

Temporal Logic Motion Control using Actor-Critic Methods ^{*}

Jing Wang[†], Xuchu Ding[‡], Morteza Lahijanian[§]

Ioannis Ch. Paschalidis[†], and Calin A. Belta[†]

December 6, 2014

Abstract

This paper considers the problem of deploying a robot from a specification given as a temporal logic statement about some properties satisfied by the regions of a large, partitioned environment. We assume that the robot has noisy sensors and actuators and model its motion through the regions of the environment as a Markov Decision Process (MDP). The robot control problem becomes finding the control policy which maximizes the probability of satisfying the temporal logic task on the MDP. For a large environment, obtaining transition probabilities for each state-action pair, as well as solving the necessary optimization problem for the optimal policy, are computationally intensive. To address these issues, we propose an approximate dynamic programming framework based on a least-square temporal difference learning method of the actor-critic type. This framework operates on sample paths of the robot and optimizes a randomized control policy with respect to a small set of parameters. The transition probabilities are obtained only when needed. Simulations confirm that convergence of the parameters translates to an approximately optimal policy.

1 Introduction

One major goal in robot motion planning and control is to specify a mission task in an expressive and high-level language and to convert the task automatically to a control strategy for the robot. The robot is subject to mechanical constraints, actuation and measurement noise, and limited communication and sensing capabilities. The challenge in this area is the development of a computationally efficient framework accommodating both the robot constraints and the uncertainty of the environment, while allowing for a large spectrum of task specifications.

In recent years, temporal logics such as Linear Temporal Logic (LTL) and Computation Tree Logic (CTL) have been promoted as formal task specification languages for robotic applications (Kress-Gazit et al., 2007; Karaman and Frazzoli, 2009; Loizou and Kyriakopoulos, 2004; Quottrup et al., 2004; Wongpiromsarn et al., 2009; Bhatia et al., 2010). They are appealing due to their high

^{*}Research partially supported by the NSF under grants CNS-1239021 and IIS-1237022, by the ARO under grants W911NF-11-1-0227 and W911NF-12-1-0390, and by the ONR under grants N00014-09-1051 and N00014-10-1-0952.

[†]Jing Wang, Ioannis Ch. Paschalidis, and Calin A. Belta are with the Division of System Eng., Dept. of Mechanical Eng., and Dept. of Electrical & Computer Eng., Boston University, Boston, MA 02215 (<{wangjing, yannisp, cbelta}@bu.edu>), respectively.

[‡]Xuchu Ding is with the Embedded Systems and Networks group, United Technologies Research Center, East Hartford, CT 06108 (dingx@utrc.utc.com).

[§]Morteza Lahijanian is with Dept. of Computer Science, Rice University, Houston, TX, 77005, (morteza@rice.edu).

expressivity and closeness to human language. (Reviewer 2: the authors claim that LTL and CTL are close to natural language. This reviewer believes this statement to be untrue. What is true is that these languages are often able to formally capture informal requirements specified in natural language.) Moreover, several existing formal verification (Clarke et al., 1999; Baier et al., 2008) and synthesis (Baier et al., 2008; Liu et al., 2013; Luna et al., 2014) tools can be adapted to generate motion plans and provably correct control strategies for the robots.

In this paper, we assume that the motion of the robot in the environment is described by a (finite) Markov Decision Process (MDP). In this model, the robot can precisely determine its current state, and by applying an action (corresponding to a motion primitive) enabled at each state, it triggers a transition to an adjacent state with a fixed probability. We are interested in controlling the robot such that it maximizes the probability of satisfying a temporal logic formula over a set of properties satisfied at the states of the MDP.

By adapting existing probabilistic model checking (Baier et al., 2008; De Alfaro, 1997; Vardi, 1999) and synthesis (Courcoubetis and Yannakakis, 1990; Baier et al., 2004) algorithms, we (Ding et al., 2011; Lahijanian et al., 2010) and others (Wolff et al., 2012) recently developed such computational frameworks for formulae of LTL and a fragment of probabilistic CTL. With the above approaches, an optimal control policy can be generated to maximize the satisfaction probability, given that the transition probabilities are known for each state-action pair of the MDP, which can be computed by using a Monte-Carlo method and repeated forward simulations.

In real applications, the size of the state space of the MDP is usually very large. Our previous approaches are not suitable for large-sized problems due to the following limitations. First, they require transition probabilities for all state-action pairs, which are costly to obtain even if an accurate simulator of the robot in the environment is available. Second, the optimal policy is calculated by solving a large Linear Programming (LP) problem on the product between the original MDP and a Rabin automaton. The existing LP solvers are not efficient enough and their memory usages increase rapidly with problem size.

In this paper, we show that approximate dynamic programming (Si, 2004; Bertsekas and Tsitsiklis, 1996) can be effectively used to address the above limitations. For large dynamic programming problems, an approximately optimal solution can be provided using actor-critic algorithms (Barto et al., 1983). By approximating both the policy and state-action value function with a parameterized structure, actor-critic algorithms use much less memory compared with other dynamic programming techniques. In particular, actor-critic algorithms with Least Squares Temporal Difference (LSTD) learning have been shown recently to be a powerful tool for large-sized problems(Moazzez Estanjini et al., 2011b; Estanjini et al., 2012; Konda and Tsitsiklis, 2003).

In (Ding et al., 2011), the authors show that a motion control problem with temporal logic specifications could be converted to a maximal reachability probability (MRP) problem, *i.e.*, maximizing the probability of reaching a set of states. In (Moazzez Estanjini et al., 2011b), we show that the MRP problem is equivalent to a Stochastic Shortest Path (SSP) problem and propose an actor-critic method to solve the SSP problem. In (Ding et al., 2012), we apply the actor-critic method in (Moazzez Estanjini et al., 2011b) to find a control policy that maximizes the probability of satisfying a temporal logic specification. Our proposed algorithm produces a *Randomized Stationary Policy* (RSP), which gives a probability distribution over enabled actions at a state. Our method requires transition probabilities to be generated only along sample paths, and is therefore particularly suitable for robotic applications. To the best of our knowledge, this is the first attempt to combine temporal logic formal synthesis with actor-critic type methods.

This paper is based on the work of (Moazzez Estanjini et al., 2011b) and (Ding et al., 2012) and makes the following improvement:

- (i) We provide a proof for the equivalence of MRP and SSP problems.

- (ii) We provide a proof of convergence for our actor-critic method on the SSP problem.
- (iii) Compared with (Ding et al., 2012), we propose a more accurate *safety score*, which helps simplify the RSP structure.
- (iv) We include a case study with more complex temporal logic specifications and present results showing that the specifications are satisfied in a sample path generated by the actor-critic method.
- (v) We analyze the time and memory usages of the actor-critic method in (Ding et al., 2011) for problems with different sizes. The results show that the actor-critic method uses much less time and memory for large problems.

The remainder of the paper is organized as follows. In Sec. 2, we formulate the motion control problem with temporal logic specifications. Sec. 3 describes our method to solve this problem. Sec. 4 illustrates our experiment setups and presents the results accordingly. Sec. 5 concludes the paper.

Notation We use bold letters to denote sequences and vectors. Vectors are assumed to be column vectors. Transpose of a vector \mathbf{y} is denoted by \mathbf{y}^T . $|S|$ denotes the cardinality of a set S .

2 Problem Formulation

Markov Decision Processes (MDPs) are mathematical frameworks for analyzing partially controllable systems and are commonly used to model robot movements.

In this paper, we consider a robot moving in an environment partitioned into regions such as the Robotic Indoor Environment (RIDE) (see Fig. 1) Please add citation for RIDE. Each region in the environment is associated with a set of properties. Properties can be **Un** for unsafe regions, or **Up** for a region where the robot can upload data. We assume that the robot can detect its current region. Moreover, the robot is programmed with a set of motion primitives allowing it to move from a region to an adjacent region. To capture noise in actuation and sensing, we make the natural assumption that, at a given region, a motion primitive designed to take the robot to a specific adjacent region may take the robot to a different adjacent region.

Such a robot model naturally leads to a labeled Markov Decision Process (MDP), which is defined below.

Definition 2.1 (Labeled Markov Decision Process). *A labeled Markov Decision Process (MDP) is a tuple $\mathcal{M} = (Q, q_0, U, A, P, \Pi, h)$, where*

- (i) $Q = \{1, \dots, n\}$ is a finite set of states;
- (ii) $q_0 \in Q$ is the initial state;
- (iii) U is a finite set of actions;
- (iv) $A : Q \rightarrow 2^U$ maps state $q \in Q$ to actions enabled at q ;
- (v) $P : Q \times U \times Q \rightarrow [0, 1]$ is the transition probability function such that for all $q \in Q$, $\sum_{q' \in Q} P(q, u, q') = 1$ if $u \in A(q)$, and $P(q, u, q') = 0$ for all $q' \in Q$ if $u \notin A(q)$;
- (vi) Π is a set of properties;

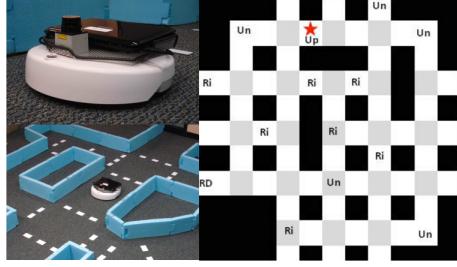


Figure 1: Robotic InDoor Environment (RIDE) platform. **(Left:)** An iCreate mobile platform moving autonomously through the corridors and intersections of an indoor-like environment. The robot is equipped with a RFID reader that can correctly identify cards placed on the floor and with a laser range finder that is used to implement motion primitives such as GoLeft and GoForward in an intersection, etc. **(Right:)** An example schematic of the environment. The black blocks represent walls, and the grey and white regions are intersections and corridors, respectively. The labels inside a region represent properties associated with regions, such as **Un** (unsafe regions) and **Ri** (risky regions).

(vii) $h : Q \rightarrow 2^\Pi$ is the property map.

Note that although this paper focuses on robot control in RIDE platform, the label MDP definition is general and can be used to in applications such as [add some examples and citations here](#)

In RIDE platform, each state of the MDP \mathcal{M} modeling the robot in the environment corresponds to an ordered set of regions in the environment, while the actions label the motion primitives that can be applied at a region. For example, a state of \mathcal{M} may be labelled as $I_1\text{-}C_1$, which means that the robot is currently at region C_1 , coming from region I_1 . Each ordered set of regions corresponds to a recent history of the robot trajectory, and is needed to ensure the Markov property (more details on such MDP abstractions of the robot in the environment can be found in *e.g.*, (Lahijanian et al., 2010)). The transition probability function P can be obtained through extensive simulations of the robot in the environment. We assume that there exists an accurate simulator that is capable of generating (computing) the transition probability $P(q, u, \cdot)$ for each state-action pair $q \in Q$ and $u \in A(q)$. In our previous work (Lahijanian et al., 2010), we developed such a simulator for the robot shown in Fig. 1. More details on the construction of the MDP model for a robot in the RIDE platform are included in Sec. 4.

