Transferring State Abstractions Between MDPs

Thomas J. Walsh   THOMASWA@CS.RUTGERS.EDU
Lihong Li   LIHONG@CS.RUTGERS.EDU
Michael L. Littman   MLITTMAN@CS.RUTGERS.EDU
Department of Computer Science, Rutgers University, 110 Frelinghuysen Road, Piscataway, NJ 08854

Abstract

Decision makers that employ state abstraction (or state aggregation) usually find solutions faster by treating groups of states as indistinguishable by ignoring irrelevant state information. Identifying irrelevant information is essential for the field of knowledge transfer where learning takes place in a general setting for multiple domains. We provide a general treatment and algorithm for transferring state abstractions in MDPs.

1. Introduction

The fields of "transfer learning" and state abstraction are closely related. In transfer learning, an agent attempts to reuse knowledge from several source domains in one or more target domains that may not be represented in the set of source domains. This goal requires some notion of a "generalized" or "abstract" state space where reasoning can be done and applied to any domain in the set. Thus, deciding what knowledge to transfer between domains can be construed as determining the correct state *abstraction* scheme for a set of source domains and then applying this compaction to a target domain. In this paper, we present an algorithm for accomplishing this abstraction transfer in structured Markov Decision Processes (MDPs) (Puterman, 1994) based on prior work that formalized abstraction in MDPs (Li et al., 2006) and prior work in transferring information for acting optimally in multiple MDPs (Guestrin et al., 2003; Jong & Stone, 2005).

2. Definitions

MDPs can be described as a five-tuple $\langle S, A, P, R, \gamma \rangle$ where S is a finite set of states; A is a finite set of actions; P is the transition function; R is a bounded reward function; and $\gamma \in [0, 1]$ is a discount factor. In this paper we deal with structured MDPs, where the state space is comprised of n multi-valued features $f_1...f_n$. A policy is a mapping from states to actions: $\pi : \mathcal{S} \mapsto \mathcal{A}$. The state-action value function, $Q^\pi(s, a)$, is the expected cumulative reward received by taking action a in state s and following π thereafter. A reinforcement-learning agent (Sutton & Barto, 1998) attempts to *learn* an optimal policy π^* with value function $Q^*(s, a)$. An abstraction will be defined as a function $\phi : S \mapsto \bar{S}$; $\phi(s) \in \bar{S}$ where S is the *ground* state space and $\bar{S}$ the *abstract* state space.

We define "transferring learning" in structured MDPs in terms of a distribution D over possible *target* MDPs, which share some common features (such as states defined in terms of the same variables or relations). Using a set $\mathcal{M}_s$ of m *source* MDPs sampled from D, we wish to uncover some information that allows us to maximize the following *speedup ratio*:

$$\frac{\mathbf{E}_{M_t \sim D} \{T(M_t)\}}{\mathbf{E}_{\mathcal{M}_s \sim D, M_t \sim D} \{T(M_t | \mathcal{M}_s)\}}$$

where $T(M_t)$ is the time needed to find an optimal policy in a target MDP, M_t, and $T(M_t | \mathcal{M}_s)$ is the time needed to find an optimal policy if information is transferred from $\mathcal{M}_s$. In this paper, we focus on learning and transferring MDP abstractions.

3. Prior Work

In this section we review several intuitive definitions of abstraction for MDPs, present several key properties of these abstractions and discuss previous attempts to achieve knowledge transfer via abstractions.

3.1. Five Types of Abstraction

Defining state-abstraction rules for MDPs has been the focus of previous work, culminating in a unified treatment of the problem (Li et al., 2006), which provides an exhaustive list of previously proposed ab-

straction techniques as well as defining five abstraction schemes based on seemingly important features of MDPs: : ϕ_{model}, ϕ_{Q^π}, ϕ_{Q^*}, ϕ_{a^*}, and ϕ_{π^*}. Intuitively, ϕ_{model} preserves the one-step model (e.g., bisimulation (Givan et al., 2003)); ϕ_{Q^π} preserves the state-action value function for all policies; ϕ_{Q^*} preserves the optimal state-action value function (e.g., stochastic dynamic programming (Boutilier et al., 2000), the G-algorithm (Chapman & Kaelbling, 1991)), Symbolic Dynamic Programming (Boutilier et al., 2001; Großmann et al., 2002), or Q-RRL (Dzeroski et al., 2001); ϕ_{a^*} preserves the optimal action and its value (e.g., utile distinction (McCallum, 1995)); and ϕ_{π^*} preserves the optimal action (e.g. Policy Irrelevance (Jong & Stone, 2005) or P-RRL (Dzeroski et al., 2001)). [1]

3.2. Abstraction Properties

Here, we present several known theoretical results for the five abstraction schemes, focusing on how they affect planning and learning algorithms. We refer the reader to previous work (Li et al., 2006) for the proofs.

