

Inference strategies for solving semi-Markov decision processes

Matthew Hoffman and Nando de Freitas

Department of Computer Science
University of British Columbia, Canada

Abstract

Semi-Markov decision processes are used to formulate many control problems and also play a key role in hierarchical reinforcement learning. In this chapter we show how to translate the decision making problem into a form that can instead be solved by inference and learning techniques. In particular, we will establish a formal connection between planning in semi-Markov decision processes and inference in probabilistic graphical models, then build on this connection to develop an expectation maximization (EM) algorithm for policy optimization in these models.

Introduction

Researchers in machine learning have long attempted to join the fields of inference and learning with that of decision making. Influence diagrams, for example, explicitly cast the decision making process as inference in a graphical model (see e.g. Cooper, 1988; Shachter, 1988). However, while these methods are a straight-forward application of inference techniques they only apply to finite-horizon problems and only learn non-stationary policies.

For goal-directed decision problems, more general techniques such as that of Attias (2003) exist for finding the *maximum a posteriori* action sequence. (This technique was later extended by Verma & Rao (2006) to compute the *maximal probable explanation*.) It is crucial to note, however, that these approaches are not optimal in an expected reward sense. Instead, they can be interpreted as maximizing the probability of reaching the goal.

While it is well known in the optimal control literature that there exists a fundamental duality between inference and control for the special case of linear-quadratic Gaussian models (Kalman, 1960), this result does not hold in general. Extending these ideas to more general models has been attempted by locally approximating the optimal solution (see e.g. Toussaint, 2009; Todorov & Li, 2005).

A key step in realizing general inference-based approaches while still maintaining optimality with respect to expected rewards was originally addressed by Dayan & Hinton (1997) for immediate reward decision problems. In particular this work proposes an expectation maximization (EM) approach to the problem which works by optimizing a lower bound on the expected rewards. This technique was then greatly formalized by Toussaint & Storkey (2006) who extend it to the infinite-horizon case (see also Toussaint et al., 2006).

This line of research has since enjoyed substantial success in the field of robotics (Peters & Schaal, 2007; Kober & Peters, 2008; Vijayakumar et al., 2009), where empirical evidence has indicated that these methods can often outperform traditional stochastic planning and control methods as well as more recent policy gradient schemes.

The focus of this chapter is two-fold: to act as an introduction to the “planning as inference” methodology and to show how to extend these techniques to semi-Markov Decision Processes (SMDPs). SMDPs are an extension of the MDP formalism that generalize the notion of time—in particular, by allowing the time-intervals between state transitions to vary stochastically. This allows us to handle tradeoffs not only between actions based on their expected rewards, but also based on the amount of time that each action takes to perform.

SMDPs are interesting problems in their own right, with applications to call admission control and queueing systems (see e.g. Singh et al., 2007; Das et al., 1999). This formalism also serves as a natural platform in robotics for building complex motions from sequences of smaller motion “templates” as evidenced by Neumann et al. (2009). Finally, SMDPs are a crucial building block for hierarchical reinforcement learning methods (see e.g. Ghavamzadeh & Mahadevan, 2007; Sutton et al., 1998; Dietterich, 2000). While this chapter serves as an introductory text to the paradigm of inference and learning, and its application to SMDPs, we hope that future work in this area will leverage advances in structured inference techniques for hierarchical tasks of this nature.

The first section of this work will describe the basic mixture of MDPs model that we build on while the second section will show how to extend this to the SMDP formalism. We then describe an EM algorithm for solving these problems. Finally, in the last section we apply this approach to a small SMDP example.

A mixture of finite-time MDPs

Following the notation of Hoffman, de Freitas, et al. (2009) an MDP can be succinctly described via the following components:

- an initial state model $p(x_0)$,
- a state transition model $p(x_{n+1}|x_n, u_n)$,
- an immediate reward model $r(x_n, u_n)$,
- and finally a stochastic policy $\pi_\theta(u_n|x_n)$.

In this model, $n = 1, 2, \dots$ is a discrete-time index, $\{x_n\}$ is the state process, and $\{u_n\}$ is the action process. The model further assumes a randomized policy, but one can also easily adopt a deterministic policy $\pi_\theta(u|x) = \delta_{\phi_\theta(x)}(u)$, where δ denotes the Dirac function and ϕ is a deterministic mapping from states to actions. (By this same reasoning we can also encode knowledge of the initial state using a Dirac mass.) We will assume that the policy-parameters are real-valued, i.e. $\theta \in \mathbb{R}^d$.

Having defined the model, our objective is to maximize the expected future reward with respect to the parameters of the policy θ :

$$J(\theta) = \mathbb{E} \left[\sum_{n=0}^{\infty} \gamma^n r(x_n, u_n) | \theta \right], \quad (1)$$

where $0 < \gamma < 1$ is a discount factor. In order to ease notation later we will also note that for a given θ this model induces a Markov chain over state/action pairs $z_n = (x_n, u_n)$. The transition probabilities for this “extended state space” can then be written as

$$p_\theta(z_0) = p(x_0) \pi_\theta(u_0|x_0) \text{ and}$$

$$p_\theta(z_{n+1}|z_n) = p(x_{n+1}|x_n, u_n) \pi_\theta(u_{n+1}|x_{n+1}),$$

where the joint distribution over any finite k -length sequence of state/action pairs is defined as

$$p_\theta(z_{0:k}|k) = p_\theta(z_0) \prod_{n=1}^k p_\theta(z_n|z_{n-1}). \quad (2)$$

Finally, we will also write the rewards as $r(z) = r(x, u)$.

