

# Unifying Nondeterministic and Probabilistic Planning Through Imprecise Markov Decision Processes<sup>\*</sup>

Felipe W. Trevisan<sup>1,\*\*</sup>, Fábio G. Cozman<sup>2</sup>, and Leliane N. de Barros<sup>1</sup>

<sup>1</sup> Instituto de Matemática e Estatística, Universidade de São Paulo  
Rua do Matão, 1010, Cidade Universitária – 05508-090 São Paulo, SP, Brazil  
{trevisan, leliane}@ime.usp.br

<sup>2</sup> Escola Politécnica, Universidade de São Paulo  
Av. Prof. Mello Moraes, 2231, Cidade Universitária – 05508-900 São Paulo, SP, Brazil  
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 minimax (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 planning under uncertainty.

## 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]. 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 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 various strands of planning, from classical to probabilistic, including variants of nondeterministic 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.

---

\* Project funded by Fundação de Amparo a Pesquisa do Estado de São Paulo (FAPESP) process number 04/09568-0.

\*\* Funded by Conselho Nacional de Pesquisa (CNPq).

The just mentioned general formulation takes probabilistic and nondeterministic approaches as two extreme and unrelated positions concerning planning. 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 minimax). 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 minimax 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 minimax 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.

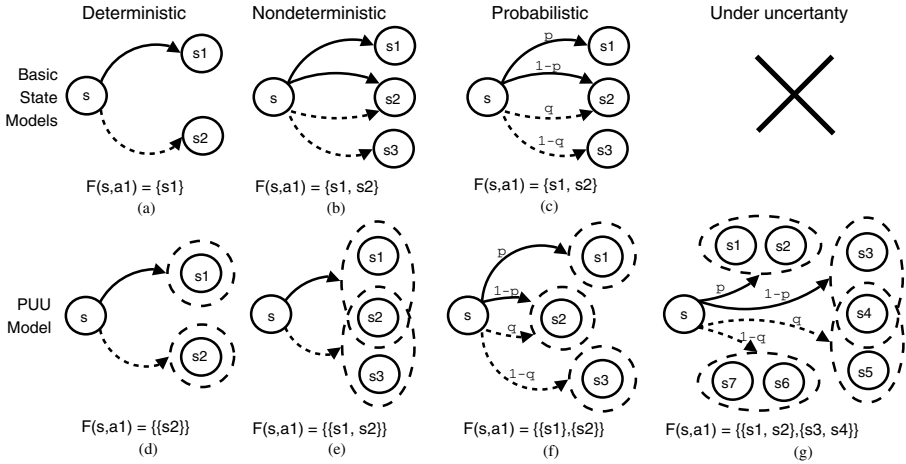
In this paper we propose a unifying formulation for planning problems, where we can smoothly transition between probabilistic and nondeterministic planning. These two approaches are viewed as simple special cases, and our analysis reveals a spectrum of new planning problems that has not been considered 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 planning under uncertainty.

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 planning under uncertainty (PUU), named **PUU model**. In Section 5 we demonstrate that the PUU model is a variant of Markov Decision Processes with Imprecise Probabilities (called **MDPIPs** in the literature). Section 6 adapts MDPIP algorithms for PUU models. Finally, in Section 7 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 (partial observability can be addressed with minor changes in the framework). Every state model that we consider can be defined in terms of the following basic state model [6]:

- 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 non-empty 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  $s$ .



**Fig. 1.** This figure illustrates the structure of transition function in the different models discussed in this paper. Solid and dotted arcs represents different actions. Solid circles are states and dashed circles indicates *possible-state sets*. Note that under uncertainty scenarios cannot be represented by the basic state model.