If the exact transition probabilities are not known, \mathcal{M} can be seen as a labeled Non-Deterministic Transition System (NTS) $\mathcal{M}^N = (Q, q_0, U, A, P^N, \Pi, h)$, where P in \mathcal{M} is replaced by $P^N : Q \times U \times Q \rightarrow \{0, 1\}$, and $P^N(q, u, q') = 1$ indicates a possible transition from q to q' applying an enabled action $u \in A(q)$; if $P^N(q, u, q') = 0$, then the transition from q to q' is not possible under u .

Reviewer 2's comment: - p.p. 4 line 47: In the definition of a path, the second region in the label associated with each state matches the first region of the following state. Only then will a path correspond(s) to a sequence of regions in the environment.

A path on \mathcal{M} is a sequence of states $\mathbf{q} = q_0 q_1 \dots$ such that for all $k \geq 0$, there exists $u_k \in A(q_k)$ such that $P(q_k, u_k, q_{k+1}) > 0$. Along a path $\mathbf{q} = q_0 q_1 \dots, q_k$ is said to be the state at time k . The trajectory of the robot in the environment is represented by a path \mathbf{q} on \mathcal{M} (which corresponds to a sequence of regions in the environment). A path $\mathbf{q} = q_1 q_2 \dots$ generates a sequence of properties

$\mathbf{h}(\mathbf{q}) := o_1 o_2 \dots$, where $o_k = h(q_k)$ for all $k \geq 0$. We call $\mathbf{o} = \mathbf{h}(\mathbf{q})$ the word generated by \mathbf{q} .

Definition 2.2 (Policy). A control policy for an MDP \mathcal{M} is an infinite sequence $M = \mu_0 \mu_1 \dots$, where $\mu_k : Q \times U \rightarrow [0, 1]$ is such that $\sum_{u \in A(q)} \mu_k(q, u) = 1$, for all $k \geq 0$ and $\forall q \in Q$.

Namely, at time k , $\mu_k(q, \cdot)$ is a discrete probability distribution over $A(q)$. If $\mu_k = \mu$ for all $k \geq 0$, then $M = \mu \mu \dots$ is called a stationary policy. If for all $k \geq 0$ and $\forall q \in Q$, $\mu_k(q, u) = 1$ for some u , then M is deterministic; otherwise, M is randomized. Given a policy M , we can then generate a sequence of states on \mathcal{M} , by applying u_k with probability $\mu_k(q_k, u_k)$ at state q_k for all time k . Such a sequence is called a path of \mathcal{M} under policy M .

We require the trajectory of the robot in the environment to satisfy a rich task specification given as a Linear Temporal Logic (LTL) (see, e.g., (Baier et al., 2008; Clarke et al., 1999)) formula over a set of properties Π . An LTL formula over Π is evaluated over an (infinite) sequence $\mathbf{o} = o_0 o_1 \dots$ (e.g., a word generated by a path on \mathcal{M}), where $o_k \subseteq \Pi$ for all $k \geq 0$. We denote $\mathbf{o} \models \phi$ if word \mathbf{o} satisfies the LTL formula ϕ , and we say \mathbf{q} satisfies ϕ if $\mathbf{h}(\mathbf{q}) \models \phi$. Roughly, ϕ can be constructed from a set of properties Π , Boolean operators \neg (negation), \vee (disjunction), \wedge (conjunction), \rightarrow (implication), and temporal operators \mathbf{X} (next), \mathbf{U} (until), \mathbf{F} (eventually), \mathbf{G} (always). A variety of robotic tasks can be easily translated to LTL formulas. For example, the following complex task command in natural language: “Gather data at locations **Da** infinitely often. Only reach a risky region **Ri** if valuable data **VD** can be gathered, and always avoid unsafe regions (**Un**)” can be translated to the LTL formula:

$$\phi := \mathbf{G} \mathbf{F} \mathbf{D}\mathbf{a} \wedge \mathbf{G} (\mathbf{R}\mathbf{i} \rightarrow \mathbf{V}\mathbf{D}) \wedge \mathbf{G} \neg \mathbf{U}\mathbf{n}.$$

In (Ding et al., 2011) (see also (Rutten et al., 2004)), we consider the following problem.

Problem 2.3. Given a labeled MDP $\mathcal{M} = (Q, q_0, U, A, P, \Pi, h)$ and an LTL formula ϕ , find a control policy that maximizes the probability of its path satisfying ϕ .

The probability that paths generated under a policy μ satisfy an LTL formula ϕ is well defined with a suitable measure over the set of all paths generated by μ (Baier et al., 2008).

In (Ding et al., 2011), we proposed a computational framework to solve Prob. 2.3, by adapting methods from the area of probabilistic model checking (De Alfaro, 1997; Baier et al., 2008; Vardi, 1999). However, this framework relies upon the fact that the transition probabilities are known for all state-action pairs. These transition probabilities are typically not available for robotic applications and computationally expensive to compute. Moreover, even if the transition probabilities are obtained for each state-action pair, this method still requires solving a linear program on the product of the MDP and the automata representing the formula, which can be very large (thousands or even millions of states).

In many robotic applications, the NTS model $\mathcal{M}^N = (Q, q_0, U, A, P^N, \Pi, h)$ can be quickly constructed for the robot in the environment and a simulator is available to generate transition probabilities on the fly. In this paper, we focus on the following problem.

Problem 2.4. Given a labeled NTS $\mathcal{M}^N = (Q, q_0, U, A, P^N, \Pi, h)$, an LTL formula ϕ , and an accurate simulator to compute transition probabilities $P(q, u, \cdot)$ given a state-action pair (q, u) , find a control policy that maximizes the probability of its path satisfying ϕ .

Transition probabilities for all state-action pairs are necessary for exact optimal solution of Prob. 2.4. In this paper, we propose an approximate method that only needs transition probabilities for a portion of state-action pairs. Our approach to Prob. 2.4 can be summarized as follows. First, we formulate the problem as a Maximal Reachability Probability (MRP) problem using \mathcal{M}^N and ϕ (Sec. 3.1), and convert the MRP problem into a Stochastic Shortest Path (SSP) problem. We then use an actor-critic framework to find a randomized policy giving an approximate solution to the SSP problem (Sec. 3.3). The randomized policy is constructed to be a function of a small set of

parameters and we find a policy that is locally optimal with respect to these parameters. The construction of a class of policies suitable for SSP problems is explained in Sec. 3.4. The algorithmic framework presented in this paper is summarized in Sec. 3.5.

3 Control Synthesis

3.1 Formulation of the MRP Problem

- Reviewer 2: p.p. 6 line 11: what are the modifications if needed? This sentence needs rewording. The formulation of the MRP problem is based on (Ding et al., 2011; De Alfaro, 1997; Baier et al., 2008; Vardi, 1999) with modification if needed when using the NTS \mathcal{M}_N instead of \mathcal{M} . We start by converting the LTL formula ϕ over Π to a so-called deterministic *Rabin automaton*, which is defined as follows.

Definition 3.1 (Deterministic Rabin Automaton). *A deterministic Rabin automaton (DRA) is a tuple $\mathcal{R} = (S, s_0, \Sigma, \delta, F)$, where*

- (i) S is a finite set of states;
- (ii) $s_0 \in S$ is the initial state;
- (iii) Σ is a set of inputs (alphabet);
- (iv) $\delta : S \times \Sigma \rightarrow S$ is the transition function;
- (v) $F = \{(L(1), K(1)), \dots, (L(M), K(M))\}$ is a set of pairs of sets of states such that $L(i), K(i) \subseteq S$ for all $i = 1, \dots, M$.

A run of a Rabin automaton \mathcal{R} , denoted by $\mathbf{r} = s_0 s_1 \dots$, is an infinite sequence of states in \mathcal{R} such that for each $k \geq 0$, $s_{k+1} \in \delta(s_k, \alpha)$ for some $\alpha \in \Sigma$. A run \mathbf{r} is *accepting* if there exists a pair $(L, K) \in F$ such that \mathbf{r} intersects with L finitely many times and K infinitely many times. For any LTL formula ϕ over Π , one can construct a DRA (which we denote by \mathcal{R}_ϕ) with input alphabet $\Sigma = 2^\Pi$ accepting all and only words over Π that satisfy ϕ (see (Gradel et al., 2002)).

We then obtain an MDP as the product of a labeled MDP \mathcal{M} and a DRA \mathcal{R}_ϕ , which captures all paths of \mathcal{M} satisfying ϕ . Note that this product MDP can only be constructed from an MDP and a deterministic automaton, this is why we require a DRA instead of, e.g., a (generally non-deterministic) Büchi automaton (see (Baier et al., 2008)).

Reviewer 2: - p.p. 6 line 53: F_P should be part of the definition of the product MDP.

Definition 3.2 (Product MDP). *The product MDP $\mathcal{M} \times \mathcal{R}_\phi$ between a labeled MDP $\mathcal{M} = (Q, q_0, U, A, P, \Pi, h)$ and a DRA $\mathcal{R}_\phi = (S, s_0, 2^\Pi, \delta, F)$ is an MDP $\mathcal{P} = (S_\mathcal{P}, s_{\mathcal{P}0}, U_\mathcal{P}, A_\mathcal{P}, P_\mathcal{P}, h_\mathcal{P})$, where*

- (i) $S_\mathcal{P} = Q \times S$ is a set of states;
- (ii) $s_{\mathcal{P}0} = (q_0, s_0)$ is the initial state;
- (iii) $U_\mathcal{P} = U$ is a set of actions inherited from \mathcal{M} ;
- (iv) $A_\mathcal{P}$ is also inherited from \mathcal{M} and $A_\mathcal{P}((q, s)) := A(q)$;
- (v) $P_\mathcal{P}$ gives the transition probabilities:

$$P_\mathcal{P}((q, s), u, (q', s')) = \begin{cases} P(q, u, q'), & \text{if } q' = \delta(s, h(q)), \\ 0, & \text{otherwise.} \end{cases}$$

Note that $h_{\mathcal{P}}$ is not used in the product MDP. Moreover, \mathcal{P} is associated with pairs of accepting states (similar to a DRA) $F_{\mathcal{P}} := \{(L_{\mathcal{P}}(1), K_{\mathcal{P}}(1)), \dots, (L_{\mathcal{P}}(M), K_{\mathcal{P}}(M))\}$ where $L_{\mathcal{P}}(i) = Q \times L(i)$, $K_{\mathcal{P}}(i) = Q \times K(i)$, for $i = 1, \dots, M$.

The product MDP is constructed such that, given a path $(s_0, q_0)(s_1, q_1) \dots$, the corresponding path $s_0s_1 \dots$ on \mathcal{M} satisfies ϕ if and only if there exists a pair $(L_{\mathcal{P}}, K_{\mathcal{P}}) \in F_{\mathcal{P}}$ satisfying the Rabin acceptance condition, *i.e.*, the set $K_{\mathcal{P}}$ is visited infinitely often and the set $L_{\mathcal{P}}$ is visited finitely often.