In order to transform the problem into one that is more amenable to inference methods we will first note that any maximum of $(1 - \gamma)J(\theta)$ is also a maximum of $J(\theta)$, as this extra multiplicative term just rescales the expected reward. Now, by expanding (1) we can write the (rescaled) expected reward as

$$(1 - \gamma)J(\theta) = (1 - \gamma) \int \left[p_\theta(z_0) \prod_{n=1}^{\infty} p_\theta(z_n|z_{n-1}) \right] \left[\sum_{k=0}^{\infty} \gamma^k r(z_k) \right] dz_{0:\infty}$$

and by exchanging the order of integration and summation we arrive at

$$\begin{aligned} &= \int (1 - \gamma) \gamma^0 p_\theta(z_0|k=0) r(z_0) dz_0 \\ &\quad + \int (1 - \gamma) \gamma^1 p_\theta(z_{0:1}|k=1) r(z_1) dz_{0:1} + \dots \\ &= \sum_{k=0}^{\infty} \int \underbrace{(1 - \gamma) \gamma^k}_{\text{time prior}} \underbrace{p_\theta(z_{0:k}|k)}_{\substack{\text{state/action} \\ \text{prior}}} r(z_k) dz_{0:k}. \end{aligned}$$

It is for this reason that the additional factor of $(1 - \gamma)$ was introduced. Under this formulation the discounting terms can be seen as a geometric distribution $p(k) = (1 - \gamma)\gamma^k$ and the expected reward under this random time-horizon can be written as

$$(1 - \gamma)J(\theta) = \mathbb{E}_{k, z_{0:k}} [r(z_k)|\theta], \quad (3)$$

$$p(k, z_{0:k}) = p(k) p(z_{0:k}|k).$$

As originally noted by Toussaint & Storkey (2006), we can now view this problem as an infinite mixture of finite horizon MDPs where rewards only occur at the end of a chain whose length is given by the random variable k . A diagram of this interpretation is shown in Figure 1. We must emphasize, however, that we have not changed the model nor the form of the expected rewards, but are instead departing from the standard *interpretation* of these types of decision problems. The classical approach to these problems is to take the expectation of an infinite-length trajectory and sum over increasingly discounted rewards (i.e. the rewards are worth less as time passes due to the discount factor). Instead we are

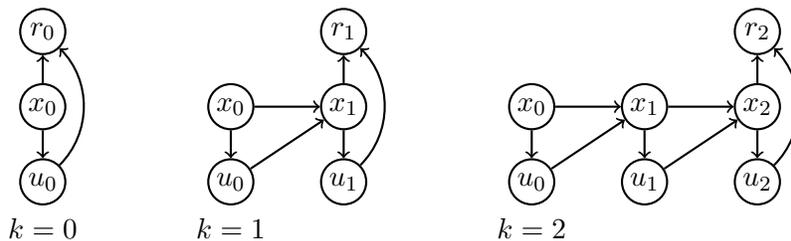


Figure 1. Illustration of the mixture of finite horizon MDPs described in Section . Expected rewards are computed by mixing over individual MDPs with probability $p(k)$ and taking reward $r_k = r(x_k, u_k)$.

taking the expectation with respect to a *finite-length* trajectory whose length is stochastic with probability equal to the discounting associated with this length. We then evaluate the expected reward only at the end of this chain, but by taking the expectation with respect to k we are essentially summing over the rewards at all such trajectory lengths. Note in particular that this is still an infinite-horizon problem!

This formulation is the basis for many inference and learning approaches to solving MDPs. In the next sections we will show how to extend this formulation in the context of SMDPs and will then describe an Expectation-Maximization (EM) algorithm for optimizing the parameters θ . We will first note, however, that these techniques are quite general and can be applied in both discrete¹ and continuous state-spaces (Toussaint & Storkey, 2006; Hoffman, de Freitas, et al., 2009). As we will briefly see later these methods can also be applied to situations where the models themselves are unknown and can only be sampled from (Kober & Peters, 2008; Vlassis & Toussaint, 2009). Finally, it is also possible to derive Markov chain Monte Carlo (MCMC) methods to optimize via sampling in parameter space. While we will omit discussion of these methods entirely we can point the interested reader towards (Hoffman et al., 2007; Hoffman, Kück, et al., 2009) for an introduction. This further enables the solution of planning problems using generic inference algorithms (see Goodman et al., 2008).

An extension to semi-MDPs

Formally we can define an SMDP as a continuous-time controlled stochastic process $z(t) = (x(t), u(t))$ consisting, respectively, of states and actions at every point in time t . In particular we will assume that the system transitions at random *arrival times* t_n and that the process is stationary in between jumps, i.e. $x(t) = x_n$ and $u(t) = u_n$ for all $t \in [t_n, t_{n+1})$. Here each of the state/action pairs evolves according to models defined in Section for MDPs. What makes this model a *semi*-MDP is the use of random transition times. In order to handle this generalization we will introduce random *sojourn times* $s_n > 0$ which represent the amount of time spent in the n th state. We will then assume the following distribution:

- a time model $p(s_n | x_n, u_n)$,
- where $t_n = t_{n-1} + s_n$ and $t_0 = 0$.

¹For discrete models the integrals become sums (i.e. integration with respect to the counting measure).

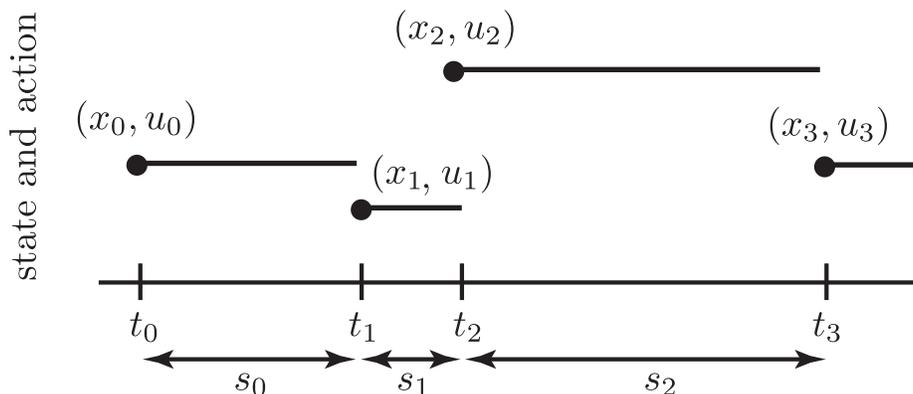


Figure 2. Relationship between arrival times t_n , sojourn times s_n , and the system state (x_n, u_n) .

