cdot|s,a) \in \mathcal{K}_s(a)} \{C(a, s) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a)V(s')\} \tag{1}$$

In [33] it has been shown that the solution to (1), called $V^*(s)$, exists and is the only one. It is also proved that the optiomal policy for an MDPIP can be expressed by a stationary policy, i.e., the same policy for any instant in time.

Proposition 1. *The NDP model is a special case of the MDPIPs model.*

Proof. Note that NDP1, NDP3, NDP4 and NDP7 are equal, respectively, to MIP1, MIP2, MIP3 and MIP5. Thus the proof is reduced to prove that NDP5 and NDP6 implies in MIP4.

Suppose, without lost of generality, that $P(k = \{s_i, \ldots, s_j\}|s, a) = p$ for a $s \in \mathcal{S}, a \in \mathcal{A}(s)$, and $k \in F(s, a)$. Therefore the following inequations hold:

$$0 \leq P(s_l|s, a) \leq p \quad \forall s_l \in \{s_i, \ldots, s_j\} \tag{2}$$

$$0 \leq \textstyle\sum_{s_l \in \{s_i, \ldots, s_j\}} P(s_l|s, a) \leq p \tag{3}$$

The set of inequations (2) and (3) describe the credal set $\mathcal{K}_s(a)$ of MIP4.

Propositon 1 not only makes the results in [32,33] valid for the NDP model, but also suggests algorithms to solve it. An example of algorithm from Operational Research used to solve a problem modelled by a NDP is given in the next section.

6 Algorithms

7 Languages

The formulation of the planning task as a state model, or a Markov Decision Process in the case of probabilistic planning, has opened up the field to try new approaches, p.e., from the operational research field.

A probabilistic planning problem can be specified through probabilistic effects of actions. The 4th International Planning Competition (IPC-4), happen in 2004, included a track on probabilistic planning for the first time. The results of the competition included an interactive system for planner evaluation, a set of benchmark problems, and some important learned lessons. It also resulted on a specification language for probabilistic planning (PPDDL), an extension of PDDL (the standard language for specifying planning domains and problems), [35].

A probabilistic effect declares an exhaustive set of probability-weighted outcomes. For isntance, in PPDDL the syntax for probabilistic effects is:

$$(\texttt{probabilistic}\, p_1 e_1 \ldots p_i e_i \ldots p_k e_k)$$

meaning that effect e_i occurs with probability p_i such as $pi \geq 0$ and $\sum_{i=1}^{k} pi = 1$. The PPDDL language incorporates syntactic extensions to PDDL that allows to specify Markov Decision Processes (MDPs).

The 5th International Planning Competition (IPC-5), generalized the probabilistic planning track to include separate subtracks for nondeterministic planning. Therefore, some extensions to PPDDL, required to model non-deterministic effects, added non-deterministic statement of the form:

$$(\texttt{oneof}\ e_1 \ldots e_i \ldots e_k)$$

where e_i are PPDDL effects. The semantics is that when executing such effect, one of the e_i is chosen and applied to the current state.

However, there is a class of planning problems that has not been considered in the IPC-5 or by the planning community: the nondeterministic and probabilistic planning. While there are significant efforts given to both research fields: probabilistic planning [2,3,4] and nondeterministic planning [5], there is not been yet a proposal of model, language and algorithms for this joint formulation of planning.

Falar rapidamente sobre a extensão de PPDDL:

$$(\texttt{probabilistic}\, p_1\ (\texttt{oneof}\ e_1^1 \ldots e_1^{j_1})\ \ldots p_k\ (\texttt{oneof}\ e_k^1 \ldots e_k^{j_k}))$$

Os exemplos na introdução seriam 'syntatic sugar' para a notação acima.

8 Conclusions and future works

Falar sobre o artigo do Thomas Dean e Robert Givan [36] onde é considerado um caso especial de MDPIP (a probabilidade de transição para cada estado é dada por um intervalo). O nosso é mais geral e ainda permite expandir o modelo para tratar restrições como $P(s1|s,a) > 2 * P(s2|s,a)$ mantendo a proposição 1 válida.

Falar sobre o artigo do Harmanec [37] para uma outra visão do problema de MDPIP.

Trabalho futuro, adaptar outros algoritmos de PO e Planejamento (RTDP e LRTDP) para NDP.