We can make a very similar product between a labeled NTS $\mathcal{M}^{\mathcal{N}} = (Q, q_0, U, A, P^{\mathcal{N}}, \Pi, h)$ and \mathcal{R}_{ϕ} . This product is also an NTS, which we denote by $\mathcal{P}^{\mathcal{N}} = (S_{\mathcal{P}}, s_{\mathcal{P}0}, U_{\mathcal{P}}, A_{\mathcal{P}}, P_{\mathcal{P}}^{\mathcal{N}}, \Pi, h_{\mathcal{P}}) := \mathcal{M}^{\mathcal{N}} \times \mathcal{R}_{\phi}$, associated with accepting sets $F_{\mathcal{P}}$. The definition (and the accepting condition) of $\mathcal{P}^{\mathcal{N}}$ is exactly the same as for the product MDP. The only difference between $\mathcal{P}^{\mathcal{N}}$ and \mathcal{P} is in $P_{\mathcal{P}}^{\mathcal{N}}$, which is either 0 or 1 for every state-action-state tuple.

From the product \mathcal{P} or equivalently $\mathcal{P}^{\mathcal{N}}$, we can proceed to construct the MRP problem. To do so, it is necessary to produce the so-called *accepting maximum end components* (AMECs). An end component is a subset of an MDP (consisting of a subset of states and a subset of enabled actions at each state) such that for each pair of states (i, j) in \mathcal{P} , there is a sequence of actions such that i can be reached from j with positive probability, and states outside the component cannot be reached (See Def. 10.117 of (Baier et al., 2008)). The definition of AMECs is as follows.

Definition 3.3 (Accepting Maximal End Components). *Given $(L_{\mathcal{P}}, K_{\mathcal{P}}) \in F_{\mathcal{P}}$, an AMEC of \mathcal{P} is the largest end component containing at least one state in $K_{\mathcal{P}}$ and no state in $L_{\mathcal{P}}$, for a pair $(K_{\mathcal{P}}, L_{\mathcal{P}}) \in F_{\mathcal{P}}$. Dennis, can you confirm (and improve) the definition of AMEC?*

Note that an AMEC always contains at least one state in $K_{\mathcal{P}}$ and no state in $L_{\mathcal{P}}$. In addition, AMECs are “absorbing” in the sense that the state does not leave once it reaches any state in an AMEC. *add citation here* A procedure to obtain all AMECs of an MDP is outlined in (Baier et al., 2008). This procedure is intended to be used for the product MDP \mathcal{P} , but it can be used without modification to find all AMECs associated with \mathcal{P} when $\mathcal{P}^{\mathcal{N}}$ is used instead of \mathcal{P} . This is because the information needed to construct the AMECs is the set of all possible state transitions at each state, and this information is already contained in $\mathcal{P}^{\mathcal{N}}$. Note that the computation of AMECs, whose time complexity is quadratic to the size of \mathcal{P} or $\mathcal{P}^{\mathcal{N}}$, cannot be avoided in any method as long as the LTL formula is not co-safe (Baier et al., 2008). As a result, we exclude the time of calculating AMECs when comparing our actor-critic algorithm with alternative methods. For a co-safe LTL formula, this computation cost can be avoided by using Deterministic Finite Automaton (DFA) instead of DRA (Baier et al., 2008).

If we denote by $S_{\mathcal{P}}^*$ as the union of all states in all AMECs associated with \mathcal{P} , it has been shown in probabilistic model checking (see *e.g.*, (Baier et al., 2008)) that maximizing the probability of satisfying the LTL formula is equivalent to a Maximal Reachability Probability (MRP) problem whose definition is as follows.

Problem 3.4. *Given a product MDP $\mathcal{P} = (S_{\mathcal{P}}, s_{\mathcal{P}0}, U_{\mathcal{P}}, A_{\mathcal{P}}, P_{\mathcal{P}}, F_{\mathcal{P}})$ and a set of states $S_{\mathcal{P}}^* \subseteq S_{\mathcal{P}}$, find the optimal policy μ that maximizes the probability of reaching the set $S_{\mathcal{P}}^*$.*

If transition probabilities are available for each state-action pair, then Problem 3.4 can be solved by a linear program (see (Puterman, 1994; Baier et al., 2008)). The resultant optimal policy is deterministic and is a huge table containing optimal controls for each state s in \mathcal{P} . In this paper, we approximate the optimal policy μ with a parameterized policy μ_{θ} , which improves computation efficiency by taking advantage of prior knowledge of the policy structure.

By definition of AMECs, if the state of the product MDP reaches $S_{\mathcal{P}}^*$, it can not leave it. We call such a state set *absorbing state set*. Intuitively, the only case when the state does not reach $S_{\mathcal{P}}^*$ is because it is “trapped” in other set of states. We present the following definition.

Definition 3.5 (Trap State). A trap state is a state on product MDP \mathcal{P} that can not reach $S_{\mathcal{P}}^*$ under any policy.

Denote by $\bar{S}_{\mathcal{P}}^*$ the set of all trap states, we have the following theorem.

Theorem 3.6. $\bar{S}_{\mathcal{P}}^*$ is an absorbing state set of product MDP \mathcal{P} . Furthermore, in product MDP \mathcal{P} , there is no absorbing state set in states $S_{\mathcal{P}} / (\bar{S}_{\mathcal{P}}^* \cup S_{\mathcal{P}}^*)$.

Proof. If we assume that there is a state $s_{\mathcal{P}}^A \in \bar{S}_{\mathcal{P}}^*$ and a state $s_{\mathcal{P}}^B \notin \bar{S}_{\mathcal{P}}^*$ such that the $s_{\mathcal{P}}^B$ is accessible from $s_{\mathcal{P}}^A$ under certain policy. By definition of $\bar{S}_{\mathcal{P}}^*$, $S_{\mathcal{P}}^*$ should not be accessible from $s_{\mathcal{P}}^A$. However, $s_{\mathcal{P}}^B$ is accessible from $s_{\mathcal{P}}^B$ thus is accessible from $s_{\mathcal{P}}^A$. This contradiction shows that no outside state is accessible from any state in $\bar{S}_{\mathcal{P}}^*$, thus $\bar{S}_{\mathcal{P}}^*$ is an absorbing state set.

In addition, if there exists any other absorbing state set that does not intersect with either $S_{\mathcal{P}}^*$ or $\bar{S}_{\mathcal{P}}^*$. By definition, states in such an absorbing state state should be trap states, thus this set should be a subset of $\bar{S}_{\mathcal{P}}^*$, which is contradictory. ■

Remark 3.7. Theorem 3.6 suggests that Problem 3.4 is equivalent to the problem of minimizing the probabilities of reaching the set $\bar{S}_{\mathcal{P}}^*$.

Remark 3.8. It is only necessary to find the optimal policy for states not in the set $S_{\mathcal{P}}^*$. This is because by construction, there exists a policy inside any AMEC that almost surely satisfies the LTL formula ϕ by reaching a state in $K_{\mathcal{P}}$ infinitely often. This policy can be obtained by simply choosing an action (among the subset of actions retained by the AMEC) at each state randomly, i.e., a trivial randomized stationary policy exists that almost surely satisfies ϕ .

3.2 Conversion from MRP to Stochastic Shortest Path (SSP) Problem

Although linear program can be used to solve Problem 3.4, it has high time and memory complexity. In particular, it requires transition probabilities for all states of the product MDP. In many situations, we only have a simulator to generate the transition probabilities on the fly, and due to the large state space, we want to avoid simulating for all states.

The so-called approximate dynamic programming (also known as neuro-dyanmic programming or reinforcement learning) can help solve these issues (Bertsekas and Tsitsiklis, 1996; Bertsekas et al., 1995; Sutton and Barto, 1998). In order to apply techniques in approximate dynamic programming, we need to convert the MRP problem into a Stochastic Shortest Path (SSP) problem. We define the following new MDP based on the product MDP.

Definition 3.9 (SSP MDP). Given the product MDP $\mathcal{P} = (S_{\mathcal{P}}, s_{\mathcal{P}0}, U_{\mathcal{P}}, A_{\mathcal{P}}, P_{\mathcal{P}}, F_{\mathcal{P}})$ and a set of states $S_{\mathcal{P}}^* \subseteq S_{\mathcal{P}}$, define a new MDP $\tilde{\mathcal{P}} = (\tilde{S}_{\mathcal{P}}, \tilde{s}_{\mathcal{P}0}, \tilde{U}_{\mathcal{P}}, \tilde{A}_{\mathcal{P}}, \tilde{P}_{\mathcal{P}})$, where

- (i) $\tilde{S}_{\mathcal{P}} = (S_{\mathcal{P}} \setminus S_{\mathcal{P}}^*) \cup \{s_{\mathcal{P}}^*\}$, where $s_{\mathcal{P}}^*$ is a “dummy” terminal state;
- (ii) $\tilde{s}_{\mathcal{P}0} = s_{\mathcal{P}0}$ (without loss of generality, we exclude the trivial case where $s_{\mathcal{P}0} \in S_{\mathcal{P}}^*$);
- (iii) $\tilde{U}_{\mathcal{P}} = U_{\mathcal{P}}$;
- (iv) $\tilde{A}_{\mathcal{P}}(s_{\mathcal{P}}) = A_{\mathcal{P}}(s_{\mathcal{P}})$ for all $s_{\mathcal{P}} \in S_{\mathcal{P}}$, and for the dummy state we set $\tilde{A}_{\mathcal{P}}(s_{\mathcal{P}}^*) = \tilde{U}_{\mathcal{P}}$;
- (v) The transition probability is redefined as follows. We then define:

$$\begin{aligned} & \tilde{P}_{\mathcal{P}}(s_{\mathcal{P}}, u, s'_{\mathcal{P}}) \\ &= \begin{cases} \sum_{s''_{\mathcal{P}} \in S_{\mathcal{P}}^*} P_{\mathcal{P}}(s_{\mathcal{P}}, u, s''_{\mathcal{P}}), & \text{if } s'_{\mathcal{P}} = s_{\mathcal{P}}^*, \\ P_{\mathcal{P}}(s_{\mathcal{P}}, u, s'_{\mathcal{P}}), & \text{if } s'_{\mathcal{P}} \in S_{\mathcal{P}} \setminus S_{\mathcal{P}}^*, \end{cases} \end{aligned}$$

for all $s_{\mathcal{P}} \in S_{\mathcal{P}} \setminus (S_{\mathcal{P}}^* \cup \bar{S}_{\mathcal{P}}^*)$ and $u \in \tilde{U}_{\mathcal{P}}$. Moreover, for all $s_{\mathcal{P}} \in \bar{S}_{\mathcal{P}}^*$ and $u \in \tilde{U}_{\mathcal{P}}$, we set $\tilde{P}_{\mathcal{P}}(s_{\mathcal{P}}^*, u, s_{\mathcal{P}}^*) = 1$ and $\tilde{P}_{\mathcal{P}}(s_{\mathcal{P}}, u, s_{\mathcal{P}0}) = 1$;

For all $s_{\mathcal{P}} \in \tilde{S}_{\mathcal{P}}$ and $u \in \tilde{U}_{\mathcal{P}}$, we define a one-step cost function $\tilde{g}_{\mathcal{P}}(s_{\mathcal{P}}, u) = 1$ if $s_{\mathcal{P}} \in \bar{S}_{\mathcal{P}}^*$, and $\tilde{g}_{\mathcal{P}}(s_{\mathcal{P}}, u) = 0$ otherwise. Then the SSP problem is defined as follows:

Problem 3.10. Given an SSP MDP $\tilde{\mathcal{P}} = (\tilde{S}_{\mathcal{P}}, \tilde{s}_{\mathcal{P}0}, \tilde{U}_{\mathcal{P}}, \tilde{A}_{\mathcal{P}}, \tilde{P}_{\mathcal{P}})$ and a one-step cost function $\tilde{g}_{\mathcal{P}}$, find a policy μ to minimize the expected total cost

$$\bar{\alpha}_{\mu} = E \left\{ \sum_{k=0}^{T^*} \tilde{g}_{\mathcal{P}}(x_k, u_k) \right\}, \quad (1)$$

where T^* is the first time when $s_{\mathcal{P}}^*$ is reached, x_k and u_k are the state and the action at time k , respectively.