Importantly, the conditional probability for sojourn times does not include the duration of the previous interval. See Figure 2 for an illustration of this process. More generally, we could also allow the sojourn times s_n to depend on the next state x_{n+1} . While the methods we will describe are fully capable of handling this situation, we will ignore this dependency for notational simplicity. Finally, we can write the joint probability over sequences of states/action pairs and sojourn times as

$$p_{\theta}(z_{0:k}, s_{0:k}) = p_{\theta}(z_0) p(s_0|z_0) \prod_{n=1}^k p(s_n|z_n) p_{\theta}(z_n|z_{n-1}) \quad (4)$$

for any choice of horizon k .

Just as in the standard MDP formulation we must also specify a reward function $r(z) = r(x, u)$ over states and actions. However, unlike in an MDP our discounting behaves differently in order to take into account the variable time in each state. In particular, we will discount the reward continuously over our entire trajectory, which because of the jump-Markov nature of our transitions will simplify to

$$\begin{aligned} J(\theta) &= \mathbb{E} \left[\int_0^{\infty} \beta e^{-\beta t} r(z(t)) dt \mid \theta \right] \\ &= \mathbb{E} \left[\sum_{n=0}^{\infty} e^{-\beta t_n} (1 - e^{-\beta s_n}) r(z_n) \mid \theta \right]. \end{aligned} \quad (5)$$

Here $\beta > 0$ is a discount rate which is analogous to the discount factor in an MDP. In fact, when the sojourn times are deterministically given by $s_n = 1$ we recover a standard infinite-horizon, discounted MDP with a discount factor of $\gamma = e^{-\beta}$. Those readers more familiar with continuous-time control problems may note that in the first line of the expected reward definition we have introduced an additional multiplicative term of β . We do this so that later we can give the discount factors a more intuitive interpretation as a probability distribution, and we point out that since this is a multiplicative constant it will not change the optimal policy parameters of the system.

Based on the intuition developed earlier for MDPs we can interpret the discount terms in (5) as a distribution over random time horizons k and write the following joint

distribution over paths and path lengths:

$$p_\theta(k, z_{0:k}, s_{0:k}) = e^{-\beta(s_0 + \dots + s_{k-1})} (1 - e^{-\beta s_k}) p_\theta(z_{0:k}, s_{0:k}). \quad (6)$$

This distribution is not, however, nearly as straightforward as that in the MDP case. Nor is there the simple division between path-lengths and paths.

Proposition 1. *The joint distribution introduced in (6) is properly defined and normalized, i.e. it integrates to 1.*

Proof. In theory we can integrate out each of the infinitely many trajectory and time variables to see that for any k the marginal over path lengths is given by

$$p_\theta(k) = \mathbb{E}[e^{-\beta s_0}] \dots \mathbb{E}[e^{-\beta s_{k-1}}] (1 - \mathbb{E}[e^{-\beta s_k}]).$$

Given this marginal we can consider an infinite sequence of non-identical Bernoulli trials, where $p_n = \mathbb{E}[e^{-\beta s_n}]$ is the probability that the n th trial fails. Note, we need not compute this quantity in practice.

Due to the restriction that sojourn times are greater than zero we know that each such quantity lies in the range $0 < p_n \leq 1$. With this in mind, we can then see that $p_\theta(k)$ is the probability that this sequence of trials has its first success after $k - 1$ failures. As a result the marginal distribution over K can be properly defined as a conjunction of Bernoulli random variables and thus the full joint (i.e. not integrating over the trajectory and time variables) must be similarly well defined and integrate to 1. \square

Finally, given the joint distribution $p_\theta(k, z_{0:k}, s_{0:k})$ we can rewrite our objective from (5) as the *expected final reward* under this distribution, i.e.

$$\begin{aligned} J(\theta) &= \int p_\theta(z_{0:\infty}, s_{0:\infty}) \left[\sum_{k=0}^{\infty} e^{-\beta t_k} (1 - e^{-\beta s_k}) r(z_k) \right] ds_{0:\infty} dz_{0:\infty} \\ &= \sum_{k=0}^{\infty} \int p_\theta(z_{0:k}, s_{0:k}) e^{-\beta t_k} (1 - e^{-\beta s_k}) r(z_k) ds_{0:k} dz_{0:k} \\ &= \sum_{k=0}^{\infty} \int p_\theta(k, z_{0:k}, s_{0:k}) r(z_k) ds_{0:k} dz_{0:k} \\ &= \mathbb{E}_{k, z_{0:k}} [r(z_k) | \theta]. \end{aligned} \quad (7)$$

Similar to the MDP case we have obtained this result by exchanging the order of integration and summation and pulling the discount factor into the previously introduced distribution from (6). In the next section we will take this one step further and treat the reward terms $r(z_k)$ as the likelihood of some “imaginary event”. We can then use this to develop an EM algorithm for finding the most likely policy parameters under this “data”.

Finally, as an aside, we should note how the formulations of this section simplify in the MDP case, and more importantly why these simplifications do not hold for SMDPs.

In particular, for an MDP it is not necessary to reason about times s_n (since these are deterministically set to 1) and the joint distribution (6) can be factorized as

$$p_\theta(k, z_{0:k}) = p(k) p_\theta(z_{0:k}|k).$$

Here the conditional distribution is given by (4) and the “time prior” $p(k)$ is given by a geometric distribution with success probability $\gamma = e^{-\beta}$, i.e. the constant discount factor. This factorization is what allowed Toussaint & Storkey (2006) to reformulate the infinite-horizon MDP problem as a mixture of finite-time MDPs, where the random variable k acts as an indicator variable. Unfortunately this interpretation does not hold in the case of more general SMDPs. By looking at the discount factors in (6) we can see that the probability of a specific trajectory length k is a function of all sojourn times $s_{0:k}$, and as a result the distribution over the random variable k depends on an infinite number of sojourn times. However, while the SMDP formalism does not have as clean of a probabilistic interpretation as MDPs, we can still apply this model by working directly with the joint distribution.

