Reasoning about Deterministic Action Sequences with Probabilistic Priors

Hannaneh Hajishirzi and Eyal Amir
Department of Computer Science
University of Illinois at Urbana-Champaign
{hajishir, eyal}@illinois.edu
Paper number: #118

Abstract

We present a novel algorithm and a new understanding of reasoning about a sequence of deterministic actions with a probabilistic prior. When the initial state of a dynamic system is unknown, a probability distribution can be still specified over the initial states. Estimating the posterior distribution over states (*filtering*) after some deterministic actions occurred is a problem relevant to AI planning, natural language processing (NLP) and robotics among others. Current approaches to filtering deterministic actions are not tractable even if the distribution over the initial system state is represented compactly. The reason is that state variables become correlated after a few steps. The main innovation in this paper is a method for sidestepping this problem by redefining state variables dynamically at each time step such that the posterior for time t is represented in a factored form. This update is done using a progression algorithm as a subroutine, and our algorithm's tractability follows when that subroutine is tractable. In particular, our algorithm is tractable for 1:1 and STRIPS actions. We apply our reasoning algorithm about deterministic actions to reasoning about sequences of probabilistic actions and improve the efficiency of the current probabilistic reasoning approaches. We demonstrate the efficiency of the new algorithm experimentally over AI-Planning data sets.

1 Introduction

Many applications in AI involve stochastic dynamic systems and answering queries about them. Examples of such applications are Natural Language Processing (NLP), robotics, AI planning, autonomous agents, speech recognition, and commonsense query answering. Inference in a stochastic dynamic system is the problem of estimating the systems' state given a sequence of actions and partial observations. This type of inference is called *filtering* in stochastic dynamic systems (it is usually called *progression* when no probabilities are involved).

When applied to real-world problems, many of these applications have very large state spaces, uncertain initial states, and uncertain effects of actions. Therefore, it is crucial to choose a representation that is compact and models the stochasticity of the domain, a representation that also enables efficient inference.

In recent years, there is a growing interest in using action-centered languages for representing dynamic systems in AI planning, NLP, and robotics. Planning Domain Definition Language (PDDL)(McDermott 2000) expresses semantics of actions by describing their preconditions and effects and represents the dynamics of the system. Situation calculus (Reiter 2001) represents changing scenarios with a series of first-order logical formulas. PDDL and Situation calculus are traditionally without probabilities. To model stochasticity, they are augmented with *probabilistic choice* actions which choose deterministic executions and a probability distribution over initial states.

Current approaches to probabilistic filtering with PDDL or situation calculus (e.g., (Reiter 2001; Bacchus, Halpern, and Levesque 1999; Hajishirzi and Amir 2008)) use traditional filtering methods as subroutines for reasoning about sequences of deterministic actions with probabilistic priors. These traditional methods are inefficient or imprecise for deterministic sequences. Some approaches marginalize over all possible initial states (exponential in the number of variables) to compute the posterior probability of a query. In others (e.g., Dynamic Bayesian Networks (DBNs) (Dean and Kanazawa 1988)) all the state variables become fully correlated after a few steps even if they are independent at time 0, resulting in a posterior representation of size exponential in the number of variables. Others use logical regression and repeat $t-1$ regressions for every new added action, so are inefficient for long sequences of actions.

The main contribution of this paper is an understanding of conditional-independence structure preservation over time in systems with deterministic actions and stochastic priors over initial states. Our new understanding leads to a new exact algorithm for reasoning about sequences of deterministic actions with a probabilistic prior over the initial states. The algorithm is tractable for 1:1 and STRIPS actions, following results of (Amir and Russell 2003).

We use a propositional version of probabilistic situation calculus that is extended with a graphical model prior for representing dynamic systems. In particular, the initial knowledge is represented with a prior distribution over state variables (in a Bayesian Network (BN) (Pearl 1988) format) and transitions are modeled naturally as stochastic choices among deterministic actions. Our algorithm uses a deterministic progression subroutine and represents the posterior at time t with a BN whose structure and conditional probabilities are identical to those of the BN of time 0, but whose nodes have a new meaning.

Specifically, every node in the BN representation of posterior at time t corresponds to a propositional logical formula that represents a set of world states. For example, when a binary node X_i (time 0) takes value 1, then it represents the set of world states that satisfy $X_i = 1$. At time t, a binary node Φ_i^t would be a logical formula over $x_1, ..., x_n$ at time t. When this binary node takes value 1, then it represents the set of world states that satisfy $\Phi_i^t = 1$. The BN comprised of such nodes at time t represents the posterior distribution over states at time t.

Finally, we apply our exact filtering algorithm to reason about sequences of probabilistic actions. Our empirical results show that our new algorithm improves the efficiency of the sampling algorithm (Hajishirzi and Amir 2008) for filtering with probabilistic actions. The improvement is due to the fact that we remove the expensive subroutine of regressing to time zero at every time step and just use progression.