Remark 3.11. In product MDP \mathcal{P} , both $S_{\mathcal{P}}^*$ and $\bar{S}_{\mathcal{P}}^*$ are absorbing state sets; When the state reaches $S_{\mathcal{P}}^*$ in \mathcal{P} , it cannot leave. In contrast, in SSP MDP $\tilde{\mathcal{P}}$, $s_{\mathcal{P}}^*$ is the only absorbing state; Whenever the state reaches $\bar{S}_{\mathcal{P}}^*$, it returns to the initial state $\tilde{s}_{\mathcal{P}0}$ (as if the process restarts). The expected total cost $\bar{\alpha}_{\mu}$ in (1) is the expected total number of falls into $\bar{S}_{\mathcal{P}}^*$ before reaching $s_{\mathcal{P}}^*$ in SSP MDP $\tilde{\mathcal{P}}$.

Let $R_{\mu}^{\mathcal{P}}$ be the reachability probability in the Problem 3.4 for the policy. The following lemma presents the relationship between $R_{\mu}^{\mathcal{P}}$ and the expected total cost $\bar{\alpha}_{\mu}$ defined in (1).

Lemma 3.12. For any RSP μ , we have $R_{\mu}^{\mathcal{P}} = 1 / (\bar{\alpha}_{\mu} + 1)$.

Proof. According to the definition of the cost function, we know that $\bar{\alpha}_{\mu}$ is the expected number of times when states in $\bar{S}_{\mathcal{P}}^*$ are visited before the termination state $s_{\mathcal{P}}^*$ is reached in SSP MDP $\tilde{\mathcal{P}}$. From the construction of $\tilde{\mathcal{P}}$, reaching $s_{\mathcal{P}}^*$ in $\tilde{\mathcal{P}}$ is equivalent to reaching one of the goal states $S_{\mathcal{P}}^*$ in \mathcal{P} . On the other hand, in the Markov chain generated by applying policy μ to \mathcal{P} , the states $S_{\mathcal{P}}^*$ and $\bar{S}_{\mathcal{P}}^*$ are the only absorbing state sets, and all other states are transient. Thus, the probability of visiting a state in $\bar{S}_{\mathcal{P}}^*$ from $s_{\mathcal{P}0}$ on \mathcal{P} is $1 - R_{\mu}^{\mathcal{P}}$, which is the same as the probability of visiting $\bar{S}_{\mathcal{P}}^*$ for each run of $\tilde{\mathcal{P}}$, due to the construction of transition probabilities in Def. 3.9. We can now consider a geometric distribution where the probability of success is $R_{\mu}^{\mathcal{P}}$. Because $\bar{\alpha}_{\mu}$ is the expected number of times when a state in $\bar{S}_{\mathcal{P}}^*$ is visited before $s_{\mathcal{P}}^*$ is reached, this is the same as the expected number of failures of Bernoulli trials (with probability of success being $R_{\mu}^{\mathcal{P}}$) before a success. This implies $\bar{\alpha}_{\mu} = \frac{1 - R_{\mu}^{\mathcal{P}}}{R_{\mu}^{\mathcal{P}}}$. ■

Lemma 3.12 states that the policy minimizing (1) for Problem 3.10 with MDP $\tilde{\mathcal{P}}$ and the termination state $s_{\mathcal{P}}^*$ is a policy maximizing the probability of reaching the set $S_{\mathcal{P}}^*$ on \mathcal{P} , i.e., a solution to Problem 3.4.

Problem 3.10 can also be constructed from the NTS \mathcal{P}^N . In this case we obtain an NTS $\tilde{\mathcal{P}}^N = (\tilde{S}_{\mathcal{P}}, \tilde{s}_{\mathcal{P}0}, \tilde{U}_{\mathcal{P}}, \tilde{A}_{\mathcal{P}}, \tilde{P}_{\mathcal{P}}^N)$, using the exact same construction as Def. 3.9, except for the definition of $\tilde{P}_{\mathcal{P}}^N$. The transition function $\tilde{P}_{\mathcal{P}}^N(s_{\mathcal{P}}, u, s'_{\mathcal{P}})$ is instead defined as:

$$\begin{aligned} & \tilde{P}_{\mathcal{P}}^N(s_{\mathcal{P}}, u, s'_{\mathcal{P}}) \\ &= \begin{cases} \max_{s''_{\mathcal{P}} \in S_{\mathcal{P}}^*} P_{\mathcal{P}}^N(s_{\mathcal{P}}, u, s''_{\mathcal{P}}), & \text{if } s'_{\mathcal{P}} = s_{\mathcal{P}}^* \\ P_{\mathcal{P}}^N(s_{\mathcal{P}}, u, s'_{\mathcal{P}}), & \text{if } s'_{\mathcal{P}} \in S_{\mathcal{P}} \setminus S_{\mathcal{P}}^* \end{cases} \end{aligned}$$

for all $s_{\mathcal{P}} \in S_{\mathcal{P}} \setminus (S_{\mathcal{P}}^* \cup \bar{S}_{\mathcal{P}}^*)$ and $u \in \tilde{U}_{\mathcal{P}}$. Moreover, for all $s_{\mathcal{P}} \in \bar{S}_{\mathcal{P}}^*$ and $u \in \tilde{U}_{\mathcal{P}}$, we set $\tilde{P}_{\mathcal{P}}^N(s_{\mathcal{P}}^*, u, s_{\mathcal{P}}^*) = 1$ and $\tilde{P}_{\mathcal{P}}^N(s_{\mathcal{P}}, u, s_{\mathcal{P}0}) = 1$.

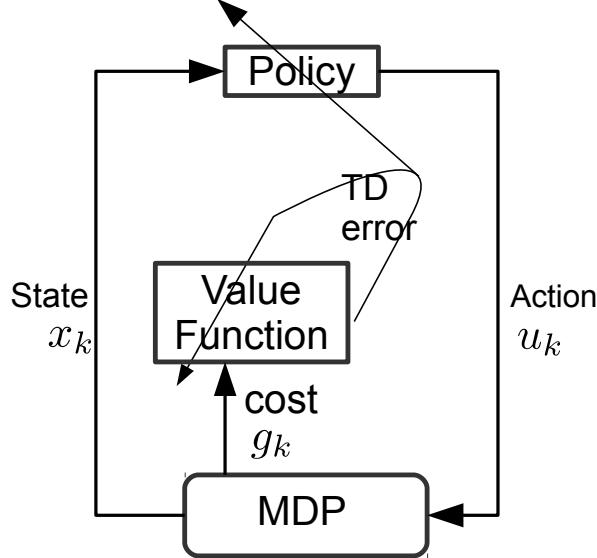


Figure 2: Illustration of actor-critic methods.

3.3 LSTD Actor-Critic Method

In this section, we propose an action-critic method that obtains a Randomized Stationary Policy (RSP) (see Def. 2.2) $M = \mu_\theta \mu_\theta \dots$, where $\mu_\theta(x, u)$ is a function of the state-action pair (x, u) and $\theta \in \mathbb{R}^n$, which is a vector of parameters. For convenience, we denote an RSP $\mu_\theta \mu_\theta \dots$ simply by μ_θ . In addition, we denote the expected total cost defined in (1) by $\bar{\alpha}(\theta)$. In this section we assume the RSP $\mu_\theta(q, u)$ to be given, and we will describe in Sec. 3.4 how to design a suitable RSP.

Given an RSP μ_θ , we apply an iterative procedure, *i.e.*, an actor-critic method, to obtain a policy that locally minimizes the cost function (1) by simulating sample paths on $\tilde{\mathcal{P}}$. Each sample path on $\tilde{\mathcal{P}}$ starts at $\tilde{s}_{\mathcal{P}0}$ and ends when the termination state $s_{\mathcal{P}}^*$ is reached. Since the probabilities are needed only along the sample path, we do not require the MDP $\tilde{\mathcal{P}}$, but only $\tilde{\mathcal{P}}^N$.

As suggested by the name, actor-critic algorithms have two learning units, an actor and a critic, interacting with each other and with the MDP during the iterations of the algorithms. At each iteration, the critic observes the state and the one-step cost from the MDP and uses the observed information to update a value function. The value function is then used to update the RSP; the actor generates the action based on the RSP and applies the action to the MDP (cf. Fig. 2). The algorithm stops when the gradient of $\bar{\alpha}(\theta)$ is small enough (*i.e.*, θ is locally optimal).

Consider the SSP MDP $\tilde{\mathcal{P}}$. Let k denote time, $x_k \in \tilde{\mathcal{S}}_{\mathcal{P}}$ and $u_k \in \tilde{\mathcal{A}}_{\mathcal{P}}(x_k)$ be the state and the action taken at time k . Under a fixed policy μ_θ , $\{x_k\}$ and $\{x_k, u_k\}$ are Markov chains with stationary distributions. We denote these stationary distributions as $\pi_\theta(x)$ and $\eta_\theta(x, u)$, respectively.

Let P_θ be an operator to take expectation after one transition, namely, for a function $f(x, u)$,

$$(P_\theta f)(x, u) = \sum_{j \in \tilde{\mathcal{S}}_{\mathcal{P}}} \sum_{\nu \in \tilde{\mathcal{A}}_{\mathcal{P}}(j)} \tilde{P}_{\mathcal{P}}(x, u, j) \mu_\theta(j, \nu) f(j, \nu).$$

Define the function Q_{θ} to be the function that satisfies the following Poisson equation:

$$Q_{\theta}(x, u) = g_{\mathcal{P}}(x, u) + (P_{\theta}Q_{\theta})(x, u). \quad (2)$$

$Q_{\theta}(x, u)$ can be interpreted as the expected total future cost after applying action u at state x and is named as the state-action value function (Sutton and Barto, 1998; Moazzez Estanjini et al., 2011a). Let

$$\psi_{\theta}(x, u) = \nabla_{\theta} \ln(\mu_{\theta}(x, u)), \quad (3)$$

where $\psi(x, i) = 0$ when x, u are such that $\mu_{\theta}(u|x) \equiv 0$ for all θ . We assume that $\psi_{\theta}(x, u)$ is bounded and continuously differentiable. For all state $x \in \tilde{S}_{\mathcal{P}}$ and action $u \in \tilde{A}_{\mathcal{P}}(x)$, $\psi_{\theta}(x, u)$ is n -dimensional vector, where n is the dimensionality of θ . We write $\psi_{\theta}(x, u) = (\psi_{\theta}^1(x, u), \dots, \psi_{\theta}^n(x, u))$. Let $\langle \cdot, \cdot \rangle$ be an inner product operator defined as following:

$$\nabla \langle f_1, f_2 \rangle = \sum_{x \in \tilde{S}_{\mathcal{P}}} \sum_{u \in \tilde{A}_{\mathcal{P}}(x)} \eta_{\theta}(x, u) f_1(x, u) f_2(x, u), \quad (4)$$

where $f_1(x, u)$ and $f_2(x, u)$ are two functions. It has been proved that the gradient of the expected total cost $\bar{\alpha}(\theta)$ is equal to (Konda, 2002)

$$\nabla \bar{\alpha}(\theta) = \langle Q_{\theta}, \psi_{\theta} \rangle. \quad (5)$$

Instead of storing Q_{θ} explicitly, which is a huge table, we approximate Q_{θ} with a linear architecture of the following form

$$Q_{\theta}^r(x, u) = \psi_{\theta}'(x, u) \mathbf{r}, \mathbf{r} \in \mathbb{R}^n, \quad (6)$$

Let $\|\cdot\|$ be the normal induced by inner product operator $\langle \cdot, \cdot \rangle$, i.e., $\|f\|^2 = \langle f, f \rangle$. The optimal coefficient \mathbf{r}^* should minimize the norm between Q_{θ} and $Q_{\theta}^{\mathbf{r}^*}$, namely,

$$\mathbf{r}^* = \arg \min_{\mathbf{r}} \|Q_{\theta} - Q_{\theta}^{\mathbf{r}}\|. \quad (7)$$

Temporal difference algorithms can be leveraged to learn the optimal coefficient \mathbf{r}^* .