An EM algorithm for SMDPs

Expectation-Maximization (EM) is an algorithm formulation that is used to compute maximum likelihood estimates in the presence of unobserved or hidden variables. In our case, the unobserved variables consist of the trajectory length k along with the state/action pairs and their corresponding sojourn times, $z_{0:k}$ and $s_{0:k}$ respectively. The observed variables for this model are then implicitly given by the reward function—i.e. we are assuming some *imaginary* random variable was observed where the likelihood of this observation conditioned on our hidden variables is given by $r(z_k)$. Note, treating the rewards as likelihoods does place some restrictions on their allowable forms. While they need not sum to 1, they must be positive. However, for finite models or models which are bounded below this is easily obtainable by adding some constant term.

Given this interpretation we can introduce the following quantities which are necessary for deriving the EM algorithm:

- The *complete data likelihood* is the likelihood of both our observed and unobserved data; here given by $r(z_k) p_\theta(k, z_{0:k}, s_{0:k})$.
- The *incomplete data likelihood* is the integral of the complete data likelihood with respect to the hidden variables; here given by $\mathbb{E}[r(z_k)|\theta]$.
- Finally, the *predictive distribution* over the hidden variables is given by the “posterior” that takes into account both the prior and likelihood, and is thus the ratio of complete and incomplete likelihoods.

In particular, we will write the predictive distribution using the following notation:

$$\tilde{p}_\theta(k, z_{0:k}, s_{0:k}) = \frac{r(z_k) p_\theta(k, z_{0:k}, s_{0:k})}{\mathbb{E}[r(z_k)|\theta]}. \quad (8)$$

Given these quantities the EM algorithm is an iterative procedure which at iteration i computes the *expected* complete data log-likelihood, parameterized by θ , under the previous

iteration's policy parameters $\theta^{(i-1)}$. This quantity, which we will denote with $Q(\theta, \theta^{(i-1)})$, is then optimized with respect to the policy parameters in order to obtain $\theta^{(i)}$. We can summarize this procedure as:

$$\begin{aligned} Q(\theta, \theta^{(i-1)}) &= \mathbb{E}[\log\{p_\theta(k, z_{0:k}, s_{0:k}) r(z_k)\} | \theta^{(i-1)}], & \text{(E-step)} \\ \theta_i &= \arg \max_{\theta} Q(\theta, \theta^{(i-1)}). & \text{(M-step)} \end{aligned}$$

This EM procedure is known to locally maximize the incomplete data likelihood (see Dempster et al., 1977; McLachlan & Krishnan, 1997) and hence will maximize our original objective $J(\theta)$.

Before deriving the E-step in more detail we will first take a brief look at the quantities that are actually needed in order to perform the maximization in the M-step. We can first let θ' denote the previous iteration's parameter estimate and rewrite the Q -function as

$$\begin{aligned} Q(\theta, \theta') &= \mathbb{E}[\log\{\prod_{n=0}^k \pi_\theta(z_n)\} | \theta'] + \\ &\quad \mathbb{E}[\log\{r(z_k) p(x_0) \prod_{n=1}^k p(x_n | z_{n-1})\} | \theta'] + \\ &\quad \mathbb{E}[\log\{\prod_{n=0}^k p(s_n | z_n)\} | \theta'] \end{aligned}$$

where we have expanded the complete data likelihood and separated those terms which do and do not depend on θ . Since only the first of these three quantities depends on θ we can drop the others and expand this expectation:

$$\begin{aligned} &= \sum_{k=0}^{\infty} \int \left[\sum_{n=0}^k \log \pi_\theta(z_n) \right] \tilde{p}_{\theta'}(k, z_{0:k}) dz_{0:k} + \text{const.} \\ &= \sum_{k=0}^{\infty} \sum_{n=0}^k \int \log \pi_\theta(z_n) \tilde{p}_{\theta'}(k, z_n) dz_n + \text{const.} \end{aligned} \tag{9}$$

Finally, in order to optimize the Q -function with respect to the policy parameters θ we will also need to evaluate the gradient

$$\nabla_{\theta} Q(\theta, \theta') = \sum_{k=0}^{\infty} \sum_{n=0}^k \int \tilde{p}_{\theta'}(k, z_n) \nabla_{\theta} \log \pi_{\theta}(z_n) dz_n. \tag{10}$$

Ultimately, this expansion informs how we will derive the steps required for the EM algorithm. In the E-step we need to construct the distribution $\tilde{p}_{\theta'}(k, z_n)$ and in the M-step we will compute the expectation of the gradient $\nabla \log \pi_{\theta}(z_n)$ under this distribution. We should note that this is the same form of the marginal distribution that would be computed in the E-step for a standard MDP. We will see shortly, however, that the computations necessary to compute this distribution are different due the integration over sojourn-times. In the next two subsections we will discuss these steps in more detail.

The E-step

As noted in the previous section, we need to construct the marginals $\tilde{p}_\theta(k, z_n)$ in order to compute the gradient of the expected complete log-likelihood. In this section, we will derive an efficient method for recursively constructing this distribution. We start by writing the marginal distribution as the integral of the predictive distribution with respect to all those terms other than k and z_n ,

$$\tilde{p}_\theta(k, z_n) = \int \tilde{p}_\theta(k, z_{0:k}, s_{0:k}) dz_{0:n-1} dz_{n+1:k} ds_{0:k}.$$

This integral can then be broken into those components that come before and after n respectively, and we can then see that the marginal distribution is proportional to

$$\begin{aligned} \tilde{p}_\theta(k, z_n) \propto & \int e^{-\beta(s_0+\dots+s_{n-1})} p_\theta(z_{0:n}, s_{0:n-1}) dz_{0:n-1} ds_{0:n-1} \quad \times \\ & \int e^{-\beta(s_n+\dots+s_{k-1})} (1 - e^{-\beta s_k}) r(z_k) p_\theta(z_{n+1:k}, s_{n:k} | z_n) dz_{n+1:k} ds_{n:k}. \end{aligned} \quad (11)$$

We should emphasize the fact that we are *not* integrating over z_n , which enables us to break the original integral into two independent integrals. Here we have also omitted the constant of proportionality, given by the expected reward $\mathbb{E}[r(z_k)]$.