References

1. Ghallab, M., Nau, D., Traverso, P.: Automated Planning: Theory & Practice. Morgan Kaufman (2004)
2. Bonet, B., Geffner, H.: Learning Depth-First Search: A unified approach to heuristic search in deterministic and non-deterministic settings, and its application to MDPs. In: Proc. 4th International Conf. on Artificial Intelligence Planning and Scheduling, To Appear (2006)
3. Bonet, B., Geffner, H.: Labeled RTDP: Improving the convergence of real-time dynamic programming. In Giunchiglia, E., Muscettola, N., Nau, D., eds.: Proc. 13th International Conf. on Automated Planning and Scheduling, Trento, Italy, AAAI Press (2003) 12–21
4. Guestrin, C., Koller, D., Parr, R., Venkataraman, S.: Efficient solution algorithms for factored MDPs. J. Artif. Intell. Res. (JAIR) **19** (2003) 399–468
5. Bertoli, P., Cimatti, A., Roveri, M., Traverso, P.: Planning in nondeterministic domains under partial observability via symbolic model checking. In: IJCAI. (2001) 473–478
6. Geffner, H.: Perspectives on artificial intelligence planning. In: Eighteenth national conference on Artificial intelligence, Menlo Park, CA, USA, American Association for Artificial Intelligence (2002) 1013–1023
7. Bonet, B., Geffner, H.: Planning with incomplete information as heuristic search in belief space. In Chien, S., Kambhampati, S., Knoblock, C., eds.: Proc. 6th International Conf. on Artificial Intelligence Planning and Scheduling, Breckenridge, CO, AAAI Press (2000) 52–61
8. Luce, D., Raiffa, H.: Games and Decisions. Dover edition, Mineola, New York (1957)
9. Berger, J.O.: Statistical Decision Theory and Bayesian Analysis. Springer-Verlag (1985)
10. Knight, F.H.: Risk, Uncertainty, and Profit. Hart, Schaffner & Marx; Houghton Mifflin Company, Boston (1921)
11. Levi, I.: The Enterprise of Knowledge. MIT Press, Cambridge, Massachusetts (1980)
12. Walley, P.: Statistical Reasoning with Imprecise Probabilities. Chapman and Hall, London (1991)
13. Walley, P.: Measures of uncertainty in expert systems. Artificial Intelligence **83** (1996) 1–58
14. Seidenfeld, T., Kadane, J.B., Schervish, M.J.: On the shared preferences of two Bayesian decision makers. The Journal of Philosophy **86**(5) (1989) 225–244
15. Seidenfeld, T., Schervish, M.: Two perspectives on consensus for (Bayesian) inference and decisions. IEEE Transactions on Systems, Man and Cybernetics **20**(1) (1990) 318–325
16. Huber, P.J.: Robust Statistics. Wiley, New York (1980)
17. Kadane, J.B., ed.: Robustness of Bayesian Analyses. Volume 4 of Studies in Bayesian econometrics. Elsevier Science Pub. Co., New York (1984)
18. Frisch, A.M., Haddawy, P.: Anytime deduction for probabilistic logic. Artificial Intelligence **69** (1994) 93–122
19. Halpern, J.Y.: Reasoning about uncertainty. MIT Press, Cambridge, Massachusetts (2003)
20. Nilsson, N.J.: Probabilistic logic. Artificial Intelligence **28** (1986) 71–87
21. Shafer, G.: A Mathematical Theory of Evidence. Princeton University Press (1976)

22. Shenoy, P.P., Shafer, G.: Propagating belief functions with local computations. IEEE Expert **1**(3) (1986) 43–52
23. Anrig, B., Bissig, R., Haenni, R., Kohlas, J., Lehmann, N.: Probabilistic argumentation systems: Introduction to assumption-based modeling with ABEL. Technical Report 99-1, Institute of Informatics, University of Fribourg (1999)
24. Kohlas, J., Shenoy, P.P. In: Computation in Valuation Algebras. Kluwer, Dordrecht (2000) 5–40
25. Cozman, F.G.: Credal networks. Artificial Intelligence **120** (2000) 199–233
26. Cozman, F.G.: Graphical models for imprecise probabilities. International Journal of Approximate Reasoning **39**(2-3) (2005) 167–184
27. Fagiuoli, E., Zaffalon, M.: 2U: An exact interval propagation algorithm for polytrees with binary variables. Artificial Intelligence **106**(1) (1998) 77–107
28. de Cooman, G., Cozman, F., Moral, S., Walley, P., eds.: Proceedings of the First International Symposium on Imprecise Probabilities and Their Applications. Imprecise Probabilities Project, Universiteit Gent, Ghent, Belgium (1999)
29. de Cooman, G., Fine, T.L., Seidenfeld, T.: Proceedings of the Second International Symposium on Imprecise Probabilities and Their Applications. Shaker Publishing, The Netherlands (2001)
30. Bernard, J.M., Seidenfeld, T., Zaffalon, M., eds.: Proceedings of the Third International Symposium on Imprecise Probabilities and Their Applications. Carleton Scientific, Lugano, Switzerland (2003)
31. Cozman, F.G., Nau, R., Seidenfeld, T.: Proceedings of the Fourth International Symposium on Imprecise Probabilities and Their Applications. SIPTA (2005)
32. III, C.C.W., Eldeib, H.K.: Markov decision processes with imprecise transition probabilities. Operations Research **42**(4) (1994) 739–749
33. Satia, J.K., Jr, R.E.L.: Markovian decision processes with uncertain transition probabilities. Operations Research **21**(3) (1973) 728–740
34. Howard, R.A.: Dynamic Porgramming and Markov Processes. MIT Press (1960)
35. Younes, H.L.S., Littman, M.L., Weissman, D., Asmuth, J.: PPDDL1.0: An extension to PDDL for expressing planning domains with probabilistic effects. Journal of Artificial Intelligence Research (2005) COMPLETAR.
36. Givan, R., Leach, S.M., Dean, T.: Bounded-parameter markov decision processes. Artificial Intelligence **122**(1-2) (2000) 71–109
37. Harmanec, D.: A generalization of the concept of markov decision process to imprecise probabilities. In: ISIPTA. (1999) 175–182