First, we consider whether standard dynamic-programming algorithms such as value iteration and policy iteration, when applied to the abstract MDP $\bar{M}$ yield an optimal policy $\bar{\pi}^*$ that maps to an optimal policy in the ground MDP.

Theorem 1 *With abstractions ϕ_{model}, ϕ_{Q^π}, ϕ_{Q^*}, and ϕ_{a^*}, the optimal abstract policy $\bar{\pi}^*$ is optimal in the ground MDP. However, examples exist where the optimal policy with abstraction ϕ_{π^*} is suboptimal in the ground MDP.*

Next, we consider the problem of *learning* the value function, where the agent estimates the optimal value function based on experience. We extend the standard Q-learning update to the abstract MDP as:

$$Q(\phi(s_t), a_t) \quad \overset{\alpha_t}{\longleftarrow} \quad r_t + \gamma \max_{a'} Q(\phi(s_{t+1}), a').$$

Theorem 2 *Q-learning with abstractions ϕ_{model}, ϕ_{Q^π}, or ϕ_{Q^*} converges and yields a policy that is optimal in the ground MDP. Q-learning with abstraction ϕ_{a^*} does not necessarily converge, but will converge with a fixed behavior policy; in either case, it yields a value function whose greedy policy is optimal in the ground MDP. Q-learning with abstraction ϕ_{π^*} can converge to an state-action value function whose greedy policy is suboptimal in the ground MDP.*

A more detailed discussion of this result, with exam-

ples and an extension to model-based reinforcement learning, is provided in previous work (Li et al., 2006). We note that since these abstractions are defined independently of the language used to represent the state space, these results are applicable to domains specified in any language, including logical ones (Dzeroski et al., 2001; Boutilier et al., 2001). The preceding results are consistent with the negative results for ϕ_{π^*} in the state abstraction (Gordon, 1996; Jong & Stone, 2005) and relational reinforcement learning (RRL) literature (Dzeroski et al., 2001).

3.3. Abstraction Transfer in Prior Work

Many of the positive results in the transfer-learning setting have appeared in the RRL literature, in environments from Blocks World (Dzeroski et al., 2001) to Freecraft (Guestrin et al., 2003). RRL uses relational or even full first-order languages to describe the state space. Although it is true that these richer representations facilitate the transfer of knowledge, the practicality of these measures is incumbent upon the abstraction techniques used to compact these potentially large representations. That is, in the relational extensions of value iteration (Boutilier et al., 2001; Großmann et al., 2002) and Q-learning (Dzeroski et al., 2001), a variant of ϕ_{Q^*} is used to keep the backups from expanding the state space far beyond the size needed to represent the value function.

Our algorithm for abstraction transfer is closely related to previous work on discovering irrelevant state variables with policy irrelevance (Jong & Stone, 2005). Our main contribution is extending this work to the complete range of abstractions discussed in Section 3.1 and developing theoretical guarantees for the behavior of traditional planning and learning algorithms in the induced abstract target domains, dependent on the choice of abstraction ϕ.

4. The GATA Algorithm

4.1. Motivation and Overview

In RRL, abstract value functions are often transferred to new domain instances, so in essence, both the value function and the abstraction are transferred. But, transferring the source values themselves often just bootstraps the value function in the target domain, where learning or planning algorithms must still be used to find an optimal policy. Previous work has indicated that transferring value functions may not be very general in the real world, as they depend critically on problem sizes (Dzeroski et al., 2001). Instead, we can induce a speedup by just transferring the learned
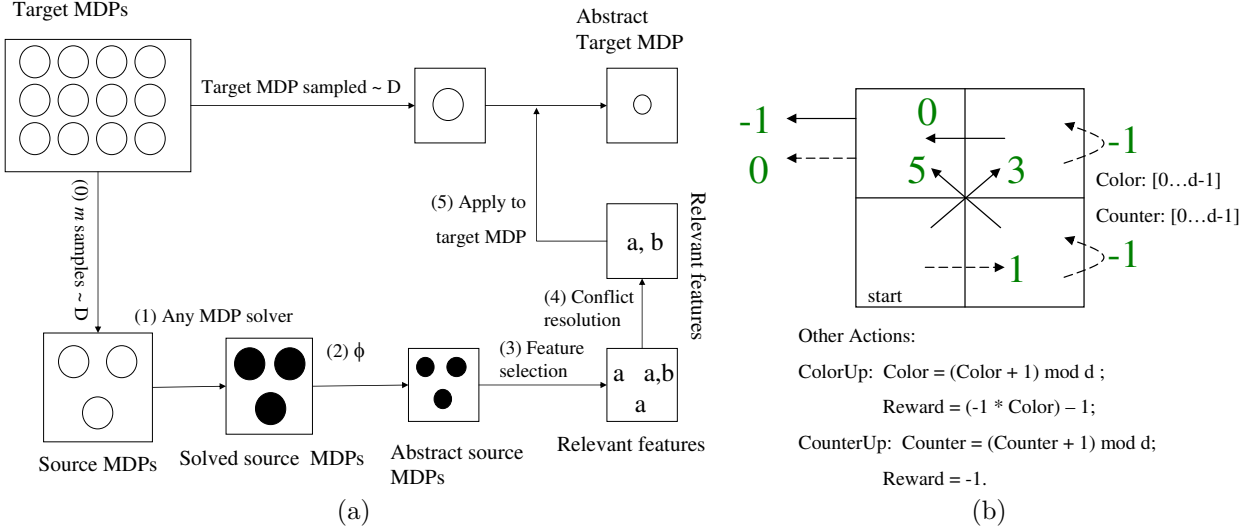
---