In an analogy to the forward-backward algorithm for hidden Markov models we will call the two sets of integrals from (11) “forward” and “backward” messages, denoting them as $\alpha_\theta(z_n)$ and $\beta_\theta(z_n | \tau = k - n)$ respectively². Using these messages we can then write the marginal distribution as

$$\tilde{p}_\theta(k, z_n) = \frac{1}{\mathbb{E}[r(z_k)]} \alpha_\theta(z_n) \beta_\theta(z_n | k - n). \quad (12)$$

Intuitively the forward messages are integrating information forward in time from the initial-state distribution where the n th such message is the state/action distribution at step n weighted by the expected discounting up to step $n - 1$. Meanwhile, the backward messages are integrating information backward from the rewards and can be seen as the expected reward at step k given the state/action pair at step n (weighted by the expected discounting between n and k). It is crucial to note, however, that these messages are not probability distributions due to the way the discount factors have been split between the forward and backward components, namely:

$$\underbrace{e^{-\beta(s_0+\dots+s_{n-1})}}_{\text{forward}} \underbrace{e^{-\beta(s_n+\dots+s_{k-1})} (1 - e^{-\beta s_k})}_{\text{backward}}.$$

This causes no technical (or conceptual) difficulties, though, because when combined in (12) these messages form the desired probability distribution. This is similar in spirit to

²The notation here is slightly confusing in that we have a term β denoting the continuous discount factor and $\beta_\theta(\cdot | \tau)$ denoting the backward messages. This confusion, however, seems unavoidable as both of these terms are unanimously used in their respective literatures. To somewhat alleviate this confusion we note that the backward messages are *always* subscripted.

techniques used to maintain numerical stability when working with hidden Markov models (Bishop, 2006).

At this point, we can also more fully describe the use of the τ term in defining the backward messages. The reason behind this notation stems from the fact that naively computing these messages for each choice of k and n turns out to involve a great deal of redundant computation. Under the predictive distribution defined in (8) the reward factors only interact with the end of a finite-length trajectory. As a result the backward messages depend only on the difference $\tau = k - n$, i.e. how far in the future the rewards are obtained. Because of this, we can instead define the backward messages purely in terms of the “time-to-go”. This notation was originally presented by Toussaint & Storkey (2006), but here the messages are generalized to account for the fact that in the semi-Markov setting the k th reward introduces a factor over the state and action z_n and duration s_n at every epoch n prior to k .

Finally, by integrating the components of the first integral in (11) successively we can recursively define the forward messages as

$$\begin{aligned} \alpha_\theta(z_0) &= \mu(x_0) \pi_\theta(u_0|x_0), \\ \alpha_\theta(z_n) &= \int \alpha_\theta(z_{n-1}) p_\theta(z_n|z_{n-1}) dz_{n-1} \times \int e^{-\beta s_{n-1}} p_\theta(s_{n-1}|z_{n-1}) ds_{n-1}. \end{aligned} \quad (13)$$

Here we can see that we have the standard MDP forward message recursion multiplied by an additional integral due to the sojourn times. Similarly, we can recursively define the backwards messages as

$$\begin{aligned} \beta_\theta(z_n|0) &= r(z_n) \int (1 - e^{-\beta s_n}) p_\theta(s_n|z_n) ds_n, \\ \beta_\theta(z_n|\tau) &= \int \beta_\theta(z_{n+1}|\tau - 1) p_\theta(z_{n+1}|z_n) dz_{n+1} \times \int e^{-\beta s_n} p_\theta(s_n|z_n) ds_n. \end{aligned} \quad (14)$$

Again we can see that these messages can be seen as two integrals, one corresponding to the standard MDP message and a sojourn time message. Given the format of these two messages we can further introduce what we call an “expected discount factor”

$$\gamma(z) = \int e^{-\beta s} T(s|z) ds \quad (15)$$

which corresponds to the integral over sojourn times noted above. We can consider this term as a generalization of the MDP formalism wherein discount factors are no longer constant and instead depend on the current state and the action taken from that state. Further, we can see that for any exponential-family distribution this integral will exist in closed form.

The M-step

The M-step requires us to maximize the Q -function with respect to the policy parameters θ . If possible we can analytically maximize this function by solving for the fixed point of $\nabla_\theta Q(\theta, \theta') = 0$. If this is not possible we can still evaluate the gradient at the current set of policy parameters $\nabla_\theta Q(\theta', \theta')$ and follow the resulting ascent direction, resulting in a generalized EM algorithm (GEM). When this procedure is iterated, both of these methods

are known to locally maximize the incomplete data likelihood (again, for more details see McLachlan & Krishnan, 1997).

While EM methods are, in general, only able to guarantee local convergence it can be shown via its relation to policy iteration that these methods exhibit global convergence for discrete models when using exact/analytic inference (see Toussaint et al. (2006) for more details). In more general continuous settings no such guarantees can be made, however as noted by Hoffman, de Freitas, et al. (2009) a sufficiently exploratory initial policy does seem to have a tempering effect. This is especially true if the exact EM updates can be used, as additional exploratory noise does not cause the dramatic increase in variance associated with sample-based methods (such as policy gradients).

