On the Impact of Belief State Representation in Planning under Uncertainty

Son Thanh To
New Mexico State University
Dept. of Computer Science
sto@cs.nmsu.edu

Planning Under Uncertainty

Planning under uncertainty is one of the most general and hardest problems considered in the area [Rintanen 2004]. Uncertainty can take the form of incomplete information, wrong information, multiple action outcomes, and varying action durations. My doctoral thesis concentrates on planning with incomplete knowledge and multiple action outcomes, specifically *conformant planning* and *contingent planning*.

Conformant planning [Goldman and Boddy 1996] is a problem of finding sequences of actions for the agent to act, with no observations, to achieve the goal from any possible initial state of the world in presence of incomplete initial information. Conformant planning is useful when observations are expensive, dangerous, and/or impossible. Conformant planning has attracted the attention of several researchers. A number of efficient and sophisticated conformant planners have been developed CFF [Brafman and Hoffmann, 2004], POND [Bryce *et al.*, 2006], `t0` [Palacios and Geffner, 2007], and CPA [Tran *et al.*, 2009].

Contingent planning is a more general problem of conformant planning that allows non-deterministic actions and observation of some properties of the world, in addition to incomplete information. The contingent plan allows the agent to act, at the execution time, conditionally depending on the observed values; and guarantees to achieve the goal no matter what the actual initial world the agent starts from and which actual action outcomes occur. Contingent planning is more practical and harder, as opposed to conformant planning. Many work have been done on contingent planning resulting in various contingent planner. Among the most competitive contingent planners are contingent-FF [Hoffmann and Brafman, 2005], POND, and CLG [Albore et al., 2009].

Previous Approaches

To deal with incomplete information about the world, the notion of *belief state* has been introduced—defined as a set of possible states. This notion is convenient for capturing the semantic of incomplete information and uncertain action effects and for defining a transition function between belief states. The use of this representation in the implementation of a planner, however, is inefficient and impractical due to its exponential size. To address this, numerous research work have been done with proposals of different approaches. Significant progress can be observed by the introduction of variety of planning systems that can solve problems of different size at different level of hardness, usually using the approach that encodes the planning problem as a search problem in the belief state space. However, the scalability of these planners, though are among the best planning systems in the literature, is still modest, mostly due to disadvantages of the methods they use to represent belief states. For example, the representation using *binary decision diagrams* (BDDs) [Bryant, 1992], used in POND, is usually very large and its size is sensitive to the order of the variables. Moreover, computing successor belief states in BDDs representation requires generation of intermediate formulae of exponential size. In contrast, the method used in CFF and contingent-FF, that encodes belief states implicitly through the action sequences leading to them from the initial belief state, needs a little space but incurs an excessive amount of repeated computation. Furthermore, checking whether a proposition holds after the execution of an action sequence is exponentially harder, compared with the case of the execution of one action, in general. `t0` and CLG transform the problem into a search problem in the state space, whose literals represent the beliefs over the original problem. This approach is fairly efficient as `t0` and CLG outperforms the others in a set of benchmarks. However, the number of literals in the translated problem can be exponential in the number of unknown literals in the original problem, making the state space extremely large and preventing the planners to scale up. Finally, the method that approximates belief states used in CPA is efficient in several problems. Yet the size of the approximated formulae, if satisfies the complete condition, explodes in many other problems.

Motivation and Approaches of this Work

This work provides a systematic methodology for dealing with planning under uncertainty, focusing on the representation of belief states that can be used in a forward search paradigm in the belief space for solutions. A good representation should be compact so that a planner implementing it can perform and scale up well as the larger the formulae, the more the computation and the more the memory consumption (i.e., the slower the system and the less the scalability). On the other hand, it should also have properties that allow for definition of an efficient transition function for computing successor belief states, e.g., checking satisfaction in a DNF formula is easy. Defining a direct complete transition function in presence of incomplete information for a general representation, other than the belief state, is particularly hard due to conditional action effects. To address this, I propose a *generic abstract algorithm*, called *GAA*, for defining such function given an arbitrary representation. For each concrete representation, however, GAA needs to be instantiated with some effort put into. The degree of effort required depends on the properties of the considered representation. The GAA algorithm also helps to evaluate the effectiveness of the representation through the corresponding transition function. By means of the GAA algorithm, my doctoral thesis investigates the properties of different logical formulae and their applicability in planning under uncertainty as a belief state representation. If a formula appears to be a promising representation, then a planning system is developed using that representation.

It is worth noting that the study of different representations is useful and beneficial as each representation is strong or weak in certain classes of problems. To this end, my research also investigates the advantages and disadvantages of alternative representations and identifies the classes of problems that promote or degrade each method. Moreover, the results obtained from this study can also be applied in other areas, since logical formulae are widely used in various areas.