[1]We note that several of these examples are not the coarsest possible implementations of their respective abstractions, including those listed for ϕ_{Q^*} and ϕ_{π^*}.

Figure 1. (a) GATA: The General Abstraction Transfer Algorithm; (b) The example world used in section 5.

abstraction scheme from the source domains to the target, and then planning or learning in this smaller space. More formally, given a set of m source MDPs as described earlier, we wish to apply an abstraction technique to compact each source domain without sacrificing the ability to plan or learn. Then, we wish to unify these source abstractions and apply this unified abstraction to the target domain. The General Abstraction Transfer Algorithm (GATA), illustrated in Figure 1 (a) implements this abstraction transfer.

4.2. Implementation

In Step (1) the source MDPs are solved using a planning or learning algorithm. In Step (2), an abstraction scheme, such as ϕ_{Q^*}, is applied to each of the solved source domains to produce abstract source MDPs. We note that if Step (2) involves an abstraction scheme that does not require solving or learning the MDP (e.g. ϕ_{model} given the MDP) then Step (1) is not necessary, but value-function-based abstraction schemes often yield greater compaction (Li et al., 2006). In Step (3), we examine each of the abstract MDPs, and determine what features the applied abstraction scheme deems *relevant*. A feature f is considered relevant with respect to abstraction ϕ if and only if the abstraction that ignores f (ϕ_f) is finer than ϕ. That is, $\phi_f(s_1) = \phi_f(s_2)$ implies $\phi(s_1) = \phi(s_2)$. Notice we are finding the minimum features consistent with ϕ, not necessarily the minimum features needed to preserve the ability to learn or plan. In Step (4), we resolve conflicts over relevance between the multiple source domains in a "safe" manner where a feature is deemed "relevant" for the target domains if it was tagged as relevant in *any* of the source domains. In Step (5), the resulting feature-based abstraction function is used to

compact a target MDP. The applicability of the optimal policy learned in such reduced target domains is discussed below.

4.3. Discussion and Theoretical Results

As long as m (the number of source domains), is large, the following theorems hold (proofs omitted due to space constraints):

Theorem 3 *Replacing ϕ in GATA with abstractions ϕ_{model}, ϕ_{Q^π}, ϕ_{Q^*}, or ϕ_{a^*}, the optimal policy in the abstract target MDP, $\bar{\pi}^*$, is optimal in the ground target MDP. However, there exist examples with abstraction ϕ_π^* in place of ϕ where $\bar{\pi}^*$ is suboptimal in the ground target MDP.*

Theorem 4 *Using Q-learning in the abstract target MDP produced by GATA we have the following cases:*

1. *If ϕ is replaced by ϕ_{model}, ϕ_{Q^π}, or ϕ_{Q^*}, Q-learning in the resulting abstract target MDP converges to the optimal state-action value function and policy in the ground target MDP.*

2. *If ϕ is replaced by ϕ_{a^*}, Q-learning in the resulting abstract target MDP does not necessarily converge. However, it converges if the behavior policy is fixed. In either case, it yields a value function whose greedy policy is optimal in the ground target MDP.*

3. *If ϕ is replaced by ϕ_{π^*} Q-learning in the resulting abstract target MDP can converge to an state-action value function whose greedy policy is suboptimal in the ground target MDP. However, policy search methods may still be effective.*

These results help justify the admittedly daunting discovery times for ϕ_{Q^*} and ϕ_{a^*} (which involve determining the optimal value function for the ground source MDPs). Essentially, we are willing to pay the high price for discovery of ϕ_{Q^*} or ϕ_{a^*} if this provides information on what state variables (in the factored case) or relations (in the RRL case) can be ignored for planning or learning in the target domains.