Monte Carlo EM

It is also possible to perform a Monte Carlo approximation during the E-step in order to optimize these problems when either the necessary distributions are unknown or the updates cannot be computed in closed form. As is derived for MDPs in (Kober & Peters, 2008; Vlassis & Toussaint, 2009), we can sample from the initial-state and transition distributions in order to approximate the Q -function. Given M trajectories $\{z_{0:k}^{(i)}, s_{0:k}^{(i)}\}_{i \leq M}$ sampled from $p_{\theta'}(z_{0:k}, s_{0:k})$ we can approximate the joint distribution for any $n < k$ with

$$\tilde{p}_{\theta'}(k, z_n) \approx \frac{1}{MA} \sum_{i=1}^M \left[e^{-\beta t_k^{(i)}} (1 - e^{-\beta s_k^{(i)}}) \right] r(z_k^{(i)}) \delta_{z_n^{(i)}}(z_n),$$

where A is a proportionality constant, given by $\mathbb{E}[r(z_k)]$ as noted earlier. If we assume some maximum time-horizon K_{\max} we can approximate the Q -function as

$$\begin{aligned} Q(\theta, \theta') &\approx \sum_{i=1}^M \sum_{k=0}^{K_{\max}} \sum_{n=0}^k \left[e^{-\beta t_k^{(i)}} (1 - e^{-\beta s_k^{(i)}}) \right] r(z_k^{(i)}) \log \pi_{\theta}(z_n^{(i)}) \\ &= \sum_{i=1}^M \sum_{n=0}^{K_{\max}} \log \pi_{\theta}(z_n^{(i)}) \sum_{k=n}^{K_{\max}} \left[e^{-\beta t_k^{(i)}} (1 - e^{-\beta s_k^{(i)}}) \right] r(z_k^{(i)}). \end{aligned}$$

This function can then be optimized using the same techniques as in the standard M-step.

Discrete models with Gamma-distributed time

The methods presented in previous sections, as with all EM-based procedures, provide a “meta-algorithm” which depends upon the exact models in use. In this section we present a simple model for the purposes of illustrating the procedure. Again, for a more advanced treatment of these methods in the context of MDPs we refer the reader to (Hoffman, de Freitas, et al., 2009; Toussaint et al., 2006; Toussaint, 2009; Kober & Peters, 2008).

We will now consider a model where the states and actions are discrete and sojourn times are given by a Gamma distribution. While simple, this model nonetheless presents an interesting scenario for planning and control domains because it can naturally be extended to cases when we want to reason about more complex distributions over the time to complete

an action. In our experiments, we define the discrete transition, initial-state, and reward models according to

$$\begin{aligned} \mu_x &= p(x_0 = x), \\ P_{xux'} &= p(x_{n+1} = x' | x_n = x, u_n = u), \text{ and} \\ R_{xu} &= r(x, u) \end{aligned}$$

where the sojourn times are Gamma-distributed according to

$$p(s_n | x_n = x, u_n = u) = \Gamma(s_n; k_{xu}, \sigma_{xu}).$$

Finally, we will assume a discrete policy where $\theta_{xu} = \pi_\theta(u|x)$.

Under this formulation the forward and backward messages will be representable as matrices, and by dropping the θ index we will let α_{xu}^n and β_{xu}^τ denote the n -step forward message and τ -step backward messages respectively. By plugging initial state, transition, and reward matrices into (13) and (14) we can explicitly write these messages as

$$\alpha_{xu}^n = \theta_{xu} \sum_{x', u'} \alpha_{x'u'}^{n-1} P_{x'u'x} \gamma_{x'u'}, \quad \alpha_{xu}^0 = \theta_{xu} \mu_x; \quad (16)$$

$$\beta_{xu}^\tau = \gamma_{xu} \sum_{x', u'} \beta_{x'u'}^{\tau-1} P_{xux'} \theta_{x'u'}, \quad \beta_{xu}^0 = R_{xu}(1 - \gamma_{xu}). \quad (17)$$

To make notation easier we will also introduce a forward message defined only over states, $\bar{\alpha}_x = \sum_u \alpha_{xu}$. In this setting the expected discount factor noted earlier can be written as the matrix

$$\begin{aligned} \gamma_{xu} &= \int \Gamma(s; k_{xu}, \sigma_{xu}) e^{-\beta s} ds \\ &= \int s^{k_{xu}-1} \frac{\exp(-(\beta + \sigma_{xu}^{-1})s)}{\Gamma(k_{xu}) \sigma_{xu}^{k_{xu}}} ds = (1 + \beta \sigma_{xu})^{-k_{xu}}. \end{aligned}$$

This particular form arises purely from the use of Gamma-distributed sojourn times, and in fact we can imagine extending this to continuous spaces using functions $k(x, u)$ and $\sigma(x, u)$.

At every iteration, for a given set of parameters θ' , the E-step consists of calculating the forward and backward messages given by Equations (16,17). By plugging these terms into the Q -function defined in (9) we can write

$$\begin{aligned} Q(\theta, \theta') &= \frac{1}{\mathbb{E}[r(z_k)]} \sum_{k=0}^{\infty} \sum_{n=0}^k \sum_{u,x} (\log \theta_{xu}) \theta'_{xu} \bar{\alpha}_x^n \beta_{xu}^{k-n} \\ &= \frac{1}{\mathbb{E}[r(z_k)]} \sum_{u,x} (\log \theta_{xu}) \theta'_{xu} \left[\sum_{n=0}^{\infty} \bar{\alpha}_x^n \right] \left[\sum_{\tau=0}^{\infty} \beta_{xu}^\tau \right], \end{aligned}$$

where the second equality can be obtained by rearranging the sums over k and n . This alternate formulation is particularly useful in discrete models where the sum over forward and backward messages can be expressed as finite quantities, i.e. a vector and a matrix respectively. Further, given this formulation we can optimize the Q -function for each state x

individually, which is possible because in discrete domains we can find the optimal action to take for each state regardless of the probability of visiting that state. By taking the gradient of the log-policy $\nabla \log \pi_\theta(u|x) = \theta_{xu}^{-1}$ and solving $\nabla Q(\theta, \theta') = 0$ for θ , subject to the constraint that $\sum_u \theta_{xu} = 1$ for each x , we arrive at the following solution methods:

$$\theta_{xu} \propto \theta'_{xu} \sum_{\tau=0}^{\infty} \beta_{xu}^\tau, \quad (\text{EM})$$

$$\theta_{xu} = \delta_{m(x)}(u) \quad \text{where } m(x) = \arg \max_{u'} \sum_{\tau=0}^{\infty} \beta_{xu'}^\tau. \quad (\text{greedy-EM})$$

Here the EM solution is performing exactly the optimization described above, while the greedy-EM solution, however, myopically chooses for each state x the one action u that maximizes the total future rewards when taken from that state. In particular, the greedy solution can be seen as iterations which correspond to repeated M-steps which skip intermediate E-steps as is shown by Toussaint et al. (2006), and in this sense this method is equivalent (again only for discrete models) to policy iteration. In larger discrete models, however, the EM approach has the advantage over policy iteration in that it is possible to prune computation (by using the forward messages α) in a principled manner; see Toussaint et al. (2006) for more details.