In this paper, we present an actor-critic algorithm that uses Least Square Temporal Difference learning to estimate the \mathbf{r}^* . We summarize the algorithm, which is referred to as LSTD actor-critic algorithm, in Alg. 1, and we note that it does not depend on the form of RSP μ_{θ} . In each iteration, we first compute the transition probabilities $\tilde{P}_{\mathcal{P}}(x_k, u_k, \cdot)$ using a simulator (Step 3 of Alg. 1). Note that different with linear programming method (Ding et al., 2011), the transition probabilities are generated on the fly. This approach is suitable for problems with large state space and strict memory requirements. Next, we obtain the simulated next state x_{k+1} based on the transition probabilities and obtain an action u_{k+1} based on the current RSP μ_{θ_k} (Step 4, 5 of Alg. 1). Then we update our estimate of \mathbf{r} in the critic step (Step 6 of Alg. 1), and update our estimate of θ in the actor step (Step 7 of Alg. 1).

In the critic step, $\mathbf{z}_k \in \mathbb{R}^n$ represents Sutton's eligibility trace (Sutton and Barto, 1998; Konda, 2002); $\mathbf{b}_k \in \mathbb{R}^n$ maintains a sample estimate for one-step cost; $\mathbf{A}_k \in \mathbb{R}^{n \times n}$ is a sample estimate for the matrix formed by $\mathbf{z}_k(\psi_{\theta_k}(x_{k+1}, u_{k+1}) - \psi_{\theta_k}(x_k, u_k))$. \mathbf{r}_k is a least square estimate of \mathbf{r} .

In the actor step, $\mathbf{r}_k^T \psi_{\theta_k}(x_{k+1}, u_{k+1}) \psi_{\theta_k}(x_{k+1}, u_{k+1})$ is a sample estimate of the gradient $\nabla \bar{\alpha}(\theta)$ (cf. Eq. (5)). The actor step is simply a gradient descent update. The role of $\Gamma(\mathbf{r})$ is mainly to keep the actor updates bounded, and we can for instance use $\Gamma(\mathbf{r}) = \min(\frac{D}{\|\mathbf{r}\|}, 1)$ for some $D > 0$.

In Alg. 1, $\{\gamma_k\}$ controls the critic step-size, while $\{\beta_k\}$ control the actor step-size together. We leave the proof of convergence for Alg. 1 as well as the requirements for the step-sizes to the appendix.

Alg. 1 learns the critic parameters using a Least Squares Temporal Difference (LSTD) method, which has been shown to be superior to other stochastic learning methods in terms of the convergence rate (Konda and Tsitsiklis, 2003; Boyan, 1999). (Estanjini et al., 2012) proposes and establishes the convergence of an LSTD actor-critic method similar to Alg. 1 for problems of minimizing *expected average costs*. In comparison, the goal of the Problem 3.10 in this paper is to minimize an *expected total cost* (cf. Eq. (1)).

Compared with some existing efforts of applying approximate dynamic programming techniques in robot motion control, the actor-critic algorithm used in this paper is very suitable for large-scale problems. For example, the method in (Fu and Topcu, 2014) is based on value iteration. Value iteration is a classical dynamic programming technique, but it is suitable only for MDPs with small state spaces because of its high time and memory complexity (Bertsekas et al., 1995). In addition, value iteration also requires transition probabilities of all states. The method in (Sadigh et al., 2014) is also based on temporal difference learning, however, it stores value function in a huge table. *i.e.*, explicitly stores the expected future cost for each state. The table will take a lot of memory, which makes them unsuitable for large-sized problem. In contrast, in actor-critic algorithm, only the combination coefficient \mathbf{r} needs to be stored since the value function $Q_\theta(x, u)$ is expressed as linear combination of some basis functions. The problems considered in both (Fu and Topcu, 2014) and (Sadigh et al., 2014) are much smaller than ours. Actually, in the context of approximate dynamic programming, the poor scalability of value iteration and table-based temporal difference methods is an important motivation of actor-critic methods (Konda, 2002; Bertsekas and Tsitsiklis, 1996). Dennis, I only compare with the two papers from the perspective of approximate dynamic programmings. You can add some comments from the temporal logic perspective.

3.4 Designing an RSP

In this section we describe an RSP suitable to be used in Alg. 1 for Problem 3.10. We first describe some “features”, *i.e.*, progress and safety scores for each state. Then we define some “desirability degrees” for controls in each state and propose an RSP based on the Boltzmann distribution.

We define the *overlay graph* $G_o = (V_o, E_o)$ for NTS $\tilde{\mathcal{P}}^N$ as an unweighted graph whose vertex set $V_o = \tilde{S}_{\mathcal{P}}$. For any pair of states $i, j \in \tilde{S}_{\mathcal{P}}$, we have $(i, j) \in E_o$ if $\tilde{P}_{\mathcal{P}}^N(i, u, j) = 1$ for some $u \in A_{\mathcal{P}}(i)$, *i.e.*, j is reachable from i in one step in the NTS. Let $d_g(i)$ be the shortest distance from state i to the “dummy” terminal state $s_{\mathcal{P}}^*$ in G_o . The distance $d_g(\cdot)$ can be efficiently calculated using Breath-First Search (BFS) with $O(|V_o| + |E_o|)$ time complexity and $O(|V_o|)$ space complexity (Cormen et al., 2001). Recall that $\tilde{S}_{\mathcal{P}}^*$ is the set of states on $\tilde{\mathcal{P}}^N$ that cannot reach $s_{\mathcal{P}}^*$ under any policy, it is easy to observe that $\tilde{S}_{\mathcal{P}}^* = \{i : d_g(i) = \infty\}$. Note that the LP method described in (Ding et al., 2011) also needs to calculate $\tilde{S}_{\mathcal{P}}^*$ by calculating $d_g(\cdot)$ first. We define the *progress score* of a state $i \in \tilde{S}_{\mathcal{P}}$ as

$$\text{prog}(i) := -d_g(i). \quad (8)$$

Larger *progress score* means the state is closer to $s_{\mathcal{P}}^*$ in G_o , thus, it is more likely to hit $s_{\mathcal{P}}^*$ in the near future.

Let μ_{null} be a “null policy” in which each action is chosen with equal probability. If we fix the policy to be this “null policy”, the process $\{x_k, u_k\}$ becomes a Markov chain. In the following part, we calculate the “ r -step transition probabilities” of this Markov chain, where r is a predefined parameter representing sensing range.

Algorithm 1 LSTD Actor-critic algorithm for Problem 3.10

Input: The NTS $\tilde{P}_{\mathcal{P}}^{\mathcal{N}}(\tilde{S}_{\mathcal{P}}, \tilde{s}_{\mathcal{P}0}, \tilde{U}_{\mathcal{P}}, \tilde{A}_{\mathcal{P}}, \tilde{P}_{\mathcal{P}}^{\mathcal{N}}, g_{\mathcal{P}})$ with the terminal state $s_{\mathcal{P}}^*$, the RSP μ_{θ} , and a computation tool to obtain $\tilde{P}_{\mathcal{P}}(s_{\mathcal{P}}, u, \cdot)$ for a given $(s_{\mathcal{P}}, u)$ state-action pair.

1: **Initialization:** Set all entries in $\mathbf{z}_0, \mathbf{b}_0$ and \mathbf{r}_0 to zeros. Set \mathbf{A}_0 to identity matrix. Let θ_0 take some initial value. Set initial state $x_0 := \tilde{s}_{\mathcal{P}0}$. Obtain action u_0 using the RSP μ_{θ_0} .

2: **repeat**

3: Compute the transition probabilities $\tilde{P}_{\mathcal{P}}(x_k, u_k, \cdot)$.

4: Obtain the simulated subsequent state x_{k+1} using the transition probabilities $\tilde{P}_{\mathcal{P}}(x_k, u_k, \cdot)$. If $x_k = s_{\mathcal{P}}^*$, set $x_{k+1} := x_0$.

5: Obtain action u_{k+1} using the RSP μ_{θ_k}

6: **Critic Update:**

$$\begin{aligned}\mathbf{z}_{k+1} &= \lambda \mathbf{z}_k + \psi_{\theta_k}(x_k, u_k) \\ \mathbf{b}_{k+1} &= \mathbf{b}_k + \gamma_k (\tilde{g}_{\mathcal{P}}(x_k, u_k) \mathbf{z}_k - \mathbf{b}_k) \\ \mathbf{A}_{k+1} &= \mathbf{A}_k + \gamma_k (\mathbf{z}_k (\psi_{\theta_k}^T(x_{k+1}, u_{k+1}) - \psi_{\theta_k}^T(x_k, u_k)) \\ &\quad - \mathbf{A}_k), \\ \mathbf{r}_{k+1} &= -\mathbf{A}_k^{-1} \mathbf{b}_k.\end{aligned}$$

7: **Actor Update:**

$$\theta_{k+1} = \theta_k - \beta_k \Gamma(\mathbf{r}_k) \mathbf{r}_k^T \psi_{\theta_k}(x_{k+1}, u_{k+1}) \psi_{\theta_k}(x_{k+1}, u_{k+1})$$

8: **until** $\|\nabla \bar{\alpha}(\theta_k)\| \leq \epsilon$ for some given ϵ .