5. An Example Domain

We demonstrate the power of GATA and the pitfalls promulgated in the previous two theorems using a simple grid world and three of our abstraction schemes (ϕ_{Q^*}, ϕ_{a^*}, and ϕ_{π^*}). The source instance is depicted in Figure 1 (b). All MDPs in this class have four state variables and four actions, but each instance has a different limit, d, on the dimensionality of Color and Counter. The actions CounterUP and ColorUP are available at any location. We denote the feature-based abstraction scheme developed by GATA with parameter ϕ as $A(\phi)$. The *relevant* variables uncovered by $A(\phi_{Q^*})$, $A(\phi_{a^*})$, and $A(\phi_{\pi^*})$ were [Color, Y, X], [Y, X] and [X] respectively. Since the targets can have d values for Color and Counter, the state space sizes for the abstractions are: none $O(d^2)$; $A(\phi_{Q^*})$ $O(d)$; $A(\phi_{a^*})$ $O(1)$; $A(\phi_{\pi^*})$ $O(1)$. Although $A(\phi_{\pi^*})$ actually gives the greatest compaction, when this abstraction ("only use X") is applied to a target domain, standard dynamic programming algorithms used on the abstract space will yield a policy that is suboptimal in the ground targets, forcing our speedup ratio to 0. In contrast, and consistent with Theorem 3, $A(\phi_{Q^*})$ and $A(\phi_{a^*})$ will yield abstraction schemes such that planning in the abstract target space yields a policy optimal for the ground target, and ensures our speedup ratio is no less than 1. More generally, we note that $A(\phi_{a^*})$ provides the smallest abstract models consistent with guaranteed optimal planning. Learning results are similar (Theorem 4).

6. Conclusions & Future Work

We have presented an algorithm for transferring abstractions learned in source MDPs to target MDPs. We also provided results concerning the soundness of planning and learning algorithms in the abstract target domains induced by this algorithm.

Studying the effectiveness of GATA based on the size of the sample set, m, is an area for future research. For example, when m is small, or the chosen target MDP is rare, the variables selected by GATA may not be sufficient for finding the optimal policy. Also, the feature-selection algorithm we outlined is very simplistic. Investigating the complexity and benefits of more

sophisticated feature-selection schemes is another avenue of future research. Replacing ϕ in GATA with an abstraction scheme that allows for inexact matching (e.g. Bisimulation metrics (Ferns et al., 2004)) could produce an abstraction that results in a bounded loss in the abstract value function for the target domains. Similarly, a less cautious combination of relevant features in Step (4) could result in a lossy but still generally effective abstraction scheme.

References

Boutilier, C., Dearden, R., & Goldszmidt, M. (2000). Stochastic dynamic programming with factored representations. *Artificial Intelligence*, *121*, 49–107.

Boutilier, C., Reiter, R., & Price, B. (2001). Symbolic dynamic programming for first-order MDPs. *IJCAI* (pp. 690–700).

Chapman, D., & Kaelbling, L. P. (1991). Input generalization in delayed reinforcement learning: An algorithm and performance comparisons. *IJCAI* (pp. 726–731).

Dzeroski, S., Raedt, L. D., & Driessens, K. (2001). Relational reinforcement learning. *Machine Learning*, *43*, 7–52.

Ferns, N., Panangaden, P., & Precup, D. (2004). Metrics for finite Markov decision processes. *UAI-04* (pp. 162–169).

Givan, R., Dean, T., & Greig, M. (2003). Equivalence notions and model minimization in Markov decision processes. *Artificial Intelligence*, *147*, 163–223.

Gordon, G. J. (1996). *Chattering in Sarsa(λ)* (Technical Report). CMU Learning Lab.

Großmann, A., Hölldobler, S., & Skvortsova, O. (2002). Symbolic dynamic programming within the fluent calculus. *IASTED International Conference Artificial and Computational Intelligence* (pp. 378–383).

Guestrin, C., Koller, D., Gearhart, C., & Kanodia, N. (2003). Generalizing plans to new environments in relational MDPs. *IJCAI-03*.

Jong, N. K., & Stone, P. (2005). State abstraction discovery from irrelevant state variables. *IJCAI-05*.

Li, L., Walsh, T. J., & Littman, M. L. (2006). Towards a unified theory of state abstraction for mdps. *9th International Symposium on Artificial Intelligence and Mathematics*.

McCallum, A. (1995). *Reinforcement learning with selective perception and hidden state*. Doctoral dissertation, University of Rochester, Rochester, NY.

Puterman, M. L. (1994). *Markov decision processes: Discrete stochastic dynamic programming*. New York: Wiley-Interscience.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.