Differents models can be defined adding new restrictions or modifying the statements 2, 5 and 6. Those models, depicted by the first row of Fig. 1, are:

- **Deterministic Models** (Fig. 1 (a)), 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  $\mathcal{C}(a, s) = 1 \quad \forall s \in \mathcal{S}, a \in \mathcal{A}(s)$ . A sequence of actions  $a_0, \dots, a_{n-1}$ , called plan, is a valid solution to the model if for  $0 \leq i \leq n - 1 \quad s_{i+1} \in \mathcal{F}(s_i, a_i)$ ,  $a_i \in \mathcal{A}(s_i)$  and  $s_n \in F(s_{n-1}, a_{n-1}) \cap S_G$ .
- **Nondeterministic Models** (Fig. 1 (b)), 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 and using the terminology discussed in Section 3, we actually have *planning under pure Knightian uncertainty*. Since we assume full observability, a valid solution to the model is a policy, i.e. a function  $\pi : \mathcal{S} \rightarrow \cup_{s \in \mathcal{S}} \mathcal{A}(s)$ , that is *closed and proper* with respect to  $S_0$  [6]. In this model, a policy offers guarantees about the worst-case behavior of the environment.
- **Probabilistic Models** (Fig. 1 (c)), 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  $F(s, a) \quad \forall s \in \mathcal{S}, a \in \mathcal{A}(s)$ . As in the Nondeterministic Models, a solution to the model is a policy, but in this case the objective is to maximize expected behavior — where expected behavior is quantified through a *single* probability measure.

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 Probabilities

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 [7]. We consider a set of states of nature  $\Omega$ ; then each option is a function  $f : \Omega \rightarrow \mathfrak{R}$ . If a decision maker can specify a single probability measure  $P$  over a field of events defined on  $\Omega$ , 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 dominated by expected utility — a simple criterion that leads to a rich theory [8].

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 [7]. 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 [9]). 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 uncertainty* 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 [10,11,12], either because there is no 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 [13,14]. 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 [8,15,16].

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 [17,18,19], in Dempster-Shafer theory [20], in theories of argumentation [21], and in techniques that generalize graph-theoretic models such as Bayesian networks [22,23,24].<sup>1</sup> Our contribution here is to identify the planning under uncertainty spectrum with the theory of sets of probability measures.

## 4 Planning Under Uncertainty: The Risk and Knightian Uncertainty Spectrum

The *planning under uncertainty model*, referred to here as the **PUU model**, is a more general model as it gives a precise semantics to planning tasks involving nondeterministic and probabilistic *effects* of actions. A planning problem can be solved considering, simultaneously, these two types of action's effects.

Since the PUU model has to represent nondeterministic effects, the transition function  $F(s, a)$  (from the basic state model described in Section 2) must be understood as follows. Instead of taking  $F(s, a) \subseteq \mathcal{S}$  as before, now we 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 PUU model, the transition function  $F(s, a)$  maps states  $s$  and actions  $a \in \mathcal{A}(s)$  into nonempty sets of the parts of  $\mathcal{S}$ :  $F(s, a) \subseteq 2^{\mathcal{S}}$ .*

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

---

<sup>1</sup> There is now significant literature on the theory and applications of sets of probability measures [25,26,27,28].

With the above definitions the probabilistic function  $P(k|s, a)$ ,  $k \in F(s, a)$  has a novel interpretation in the PUU model:  $P_n(k|s, a)$  represents the probability of the next state to be within one of the states in  $k$ . The PUU model is illustrated in the second row of Fig. 1 and a formal description of this model is given by:

- PUU1 a discrete and finite state space  $\mathcal{S}$ ,
- PUU2 a nonempty set of initial states  $S_0 \subseteq \mathcal{S}$ ,
- PUU3 a goal situation given by a nonempty set  $S_G \subseteq \mathcal{S}$ ,
- PUU4 a nonempty set of actions  $\mathcal{A}(s) \subseteq \mathcal{A}$  representing the actions applicable in each state  $s$ ,
- PUU5 a state transition function  $F(s, a) \subseteq 2^{\mathcal{S}}$  mapping states  $s$  and actions  $a \in \mathcal{A}(s)$  into nonempty sets of the parts of the state space,
- PUU6 a probability distribution  $P_n(\cdot|s, a)$  over  $F(s, a) \quad \forall s \in \mathcal{S}, a \in \mathcal{A}(s)$  where  $P_n(k|s, a)$  represents the probability of choosing the possible-state set  $k \subseteq \mathcal{S}$  when action  $a$  is applied in state  $s$ , and
- PUU7 a positive action cost  $C(a, s)$  for doing  $a \in \mathcal{A}(s)$  in  $s$ .