Suppose $\tilde{P}_{\mathcal{P}}^{(r)}(i, j)$ is the probability from state i to state j in r steps without reaching any state in $\bar{S}_{\mathcal{P}}^*$ under μ_{null} . $\tilde{P}_{\mathcal{P}}^{(r)}(i, u, j)$ can be calculated recursively. For $\forall i, j, u$ and $m = 2, \dots, r$,

$$\tilde{P}_{\mathcal{P}}^{(m)}(i, j) = \sum_{x \in \tilde{S}_{\mathcal{P}}} \tilde{P}_{\mathcal{P}}^{(m-1)}(i, x) \tilde{P}^{(1)}(x, j),$$

where

$$\tilde{P}^{(1)}(i, j) = \begin{cases} 1, & \text{if } i \in \bar{S}_{\mathcal{P}}^* \text{ and } i = j, \\ 0, & \text{if } i \in \bar{S}_{\mathcal{P}}^* \text{ and } i \neq j, \\ \frac{1}{|\tilde{U}_{\mathcal{P}}|} \sum_{u \in \tilde{U}_{\mathcal{P}}} \tilde{P}_{\mathcal{P}}(i, u, j), & \text{otherwise.} \end{cases}$$

$\tilde{P}_{\mathcal{P}}^{(1)}(i, j)$ is the one-step transition probability from state i to state j under μ_{null} . Define the *safety score* for state i as

$$\text{safe}(i) := \sum_j \tilde{P}_{\mathcal{P}}^{(r)}(i, j) I(j), \quad (9)$$

where $I(j)$ is an indicator function such that $I(i) = 1$ if and only if $i \in \tilde{S}_{\mathcal{P}} \setminus \bar{S}_{\mathcal{P}}^*$ and $I(i) = 0$ otherwise. The *safety score* of a state is the probability of hitting trap states $\bar{S}_{\mathcal{P}}^*$ in the following r steps under μ_{null} . Thus, a higher safety score for the current state implies that it is less likely to reach $\bar{S}_{\mathcal{P}}^*$ in the near future.

Compared with the *safety score* defined in (Ding et al., 2012), (9) is much more accurate. The *safety score* in (Ding et al., 2012) is the proportion of non-trap states in the neighborhood of a state, which is problematic in many cases. For example, for a state that has only one trap state in its neighborhood but has a large transition probability to this trap state, it will have a high *safety score* according to (Ding et al., 2012) despite that it is quite unsafe.

The *desirability degree* of an action u is defined as

$$a(\boldsymbol{\theta}, i, u) = \theta_1 \times \left(\sum_j \tilde{P}_{\mathcal{P}}(i, u, j) \text{prog}(j) - \text{prog}(i) \right) + \theta_2 \times \left(\sum_j \tilde{P}_{\mathcal{P}}(i, u, j) \text{safe}(j) - \text{safe}(i) \right), \quad (10)$$

which is a weighted sum of two contributing terms, and $\boldsymbol{\theta} := [\theta_1, \theta_2]^T$ is the vector of corresponding weights. The first term is based on the influence of this action on improving progress, and the second term is based on its influence on improving safety.

Our RSP is constructed using the Boltzmann distribution. The probability of taking action u at state i is defined as

$$\mu_{\boldsymbol{\theta}}(u|i) = \frac{\exp(a(\boldsymbol{\theta}, i, u)/T)}{\sum_{u \in \tilde{U}_{\mathcal{P}}} \exp(a(\boldsymbol{\theta}, i, u)/T)}, \quad (11)$$

where T is the temperature of the Boltzmann distribution.

As a comparison, to alleviate the influence of the poor *safety score*, the RSP in (Ding et al., 2012) “looks multiple steps ahead”, *i.e.*, considers all possible sequences of actions in the several following steps. In order to get probability of taking an action u at a state, the RSP in (Ding et al., 2012) needs to calculate the sum of the probabilities of all action sequences starting with u (See. Eq. (6) of (Ding et al., 2012)). Although “looking multiple steps ahead” helps the performance in (Ding et al., 2012), the number of the sequences that needs to be considered increases exponentially with the lookahead step number.

The *safety score* in (9) is also motivated by the concept of “looking multiple steps ahead.” You can intuitively think the *safety score* in (9) as some knowledge left in a state by a “vanguard” robot that takes policy μ_{null} and looks r steps ahead at every state. When a robot visits a state, it refers to this existing knowledge rather than checking the following r steps by itself, which reduces the computational cost.

There is a well-known tradeoff between *exploitation* and *exploration* in designing RSPs (Sutton and Barto, 1998). A RSP will have higher *exploitation* if it is greedier, *i.e.*, it is more likely to only pick the action with the highest desirability degree. However, in each step, the *exploration* for undesirable actions are necessary because they may be desirable in the long run. High *exploitation* and low *exploration* may result in sub-optimal solution. On the contrary, low *exploitation* and high *exploration* may reduce the convergence rate of the actor-critic algorithm. Based on the Boltzmann distribution, our RSP defined in (11) is flexible because tuning T in (11) can effectively adjust the weight of *exploration*. High temperature results in more *exploration* and vice versa. A large T also makes the RSP more randomized while a small T makes the RSP more deterministic.

3.5 Overall Algorithm

We now connect all the pieces together and present the overall algorithm giving a solution to Prob. 2.4.

Proposition 3.13. *Alg. 2 returns a $\boldsymbol{\theta}^*$ locally maximizing the probability of the RSP $\mu_{\boldsymbol{\theta}}$ satisfying the LTL formula ϕ .*

Proof. Theorem 5.4 shows that the actor-critic algorithm used in this paper returns a locally optimal $\boldsymbol{\theta}^*$ such that $\|\nabla \bar{\alpha}(\boldsymbol{\theta}^*)\| \leq \epsilon$ for a given ϵ . We have shown throughout the paper that the optimal policy maximizing the probability of reaching $S_{\mathcal{P}}^*$ on \mathcal{P} is a policy maximizing the probability of satisfying ϕ . We also showed throughout the paper that the SSP problem, as well as the RSP $\mu_{\boldsymbol{\theta}}$ can be constructed without the transition probabilities, and only with $\mathcal{M}^{\mathcal{N}}$. Therefore, Alg. 2 produces an RSP maximizing the probability of satisfying ϕ with respect to $\boldsymbol{\theta}$ up to a threshold ϵ . ■

Algorithm 2 Overall algorithm providing a solution to Prob. 2.4

Input: A labeled NTS $\mathcal{M}^{\mathcal{N}} = (Q, q_0, U, A, P^{\mathcal{N}}, \Pi, h)$ modeling a robot in a partitioned environment, LTL formula ϕ over Π , and a simulator to compute $P(q, u, \cdot)$ given a state-action pair (q, u) .

- 1: Translate the LTL formula ϕ to a DRA \mathcal{R}_ϕ .
- 2: Generate the product NTS $\mathcal{P}^{\mathcal{N}} = \mathcal{M}^{\mathcal{N}} \times \mathcal{R}_\phi$.
- 3: Find the union of all AMECs $S_{\mathcal{P}}^*$ associated with $\mathcal{P}^{\mathcal{N}}$.
- 4: Convert from an MRP to an SSP and generate $\tilde{\mathcal{P}}^{\mathcal{N}}$.
- 5: Obtain the RSP μ_θ with $\mathcal{P}^{\mathcal{N}}$.
- 6: Execute Alg. 1 with $\tilde{\mathcal{P}}^{\mathcal{N}}$ and μ_θ as inputs until $\|\nabla \bar{\alpha}(\theta^*)\| \leq \epsilon$ for a θ^* and a given ϵ .

Output: RSP μ_θ and θ^* locally maximizing the probability of satisfying ϕ with respect to θ up to a threshold ϵ .

4 Simulation Results

We test the algorithms proposed in this paper through simulation in the RIDE environment (as shown in Fig. 1). The transition probabilities are computed by an accurate simulator of RIDE as needed (cf. Sec. 4.2). We compare the results of our actor-critic method with the results of the method in (Ding et al., 2011), which is referred to as the LP method. We analyze the movements of robots under the policy calculated by our actor-critic method in a 21×21 scenario. We also analyze the increase of time and memory usages of both methods when the problem size increases.

4.1 Environment

In this case study, we consider environments with topologies such as the one shown in Fig. 3. Such environments are made of square blocks formed by corridors and intersections. In the case shown in Fig. 3, the corridors (C_1, C_2, \dots, C_{164}) are of one- or three-unit lengths. The three-unit corridors are used to build corners in the environment. The intersections (I_1, I_2, \dots, I_{84}) are of two types, three-way and four-way. The black regions in this figure represent the walls of the environment. Note that there is always a corridor between two intersections. Thus, we can recursively build larger scenes by concatenating smaller scenes and merging the one-length corridors on the borders.

There are five properties of interest associated with the regions of the environment. These properties are: **VD** = *ValuableData* (regions containing valuable data to be collected), **RD** = *RegularData* (regions containing regular data to be collected), **Up** = *Upload* (regions where data can be uploaded), **Ri** = *Risky* (regions that could pose a threat to the robot), and **Un** = *Unsafe* (regions that are unsafe for the robot). If the robot reaches an unsafe state, it will break down and will not finish the task. There are also some **reset** states in the scene. Whenever a robot reaches a **reset** state, it will be removed from the scene. A mission ends when the robot reaches a **reset** state, no matter whether the task specification has already been met or not. Our goal is to find a control strategy for the robot to maximize the probability of finishing a task specified as an LTL formula over the set of properties before the mission ends. Note that while **VD**, **RD**, **Up**, **Ri**, and **Un** are properties that comprise the LTL formula, **reset** states only trigger the end of a mission and should not appear in the LTL formula.

4.2 Construction of the MDP model

Let x_k, y_k be the position of the robot at time k . Denote by v_k and θ_k the speed and the heading of the robot, respectively. Then the movement of the robot can be described using the following

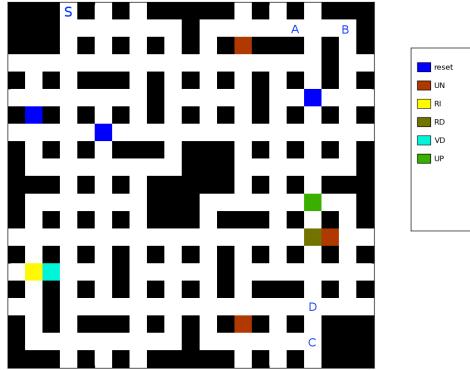


Figure 3: An example schematic representation of an environment with intersections and corridors. The black blocks represent walls, and the white regions are intersection and corridors. There are five properties of interest in the regions indicated with **VD** = *ValuableData*, **RD** = *RegularData*, **Up** = *Upload*, **Ri** = *Risky*, and **Un** = *Unsafe*. There are also some reset states in the scene. Regions with properties of interests are plotted in corresponding colors. The correspondence of properties and colors are shown in the legend. **A** and **B** are examples of corridors with length one and three units, respectively. **C** and **D** are examples of three- and four-way intersections. The initial position of the robot is marked with a blue **S**.

unicycle model.

$$\begin{aligned}
 x_{k+1} &= x_k + v_k \cos(\theta_k) + N(0, \sigma_x^2), \\
 y_{k+1} &= y_k + v_k \sin(\theta_k) + N(0, \sigma_y^2), \\
 v_{k+1} &= v_k + N(0, \sigma_v^2), \\
 \theta_{k+1} &= \theta_k + \rho(\theta_d - \theta_k) + N(0, \sigma_\theta^2),
 \end{aligned}$$

where θ_d is the desired heading, ρ is the angular feedback coefficient, and σ_θ determines the noise actuation noise. σ_v affects the speed of the robot, while σ_x and σ_y characterize the roughness of the surface in x and y directions, respectively.