Beside the study of belief state representation, this work also consider such techniques that help to improve the performance of the planners as simplification of the problem, heuristics, and search space pruning. So far, the results obtained in the presented direction is very promising.

Several Results

The first representation that was successfully investigated by this work is *minimal-DNF*, a compact disjunctive normal form formula, implemented in a conformant planner, called DNF [To *et al.*, 2009]. This representation allows for definition of an efficient transition function, i.e., polynomial under reasonable assumptions, and appears to be very effective as the performance and scalability of DNF is highly competitive in a large set of benchmarks available in the literature.

The second result obtained by the work is CNF, a conformant planner using a compact disjunctive normal form formula, called *reduced-CNF* [To *et al.*, 2010a]. CNF outperforms DNF on the problems where the disjunctive formulae representing belief states are much larger than the equivalent conjunctive formulae. The paper also identifies those problems and provides a technique, called *one-of relaxation*, to reduce the size of the CNF-formulae.

A further investigation on conjunctive formulae shows that the *prime implicates* have properties desirable for a representation, e.g., checking satisfaction is easy, the conversion of a small set of prime implicate formulae to a prime implicate formula required in the state computation is polynomial time, updating the formula w.r.t. a literal that become true after the execution of an action is very fast, etc... A new conformant planner, called PIP, using two alternative representations prime implicates and minimal-CNF has been introduced in [To *et al.*, 2010b]. The second representation is needed when the prime implicate formula is very large.

The most recent result of the work is [To *et al.*, 2011], which extends the transition function defined in [To *et al.*, 2009] to deal with non-deterministic and sensing actions and implemented it in a contingent planner, called DNF_{ct}, using a new AND/OR forward search algorithm, called PrAO. The novelty of PAO is a new pruning technique which safely prune the search space significantly. DNF_{ct} outperforms all the most competitive contingent planners on most benchmarks available in the literature.

Future Work

My dissertation continues with the effectiveness of the other investigated representations in contingent planning. DNNF formula [Darwiche, 2001] and their variants appear to be potential too, as they are compact in a certain class of problems and checking entailment in a DNNF formula is easy. Finally, looking for effective heuristics in planning under uncertainty is also a subject of this research, which is particularly hard due to incomplete information, that makes the evaluation of a heuristic function much more expensive and less accurate.

References

[Albore et al., 2009] A. Albore, H. Palacios, and H. Geffner. A Translation-based Approach to Contingent Planning. In *IJCAI*, 2009.

[Brafman and Hoffmann, 2004] R. Brafman and J. Hoffmann. Conformant planning via heuristic forward search: A new approach. In *ICAPS*, 355–364, 2004.

[Darwiche, 2001] Darwiche, A. 2001. Decomposable Negation Normal Form. *Journal of the ACM* 48(4):608–647.

[Goldman and Boddy 1996] Goldman, R. and Boddy, M. 1996. Expressive Planning and Explicit Knowledge. In *Proc. of 3rd Int. Conf. AI Planning Systems*, 110–117.

[Hoffmann and Brafman, 2005] J. Hoffmann and R. Brafman. Contingent planning via heuristic forward search with implicit belief states. In *ICAPS*, 2005.

[Bryant, 1992] R. E. Bryant. Symbolic boolean manipulation with ordered binary decision diagrams. *ACM Computing Surveys*, 24(3):293–318, September 1992.

[Bryce *et al.*, 2006] D. Bryce et al. Planning Graph Heuristics for Belief Space Search. *JAIR*, 26:35–99, 2006.

[Palacios and Geffner, 2007] H. Palacios and H. Geffner. From Conformant into Classical Planning: Efficient Translations that may be Complete Too. In *ICAPS*, 2007.

[Rintanen 2004] Rintanen, J. 2004. Complexity of planning with partial observability. In *ICAPS*, 2004.

[To *et al.*, 2009] S. T. To et al. A Conformant Planner with Explicit Disjunctive Representation of Belief States. In *ICAPS*, 2009.

[To *et al.*, 2010a] S. T. To et al. A New Approach to Conformant Planning using CNF. In *ICAPS*, 2010a.

[To *et al.*, 2010b] S. T. To et al. On the Use of Prime Implicates in Conformant Planning. In *AAAI*, 2010b.

[To *et al.*, 2011] S. T. To et al. Contingent Planning as AND/OR forward Search with Disjunctive Representation. To appear in Proc. of *ICAPS*, 2011.

[Tran *et al.*, 2009] S. D. Tran et al. Improving performance of conformant planners. In *PADL*. Springer, 2009.