We first test these algorithms on a small, 16-state, 5-action domain with randomly generated transition and sojourn parameters, as well as a randomly generated reward model. The principal challenge of this domain, over other discrete domains, is to take advantage of the structure in the sojourn time distribution. The top-left plot of Figure 3 displays convergence properties of the described algorithms as well as a comparison to a standard policy-gradient method (see e.g. Baxter & Bartlett, 2001). In particular we should note that the resulting model was densely connected which allows for quick travel across the space, and explains the very good performance of the stochastic policy gradient algorithm. Also shown for policy gradients are error-bars corresponding to one standard deviation. The other methods don't need error bars because they are deterministic.

Building upon these results, the top-right plot shows the algorithms learning in a more structured environment. In particular the model used has grid-structured transitions on a small 4-by-4 grid. This model is especially interesting because we specify different Gamma-distributions for the inner nodes than the outer nodes such that the inner nodes move much more slowly. Also, we use a sparse reward model where most states have negligible reward and one state has high reward. The most important thing to note from this sub-figure is that the policy gradient method starts to break down under this sparse transition and reward model, even though the size of the state and action spaces are the same as in the previous example.

Lastly the bottom-left plot of this figure displays the progress of these algorithms on a much larger 20-by-20 grid, i.e. one in which there are 2000 state/action pairs. Similar to the previous example there is a single (relatively) large reward in the upper right corner of the grid and inner nodes with much slower sojourn times. Here we see that the EM algorithms vastly out-perform the policy gradient method and the learned policy successfully skirts the outside edge of the state-space in order to most quickly get to the high reward. Here the policy gradient method has relatively low-variance because it is not able to make any progress (i.e. it is stuck exploring a plateau with very little gradient information).

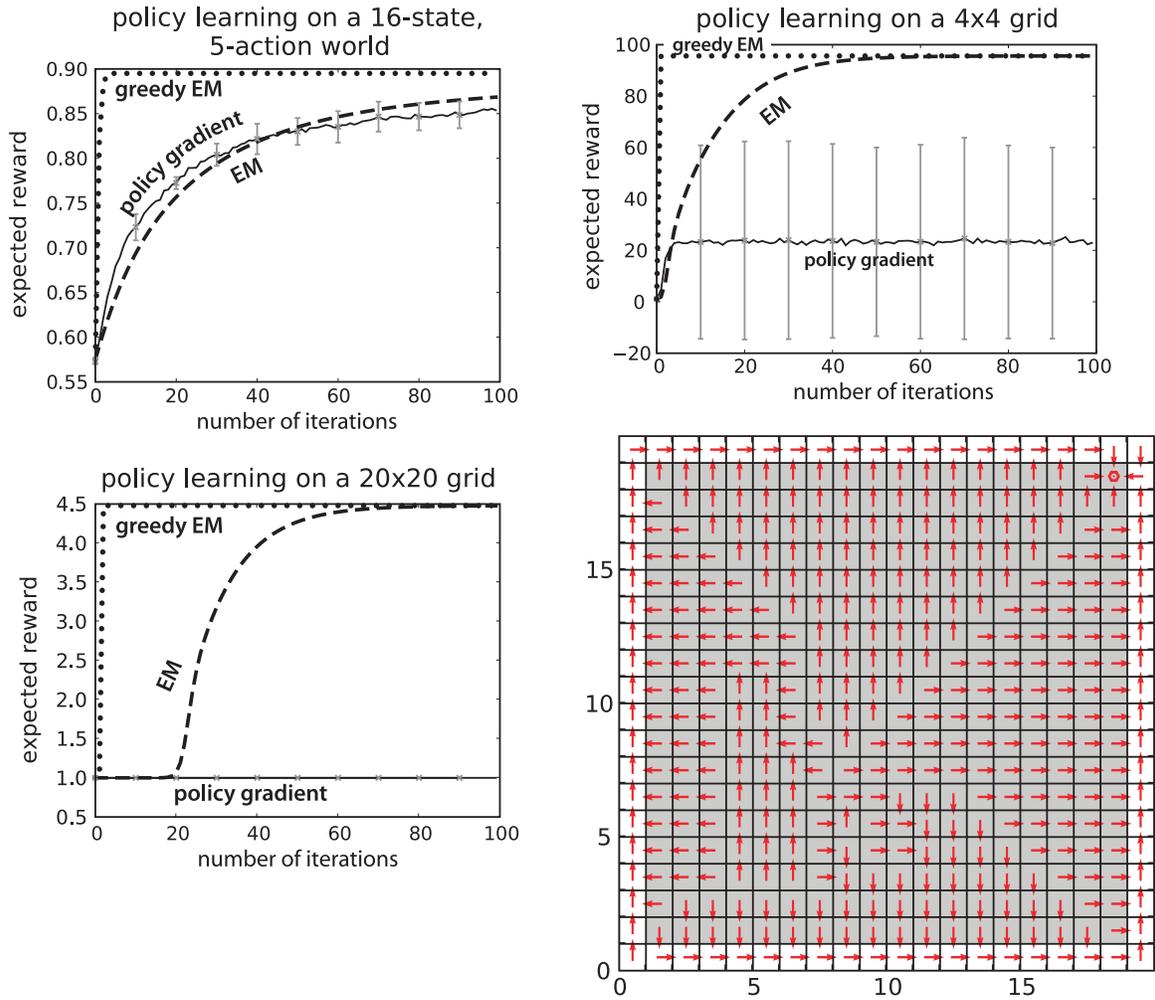


Figure 3. Results on various discrete models. The top-left plot shows convergence of the described algorithms on a randomly generated, dense model; the top-right plot shows performance on a model of the same size but with grid-like transitions and more structured transition times. The bottom-left plot shows performance on a larger grid domain. The bottom-right shows the learned policy in the larger domain. For both grid domains there is a single large reward (denoted with a dot) and transitions taken outside of the edge states (denoted with grey) have a larger expected sojourn time.