The robot is equipped with a set of feedback control primitives (actions) - GoLeft, GoForward, GoRight, and GoBackward. Due to the presence of noise in the actuators and sensors, however, the resulting motion may be different than intended. Thus, the outcome of each control primitive is characterized probabilistically.

To create an MDP model of the robot in RIDE, we define each state of the MDP as a collection of two adjacent regions (a corridor and an intersection). The first region in the pair represents the last location of the robot, and the second region in the pair represents the current location. For instance, the pairs C_1 - I_2 and I_3 - C_4 are two states of the MDP. If the robot is in C_1 - I_2 at time k , then it is in intersection I_2 at time k and was in corridor C_1 at time $k-1$. If the robot is in I_3 - C_4 at time k , then it is in corridor C_4 at time k and was in intersection I_3 at time $k-1$.

Through this pairing of regions, it was shown that the Markov property (*i.e.*, the result of an action at a state depends only on the current state) can be achieved for the motion of the robot in RIDE (Lahijanian et al., 2010). The resulting MDP has 608 states. The set of actions available at a state is the set of controllers available at the second region corresponding to the state. For example, when in state C_1 - I_2 only those actions from region I_2 are allowed. Each state of the MDP whose second region satisfies a property in Π is mapped to that property.

To obtain accurate transition probabilities, we use the unicycle model described in (12) to simulate transition probabilities of the labeled MDP. Note that each MDP state is a tuple of two regions, the second one is the current region and the first one in the previous region. To simulate transition probabilities of an MDP state to its neighboring states. We put the robot in the center of the second region in the MDP state (current region), and the initial heading θ is determined by the relative position of the previous region and the current region. A simulation ends if the robot moves out of the current region to a new region. The new MDP state will have the current region as the first element and the new region as the second element. For each action available in each MDP state, we performed a total of 1,000 Monte-Carlo simulations, and use the normalized frequency as our estimate of the transition probabilities. For example, for state C_1-I_2 and action `GoLeft`, if the robot goes to C_2, C_3, C_4 for 450, 200, 350 times, respectively, then under action `GoLeft`, the transition probabilities from state C_1-I_2 to state $I_2-C_2, I_2-C_3, I_2-C_4$ are 0.45, 0.2, 0.35, respectively.

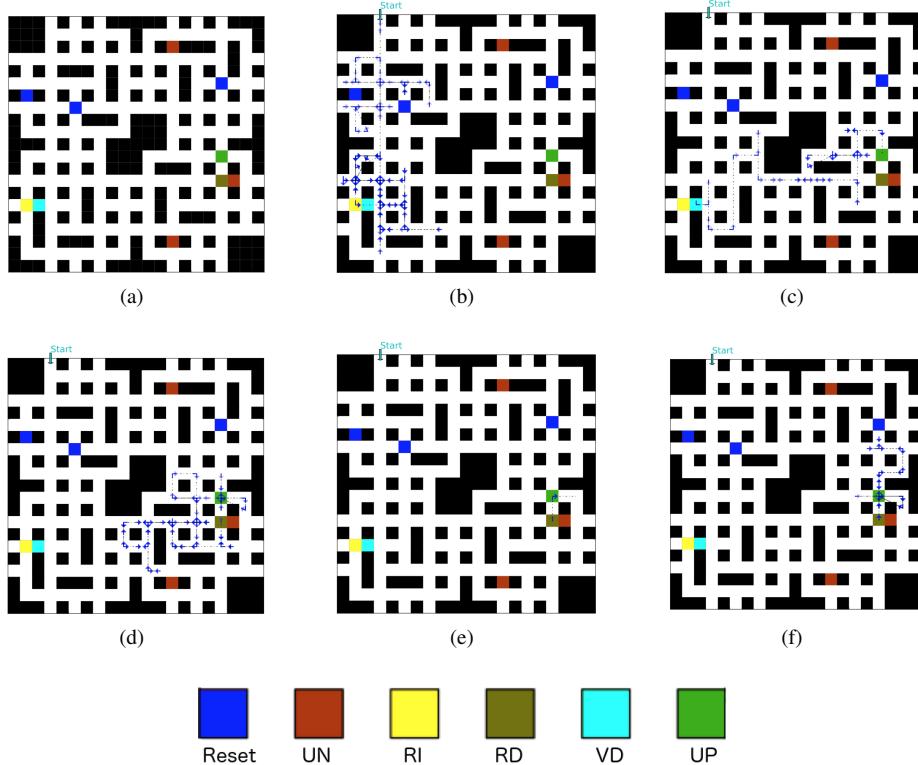


Figure 4: (a) Schematics of the 21×21 grid. An example trace: (b) The robot moves from the start point and gets valuable data (**VD**) after crossing the Risky (**Ri**) state; (c) The robot uploads (**Up**) the valuable data; (d) The robot picks up the first regular data (**RD**); (e) the robot uploads (**Up**) the first regular data; (f) The robot picks up and uploads the second regular data.

4.3 Task specification and results

We consider the following mission, composed of three sub-tasks (with no restriction on the order):

- Reach a location with *ValuableData* (**VD**), then reach *Upload* (**Up**);
- (repeat twice) Reach a location with *RegularData* (**RD**), then reach *Upload* (**Up**);

with the following requirements:

- Always avoid *Unsafe* (**Un**) regions;
- Do not reach *Risky* (**Ri**) regions unless directly preceding a location with *ValuableData* (**VD**);
- After getting either *ValuableData* (**VD**) or *RegularData* (**RD**), go to *Upload* (**Up**) before going for another *ValuableData* (**VD**) or *RegularData* (**RD**).

The above task specification can be translated to the LTL formula:

$$\begin{aligned}
\phi := & \quad \mathsf{F} \mathbf{VD} \wedge \mathsf{F} (\mathbf{RD} \wedge \mathsf{X} \mathsf{F} \mathbf{RD}) \\
& \wedge \mathsf{G} \neg \mathbf{Un} \\
& \wedge \mathsf{G} (\mathbf{Ri} \longrightarrow \mathsf{X} \mathbf{VD}) \\
& \wedge \mathsf{G} (\mathbf{VD} \vee \mathbf{RD} \longrightarrow \mathsf{X} (\neg (\mathbf{VD} \vee \mathbf{RD}) \cup \mathbf{Up})).
\end{aligned} \tag{12}$$

Note that in (12), the first line corresponds to the mission specification, and the rest correspond to the mission requirements as stated above.

We now show results of applying our algorithm for two scenarios, a 21×21 grid and a large 81×81 grid. We use the computational frameworks described in this paper to find the control strategy maximizing the probabilities of satisfying the specification. For the 21×21 grid, we visualize a movement of the robot under the policy calculated by our actor-critic method and verify that all task specifications are satisfied. We also evaluate the time and memory usage of our algorithm with the LP method in the 21×21 grid and the 81×81 grid, respectively. Compared with the LP method, our algorithm uses much less time and memory for the 81×81 grid scenario.

4.3.1 Results of a 21×21 grid scenario

In this scenario, we use the scene settings described in Sec. 4.1. Fig. 4(a) plots the properties of the states in the 21×21 grid using different colors (the legend shows the correspondence of the properties and the colors). Fig. 4 also shows an example trace, *i.e.*, a sequence of MDP states generated by the actor-critic algorithm, to verify that the robot motion is indeed correct and satisfies the mission specification and requirements. For convenience of visualization, we divide the example trace into five segments, each segment verifying a subtask in our LTL formula. In Fig. 4(b-f), paths of the robot are plotted as blue lines, and the turning directions of the robot are plotted as arrows. The transparency of the arrows are associated with the timestamps of the turns. Turns with smaller timestamps are plotted more transparently while turns with larger timestamps are plotted less transparently. We map each MDP state to the current region of the robot. For example, the MDP state $R_1 \cup I_3$ corresponds to region I_3 in the scene. The robot first moves from starting point to pick up the valuable data (**VD**) after crossing the Risky (**Ri**) state (Fig. 4(b)). After that, the robot uploads (**Up**) the valuable data (Fig. 4(c)). Later, the robot picks up the first regular data (**RD**) (Fig. 4(d)) and uploads (**Up**) the first regular data (Fig. 4(e)). The robot eventually picks up and uploads the second regular data (Fig. 4(f)).

In this case, the product MDP contains 22,533 states. There are 1,085 “goal” states and 10,582 “trap” states as defined in Def. 3.9. Note that in order to solve the probability exactly using the LP method, we need to compute transition probabilities for all 90,132 state-action pairs. In our method,

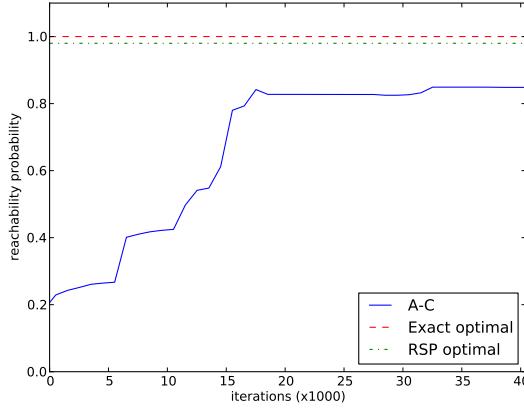


Figure 5: Reachability probability for the RSP as a function of the number of iterations applying the proposed algorithm. The exact solution (max probability of satisfying the specification) is 1. The RSP optimal solution is 0.98.

we only compute the transition probabilities along the sample path, which only consists of a very small portion of the state space. By caching the transition probabilities of previous visited states, we can furthermore reduce the computation cost for generating transition probabilities. We observe that less than 5,000 transition probabilities are needed during the simulation of the actor-critic algorithm.

The result of the LP method is a deterministic solution, which has been proved to be the exact optimal solution over all policies. The satisfaction probability is 100%. Note that our actor-critic algorithm optimizes θ over a set of parameterized RSPs to a local optimal solution. Note that $\theta = [\theta_1, \theta_2]$ is a vector of parameters, where θ_1 and θ_2 are the weights for progress and safety, respectively (cf. Eq. 10). If we assume the range of θ_1 and θ_2 , and assume both θ_1 and θ_2 are discrete values (generated by dividing the range equally), we can obtain a set of discrete values for θ . We could use a brute force method to calculate the optimal solution within all possible values of the discretized θ , which is referred to as *RSP optimal solution*. When discretization is fine enough, we could treat the result of the brute force method as the global optimal solution within the set of parameterized RSPs. In our case, we choose the difference of two consecutive discretized values for both θ_1 and θ_2 to be 0.01. The *RSP optimal solution* for the 21×21 grid scenario has 98% probability of satisfying the specification.

The difference between the exact optimal solution and the *RSP optimal solution* characterizes the expressivity of the RSP structure we used. The difference of the result of the actor-critic algorithm and the *RSP optimal solution* characterizes the effectiveness of the actor-critic algorithm in terms of optimizing the RSP.

The graph of the convergence of the actor-critic solution is shown in Fig. 5. The parameters for this examples are: $\lambda = 0.9$, and initial $\theta = [2.85, 100]^T$. The sensing radius r in the RSP is 2. The actor-critic algorithm converges after 20,000 iterations to 85% **reachability probability**. Its difference to exact and RSP optimal solutions is 15% and 13%, respectively.