Notice that there are two types of choices happening the PUU model: a probabilistic choice of a possible-state set and a nondeterministic choice of a successor state from the possible-state set. The planning under uncertainty task can thus be characterized by domains for which the action dynamics satisfies the following restrictions: (1)  $\|F(s, a)\| > 1$  and (2)  $\exists k \in F(s, a)$  s.t.  $\|k\| > 1$ , for  $s \in \mathcal{S}, a \in \mathcal{A}(s)$  (Fig. 1 (d)). If none of these requirements is true, then PUU 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 PUU model is equivalent to the nondeterministic model (Fig. 1 (e)). This is because:  $\forall s \in \mathcal{S}, a \in \mathcal{A}(s)$ , if  $\|F(s, a)\| = 1$  then  $P_n(k \in F(s, a)|s, a) = 1$ , 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 (Fig. 1 (f)). 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  $P_n(k|s, a)$  after executing  $a$  in the state  $s$ . Under this assumptions the probability distribution over  $2^{\mathcal{S}}$  is equivalent to a probability distribution over  $\mathcal{S}$ .

Finally, when both requirements are false, the model is equivalent to the deterministic model (Fig. 1 (d)) once there is no point of choice: neither in the probabilistic choice of a possible-state set nor in the nondeterministic choice of a successor state.

Furthermore, the complete PUU 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 PUU model.

## 5 The Relation Between PUU and MDPIP Model

Markov Decision Processes with Imprecise Probabilities (MDPIPs) [29,30] are an extension of Markov Decision Processes (MDPs) [31] 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  $\mathcal{K}$  [23] 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 situation given by a nonempty set  $S_G \subseteq \mathcal{S}$ ,
- MIP3 a nonempty set of actions  $\mathcal{A}(s) \subseteq \mathcal{A}$  representing the actions applicable in each state  $s$ ,
- MIP4 a nonempty 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 plannet wishes to minimize (1). Therefore, a minimax 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 [30] it has been shown that the solution to (1), called  $V^*(s)$ , exists and is unique. It is also proved that the optimal policy for an MDPIP can be expressed by a stationary policy, i.e., the same policy for any instant in time. We have the following fundamental relationship between the PUU model and MDPIPs:

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

*Proof.* Note that PUU1, PUU3, PUU4 and PUU7 are equal, respectively, to MIP1, MIP2, MIP3 and MIP5. Thus the proof is reduced to prove that PUU5 and PUU6 implies in MIP4.

First, note that PUU6 bounds the probability of being in state  $s'$  after applying the action  $a$  in the state  $s$  by (2). This is due to the definition of possible-state set: let  $k \in F(s, a)$ , if  $s' \notin k$ , then nature is not able to choose  $s'$  as a non-deterministic effect of  $a$ .

$$P_n(\{s'\}|s, a) \leq P(s'|s, a) \leq \sum_{\substack{k \in F(s,a) \\ s' \in k}} P_n(k|s, a) \leq 1 \quad \forall s' \in \mathcal{S} \tag{2}$$

Let us define the set of states  $\mathcal{D}(k, s, a)$ , for (3) . This set represents all non-deterministic effects of  $k$  that bellong only to  $k$ .

$$\mathcal{D}(k, s, a) = k \setminus \bigcup_{\substack{k' \in F(s,a) \\ k' \neq k}} k' \tag{3}$$