Conclusions

In this chapter, we have shown how it is possible to design effective planning algorithms in continuous-time domains by framing the policy optimization problem in an SMDP as one of inference in a related probabilistic graphical model. The connection between the reinforcement learning and probabilistic inference domains allows us to exploit existing developments in exact and approximate inference, and will perhaps provide us with leverage for tackling pressing problems in hierarchical reinforcement learning.

Of particular interest are inference approaches which attempt to exploit the structure of the underlying models. Preliminary work on applying such techniques to control problems includes the pruning steps of Toussaint et al. (2006). As evidenced by the small example in this chapter as well as larger continuous models in Hoffman, de Freitas, et al. (2009) we can see large performance gains by using exact inference methods. Another promising area of research would be the combination of these ideas with sample-based methods such as Kober & Peters (2008), perhaps via Rao-Blackwellization (de Freitas et al., 2004).

Acknowledgements

We would like to thank Peter Carbonetto and Arnaud Doucet for helpful contributions to an earlier version of this report. We would also like to thank Peter for his help in generating Figure 2. Finally, we would like to thank the editors and reviewers of this chapter for their helpful comments.

References

- Attias, H. (2003). Planning by probabilistic inference. In *Proceedings of the international conference on artificial intelligence and statistics*.
- Baxter, J., & Bartlett, P. (2001). Infinite-horizon policy-gradient estimation. *Journal of Artificial Intelligence Research*, 15(4).
- Bishop, C. (2006). *Pattern recognition and machine learning*. Springer-Verlag.
- Cooper, G. (1988). A method for using belief networks as influence diagrams. In *Proceedings of the conference on uncertainty in artificial intelligence*.
- Das, T., Gosavi, A., Mahadevan, S., & Marchallick, N. (1999). Solving semi-Markov decision problems using average reward reinforcement learning. *Management Science*, 45(4).
- Dayan, P., & Hinton, G. (1997). Using expectation-maximization for reinforcement learning. *Neural Computation*, 9(2).
- de Freitas, N., Dearden, R., Hutter, F., Morales-Menendez, R., Mutch, J., & Poole, D. (2004). Diagnosis by a waiter and a Mars explorer. *Proceedings of the IEEE, special issue on sequential state estimation*, 92(3).
- Dempster, A., Laird, N., & Rubin, D. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society, Series B (Statistical Methodology)*, 39(1).
- Dietterich, T. (2000). Hierarchical reinforcement learning with the MAXQ value function decomposition. *Journal of Artificial Intelligence Research*, 13(1).

- Ghavamzadeh, M., & Mahadevan, S. (2007). Hierarchical average reward reinforcement learning. *The Journal of Machine Learning Research*, 8.
- Goodman, N., Mansinghka, V., Roy, D., Bonawitz, K., & Tenenbaum, J. (2008). Church: a language for generative models. In *Proceedings of the conference on uncertainty in artificial intelligence*.
- Hoffman, M., de Freitas, N., Doucet, A., & Peters, J. (2009). An expectation maximization algorithm for continuous Markov decision processes with arbitrary reward. In *Proceedings of the international conference on artificial intelligence and statistics*.
- Hoffman, M., Doucet, A., de Freitas, N., & Jasra, A. (2007). Bayesian policy learning with trans-dimensional MCMC. In *Advances in neural information processing systems*.
- Hoffman, M., Kück, H., de Freitas, N., & Doucet, A. (2009). New inference strategies for solving Markov decision processes using reversible jump MCMC. In *Proceedings of the conference on uncertainty in artificial intelligence*.
- Kalman, R. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1).
- Kober, J., & Peters, J. (2008). Policy search for motor primitives in robotics. In *Advances in neural information processing systems*.
- McLachlan, G., & Krishnan, T. (1997). *The EM algorithm and extensions*. Wiley New York.
- Neumann, G., Maass, W., & Peters, J. (2009). Learning complex motions by sequencing simpler motion templates. In *Proceedings of the international conference on machine learning*.
- Peters, J., & Schaal, S. (2007). Reinforcement learning for operational space control. In *Proceedings of the IEEE international conference on robotics and automation*.
- Shachter, R. (1988). Probabilistic inference and influence diagrams. *Operations Research*.
- Singh, S., Tadic, V., & Doucet, A. (2007). A policy gradient method for semi-Markov decision processes with application to call admission control. *European Journal of Operational Research*, 178(3), 808–818.
- Sutton, R., Precup, D., & Singh, S. (1998). Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1).
- Todorov, E., & Li, W. (2005). A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems. In *American control conference, 2005. proceedings of the 2005*.
- Toussaint, M. (2009). Robot trajectory optimization using approximate inference. In *Proceedings of the international conference on machine learning*.
- Toussaint, M., Harmeling, S., & Storkey, A. (2006). *Probabilistic inference for solving (PO)MDPs* (Tech. Rep. No. EDI-INF-RR-0934). University of Edinburgh, School of Informatics Research.
- Toussaint, M., & Storkey, A. (2006). Probabilistic inference for solving discrete and continuous state Markov decision processes. In *Proceedings of the international conference on machine learning*.
- Verma, D., & Rao, R. (2006). Planning and acting in uncertain environments using probabilistic inference. In *Intelligent robots and systems*.

- Vijayakumar, S., Toussaint, M., Petkos, G., & Howard, M. (2009). Planning and moving in dynamic environments: A statistical machine learning approach. In *Creating brain like intelligence: From principles to complex intelligent systems*. Springer-Verlag.
- Vlassis, N., & Toussaint, M. (2009). Model-free reinforcement learning as mixture learning. In *Proceedings of the 26th international conference on machine learning*.