(Reiter 2001) and (Bacchus, Halpern, and Levesque 1999) present exact algorithms to answer a query given a sequence of actions and observations in a dynamic system represented in a probabilistic situation calculus form. Both algorithms marginalize over all the possible initial states and all the possible deterministic sequences to compute the probability of a world state at time step t. Both algorithms assign probability to every world state individually, while our method uses a BN to compactly represent the prior distribution.

First order MDPs (Boutilier, Reiter, and Price 2001) use probabilistic situation calculus to represent the dynamics of the system. They introduce a dynamic programming approach for solving MDPs by describing the optimal value function and policies in a logical format. Their approach uses a logical regression subroutine which results in a combinatorial explosion even for simple deterministic actions.

(Pasula, Zettlemoyer, and Kaelbling 2004) learn the probability distribution over deterministic executions of probabilistic actions which is different from filtering with action sequences. Their representation does not include a compact prior over the initial states.

A DBN compactly represents a dynamic system using a BN for time 0 and a graphical representation of a transition model between times t and $t + 1$. DBNs focus on conditional independence assumption whereas our representation focuses on decomposition of actions into deterministic actions. Traditional methods for exact inference algorithms in DBNs (Murphy 2002) are not tractable because all the state variables become correlated after a few steps even for deterministic transitions. (Pfeffer 2001) presents an exact tractable inference algorithm for a class of DBNs with no observations. This method assumes that the DBN is decomposed into separable subsystems. In contrast, our exact inference is applicable to deterministic inseparable DBNs.

2 Probabilistic Action Model

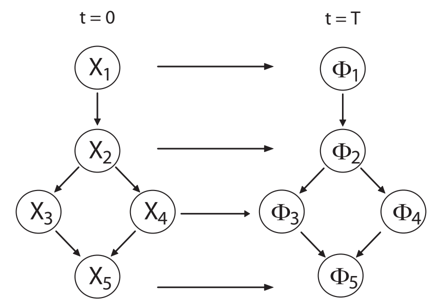

Figure 1: (*left*) *BN* at time 0 with state variables $X_1 \dots X_n$, (*right*) New *BN* constructed at time T with new BN bases $\Phi_1 \dots \Phi_n$. *BN*^{0} and *BN*^{T} have identical structure.

References

Amir, E., and Russell, S. 2003. Logical filtering. In *IJCAI '03*, 75–82. MK.

Bacchus, F.; Halpern, J. Y.; and Levesque, H. J. 1999. Reasoning about noisy sensors and effectors in the situation calculus. *AIJ* 111(1–2):171–208.

Boutilier, C.; Reiter, R.; and Price, B. 2001. Symbolic dynamic programming for fo mdps. In *IJCAI*.

Chavira, M., and Darwiche, A. 2008. On probabilistic inference by weighted model counting. *AIJ* 172(6-7).

Dean, T., and Kanazawa, K. 1988. Probabilistic temporal reasoning. In *Proc. AAAI '88*, 524–528.

Doucet, A.; de Freitas, N.; and Gordon, N. 2001. *Sequential Monte Carlo Methods in Practice*. Springer, 1st edition.

Gomes, C.; Sabharwal, A.; and Selman, B. 2006. Model counting: A new strategy for obtaining good bounds. In *AAAI*.

Hajishirzi, H., and Amir, E. 2007. Stochastic filtering in a probabilistic action model. In *AAAI*.

Hajishirzi, H., and Amir, E. 2008. Sampling first order logical particles. In *UAI*.

Lavasque, H. 1998. A completeness result for reasoning with incomplete knowledge bases. In *KR*.

McDermott, D. 2000. The 1998 AI planning systems competition. *AI Magazine* 21(2).

Murphy, K. 2002. *Dynamic Bayesian Networks: Representation, Inference and Learning*. Ph.D. Dissertation, Berkeley.

Pasula, H. M.; Zettlemoyer, L. S.; and Kaelbling, L. P. 2004. Learning probabilistic relational planning rules. In *Proc. ICAPS'04*.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems : Networks of Plausible Inference*. Morgan Kaufmann.

Pfeffer, A. 2001. Sufficiency, separability and temporal probabilistic models. In *Proc. UAI '01*, 421–428. MK.

Rabiner, L. R. 1989. A tutorial on hmm and selected applications in speech recognition. *IEEE* 77(2).

R.E.Bryant. 1992. Symbolic boolean manipulation with ordered BDDs. *ACM* 24(3).

Reiter, R. 2001. *Knowledge In Action*. MIT Press.

Shahaf, D., and Amir, E. 2007. Logical circuit filtering. In *IJCAI-07*.