From PUU5 and PUU6 it is possible to bound the sum of the probabilities of each state in a possible-state set  $k \in F(s, a)$  and in the associated set  $\mathcal{D}(k, s, a)$ . These bounds are presented in (4).

$$0 \leq \sum_{s' \in \mathcal{D}(k, s, a)} P(s'|s, a) \leq P_n(k|s, a) \leq \sum_{s' \in k} P(s'|s, a) \leq 1 \quad (4)$$

The set of inequations (2) and (4) for a state  $s \in \mathcal{S}$  and an action  $a \in \mathcal{A}(s)$  describe a possible credal set  $\mathcal{K}_s(a)$  for MIP4.  $\square$

Proposition 1 not only makes the results in [29,30] valid for the PUU model, but also suggests algorithms to solve it. An algorithm that can solve a PUU problem, inspired by previous algorithms for MDPIPs, is given in the next section.

## 6 Algorithms

Due to Proposition 1, every algorithm for MDPIPs can be applied to solve a PUU problem. However, the process is not immediate as it is necessary to adapt several key concepts. To illustrate this, we selected a modified version of the policy-iteration algorithm [31] given by [30]. This algorithm is divided in two phases, *policy evaluation*, where the expected utility of a policy is calculated, and *policy improvement* in which a better policy is built based on the values received from the *policy evaluation* phase.

The algorithm presented bellow receives an tuple  $\langle \mathcal{S}, \mathcal{A}, K_s(a) \rangle$  (constructed by proposition 1),  $C(a, s)$ , a discount factor  $\gamma$  and an initial, possible random, policy  $\pi$  as input and return the optimal policy for this PUU problem.

*Policy evaluation.* For every state  $s \in \mathcal{S}$ :

- (a) Select a probability distribution  $P(\cdot|s, \pi(s)) \in \mathcal{K}_s(\pi(s))$ .
- (b) Use  $P(\cdot|s, \pi(s))$  to solve (5).

$$V_\pi(s) = C(\pi(s), s) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a) V_\pi(s') \quad (5)$$

- (c) Find the probability distribution  $P'(\cdot|s, \pi(s)) \in \mathcal{K}_s(\pi(s))$  s.t.

$$\max_{P'(\cdot|s, \pi(s)) \in \mathcal{K}_s(\pi(s))} \{C(\pi(s), s) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, \pi(s)) V_\pi(s')\}, \quad (6)$$

where  $V_\pi(s')$  are solutions to (5).

- (d) If  $V_\pi(s)$  is equal to  $C(\pi(s), s) + \gamma \sum_{s' \in \mathcal{S}} P'(s'|s, a) V_\pi(s')$  then proceed to the *policy improvement* phase, otherwise, return to step (b) with  $P(\cdot|s, \pi(s)) \leftarrow P'(\cdot|s, \pi(s))$ .

*Policy improvement.* For every state  $s \in \mathcal{S}$ , find  $\pi'(s)$  s.t.

$$\pi'(s) = \operatorname{argmin}_{a \in \mathcal{A}(s)} \max_{P(\cdot|s, a) \in \mathcal{K}_s(a)} \{C(a, s) + \gamma \sum_{s' \in \mathcal{S}} P(s'|s, a) V_\pi(s')\}, \quad (7)$$



using the value  $V_\pi(s')$  from (5). If  $\pi'(s)$  is different from  $\pi(s)$  for some state  $s \in \mathcal{S}$ , return to the *policy evaluation* phase with  $\pi \leftarrow \pi'$ ; otherwise,  $\pi$  is the minimax optimal solution for the PUU problem.

## 7 Conclusions and Future Work

We can list the following contributions of this work. First, we identified a general formulation that encompasses probabilistic and nondeterministic planning, *and that includes* a continuum of planning problems between these extremes. In fact, the probabilistic/nondeterministic actions we define open new types of planning scenarios that go beyond existing planning problems. For example, one may have actions whose transitions are specified by general sets of probabilities, not just the set-valued consequences discussed in this paper. Second, we have shown how our proposal fits within the MDPIP framework. Proposition 1 establishes the link between planning under risk and Knightian uncertainty and the previous literature on MDPIPs. We have then shown how to transfer algorithms previously developed for MDPIP to our proposal.

This paper should open a profitable avenue for future research in a variety of directions. It would be interesting to explore the many possible combinations of probabilistic and nondeterministic actions, and more general transitions defined by sets of probability measures. However, we feel that the most fruitful task for the near future is to adapt existing online algorithms for existing planning problems, such as RTDP and LRTDP [3], to the full generality of planning under risk and Knightian uncertainty.

## 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: ICAPS, To Appear (2006)
3. Bonet, B., Geffner, H.: Labeled RTDP: Improving the convergence of real-time dynamic programming. In: ICAPS, 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. Bonet, B., Geffner, H.: Planning with incomplete information as heuristic search in belief space. In: ICAPS, Breckenridge, CO, AAAI Press (2000) 52–61
7. Luce, D., Raiffa, H.: Games and Decisions. Dover edition, Mineola (1957)
8. Berger, J.O.: Statistical Decision Theory and Bayesian Analysis. Springer-Verlag (1985)
9. Knight, F.H.: Risk, Uncertainty, and Profit. Hart, Schaffner & Marx; Houghton Mifflin Company, Boston (1921)

10. Levi, I.: *The Enterprise of Knowledge*. MIT Press (1980)
11. Walley, P.: *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London (1991)
12. Walley, P.: Measures of uncertainty in expert systems. *AI* **83** (1996) 1–58
13. 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
14. 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
15. Huber, P.J.: *Robust Statistics*. Wiley, New York (1980)
16. Kadane, J.B., ed.: *Robustness of Bayesian Analyses*. Volume 4 of *Studies in Bayesian econometrics*. Elsevier Science Pub. Co., New York (1984)
17. Frisch, A.M., Haddawy, P.: Anytime deduction for probabilistic logic. *Artificial Intelligence* **69** (1994) 93–122
18. Halpern, J.Y.: *Reasoning about uncertainty*. MIT Press (2003)
19. Nilsson, N.J.: Probabilistic logic. *Artificial Intelligence* **28** (1986) 71–87
20. Shafer, G.: *A Mathematical Theory of Evidence*. Princeton University Press (1976)
21. 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)
22. Cozman, F.G.: Credal networks. *AI* **120** (2000) 199–233
23. Cozman, F.G.: Graphical models for imprecise probabilities. *International Journal of Approximate Reasoning* **39**(2-3) (2005) 167–184
24. Fagioli, E., Zaffalon, M.: 2U: An exact interval propagation algorithm for polytrees with binary variables. *Artificial Intelligence* **106**(1) (1998) 77–107
25. de Cooman, G., Cozman, F., Moral, S., Walley, P., eds.: *Proceedings of the First International Symposium on Imprecise Probabilities and Their Applications (SIPTA)*. Universiteit Gent, Ghent, Belgium (1999)
26. de Cooman, G., Fine, T.L., Seidenfeld, T.: *Proceedings of the 2nd International SIPTA* Shaker Publishing, The Netherlands (2001)
27. Bernard, J.M., Seidenfeld, T., Zaffalon, M., eds.: *Proceedings of the 3rd International SIPTA* Carleton Scientific, Lugano, Switzerland (2003)
28. Cozman, F.G., Nau, R., Seidenfeld, T.: *Proceedings of the Fourth International Symposium on Imprecise Probabilities and Their Applications*. SIPTA (2005)
29. White III, C.C., Eldeib, H.K.: Markov decision processes with imprecise transition probabilities. *Operations Research* **42**(4) (1994) 739–749
30. Satia, J.K., Jr, R.E.L.: Markovian decision processes with uncertain transition probabilities. *Operations Research* **21**(3) (1973) 728–740
31. Howard, R.A.: *Dynamic Programming and Markov Processes*. MIT Press (1960)