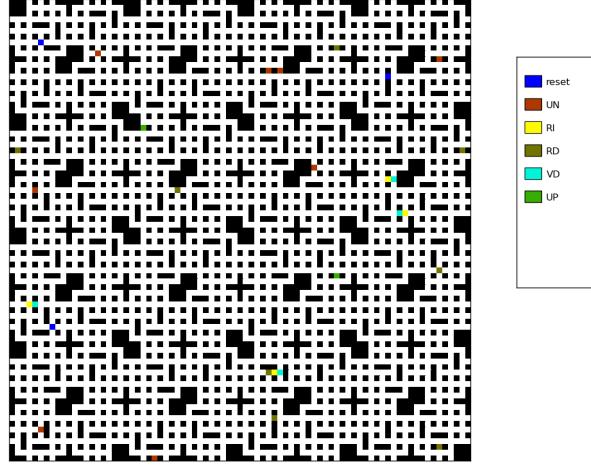


Figure 6: Schematics of the 81×81 grid. The states with properties are plotted using corresponding color (see legend for the correspondence of properties and colors).

4.3.2 Results of a 81×81 grid scenario

The 81×81 grid is shown in Fig. 6. The grid is created by concatenating 16 smaller 21×21 grids and merging the borders between the adjacent small grids. The states with properties are generated randomly. In this scenario, the product MDP contains 359,973 states. There are 19,135 “goal” states and 166,128 “trap” states as defined in Def. 3.9. By applying the LP method, we calculate the exact optimal solution and found a maximum satisfaction probability of 99%. The RSP optimal solution leads to satisfaction probability of 92%.

The initial $\theta = [1, 110]^T$. The sensing radius r in the RSP is 2. The graph of the convergence of the actor-critic solution is shown in Fig. 7. The actor-critic algorithm converges to 80% after 68,000 iterations. Its difference to exact and RSP optimal solutions is 19% and 12%, respectively.

In Problem 3.10, the period cost is zero unless the states reaches \bar{S}_P^* , in which case the cost is 1. As a result, when \bar{S}_P^* is reached, the algorithm will receive a strong penalty (bad feedback) for current policy thus make a big policy update, which results in the jumps of reachability probabilities in Fig. 7.

Grid Size	CPU Time (A-C)	CPU Time (LP)	Memory (A-C)	Memory (LP)
21×21	25 sec	6.3 sec	$\sim 0.05G$	$\sim 0.10G$
81×81	56 sec	110.4 sec	$\sim 0.22G$	$\sim 1.5G$

Table 1: Comparison between our method and exact methods.

Finally, in order to demonstrate the usefulness of our approach, in Table 1 we compare the computation time and memory needed to execute our algorithm versus using the CPLEX LP solver. Both algorithms are run on a machine with 2.50GHz 4-core Intel i5 CPU and 2GB memory. Note

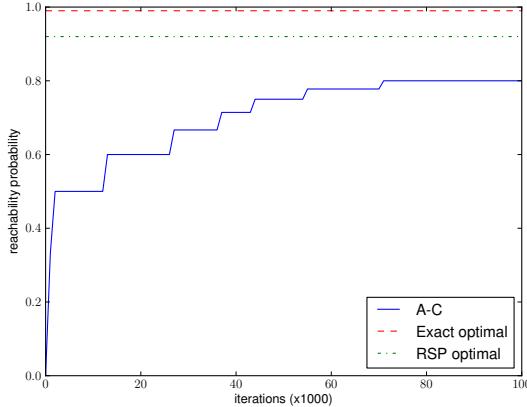


Figure 7: Reachability probability for the RSP as a function of the number of iterations for the large environment. The exact solution (maximal probability of satisfying the specification) is 0.99. The RSP optimal solution is 0.92.

that the time in the table does not include the time of calculating AMECs. For the 21×21 grid, the LP method is faster. The actor-critic algorithm needs considerable number of iterations to learn the correct parameters, even if the problem size is small, which makes it unfavorable for small problems. However, the time required for the LP method increases rapidly when the problem size increases, while the time for the actor-critic method only increases modestly. In the 81×81 grid, the LP method takes about 2 times more CPU time than our algorithm. The actor-critic method also uses much less memory than the LP method in both scenarios because only a small percentage of transition probabilities needs to be stored in the actor-critic algorithm. For the 21×21 grid, this percentage is only around 5%. Although actor-critic algorithm also needs to store some additional information like progress and safety score, the overall memory usage is still smaller. The LP method uses 15 times more memory in the 81×81 scenario than in the 21×21 scenario. At the same time, the memory usage of our method only increases by 4 times. These results show that the compared with the LP method, the actor-critic method is suitable for large-scale problems because of its reduced time and space complexity.

5 Conclusions and Future Work

We presented a framework that brings together an approximate dynamic programming computational method of the actor-critic type, with formal control synthesis for Markov Decision Processes (MDPs) from temporal logic specifications. We showed that this approach is particularly suitable for problems where the transition probabilities of the MDP are difficult or computationally expensive to compute, such as for many robotic applications. We demonstrated that this approach effectively finds an approximate optimal policy within a class of randomized stationary policies maximizing the probability of satisfying the temporal logic formula. Because our experiment setup is based on RIDE platform, the results in this paper only focus on maze-like environment. However, the techniques presented in this paper, including the conversion from MRP to SSP and the LSTD actor-critic algorithm, are general and should work beyond RIDE as long as a label MDP can be properly defined in

the environment. As a future work, we can evaluate our approach in more realistic environment, for example city roads. In addition, the approach presented in this paper is designed for controlling one robot. Another future work is to extend our approach to multi-robot teams.

Appendix: Convergence of the LSTD Actor-Critic Algorithm

We first cite the theory of *linear stochastic approximation driven by a slowly varying Markov chain* (Konda, 2002) (with simplifications).

Let $\{\mathbf{y}_k\}$ be a finite Markov chain whose transition probabilities $p_{\theta}(\cdot|\cdot)$ depend on a parameter $\theta \in \mathbb{R}^n$. Consider a generic iteration of the form

$$\mathbf{s}_{k+1} = \mathbf{s}_k + \gamma_k (\mathbf{h}_{\theta_k}(\mathbf{y}_{k+1}) - \mathbf{G}_{\theta_k}(\mathbf{y}_{k+1})\mathbf{s}_k) + \gamma_k \mathbf{\Xi}_k \mathbf{s}_k, \quad (13)$$

where $\mathbf{s}_k \in \mathbb{R}^m$, and $\mathbf{h}_{\theta}(\cdot) \in \mathbb{R}^m$, $\mathbf{G}_{\theta}(\cdot) \in \mathbb{R}^{m \times m}$ are θ -parameterized vector and matrix functions, respectively. We first present the following conditions.

Condition 5.1. (1) *The sequence $\{\gamma_k\}$ is deterministic, non-increasing, and*

$$\sum_k \gamma_k = \infty, \quad \sum_k \gamma_k^2 < \infty.$$

(2) *The random sequence $\{\theta_k\}$ satisfies $\|\theta_{k+1} - \theta_k\| \leq \beta_k H_k$ for some process $\{H_k\}$ with bounded moments, where $\{\beta_k\}$ is a deterministic sequence such that*

$$\sum_k \left(\frac{\beta_k}{\gamma_k} \right)^d < \infty \quad \text{for some } d > 0.$$

(3) *$\mathbf{\Xi}_k$ is an $m \times m$ -matrix valued martingale difference with bounded moments.*

(4) *For each θ , there exist $\bar{\mathbf{h}}(\theta) \in \mathbb{R}^m$, $\bar{\mathbf{G}}(\theta) \in \mathbb{R}^{m \times m}$, and corresponding m -vector and $m \times m$ -matrix functions $\hat{\mathbf{h}}_{\theta}(\cdot)$, $\hat{\mathbf{G}}_{\theta}(\cdot)$ that satisfy the Poisson equation. That is, for each \mathbf{y} ,*

$$\hat{\mathbf{h}}_{\theta}(\mathbf{y}) = \mathbf{h}_{\theta}(\mathbf{y}) - \bar{\mathbf{h}}(\theta) + \sum_{\mathbf{z}} p_{\theta}(\mathbf{z}|\mathbf{y}) \hat{\mathbf{h}}_{\theta}(\mathbf{z}),$$

$$\hat{\mathbf{G}}_{\theta}(\mathbf{y}) = \mathbf{G}_{\theta}(\mathbf{y}) - \bar{\mathbf{G}}(\theta) + \sum_{\mathbf{z}} p_{\theta}(\mathbf{z}|\mathbf{y}) \hat{\mathbf{G}}_{\theta}(\mathbf{z}).$$

Denote by $E_{\theta}[\cdot]$ the expectation with respect to the stationary distribution of the finite Markov chain $\{\mathbf{y}_k\}$, then $\bar{\mathbf{G}}(\theta_k) = E_{\theta}[\mathbf{G}_{\theta_k}(\mathbf{y})]$ and $\bar{\mathbf{h}}(\theta_k) = E_{\theta}[\mathbf{h}_{\theta}(\mathbf{y})]$.

- (5) *For some constant C and for all θ , we have $\max(\|\bar{\mathbf{h}}(\theta)\|, \|\bar{\mathbf{G}}(\theta)\|) \leq C$.*
- (6) *For any $d > 0$, there exists $C_d > 0$ such that $\sup_k \mathbf{E}[\|\mathbf{f}_{\theta_k}(\mathbf{y}_k)\|^d] \leq C_d$, where $\mathbf{f}_{\theta}(\cdot)$ represents any of the functions $\hat{\mathbf{h}}_{\theta}(\cdot)$, $\mathbf{h}_{\theta}(\cdot)$, $\hat{\mathbf{G}}_{\theta}(\cdot)$ and $\mathbf{G}_{\theta}(\cdot)$.*
- (7) *For some constant $C > 0$ and for all $\theta, \bar{\theta} \in \mathbb{R}^n$, $\max(\|\bar{\mathbf{h}}(\theta) - \bar{\mathbf{h}}(\bar{\theta})\|, \|\bar{\mathbf{G}}(\theta) - \bar{\mathbf{G}}(\bar{\theta})\|) \leq C\|\theta - \bar{\theta}\|$.*
- (8) *There exists a positive measurable function $C(\cdot)$ such that for every $d > 0$, $\sup_k \mathbf{E}[C(\mathbf{y}_k)^d] < \infty$, and $\|\mathbf{f}_{\theta}(\mathbf{y}) - \mathbf{f}_{\bar{\theta}}(\mathbf{y})\| \leq C(\mathbf{y})\|\theta - \bar{\theta}\|$.*

(9) There exists $a > 0$ such that for all $\mathbf{s} \in \mathbb{R}^m$ and $\boldsymbol{\theta} \in \mathbb{R}^n$

$$\mathbf{s}' \bar{\mathbf{G}}(\boldsymbol{\theta}) \mathbf{s} \geq a \|\mathbf{s}\|^2.$$

It has been shown in (Konda, 2002) that the critic in (13) converges if Condition 5.1(1-9) are met.

Theorem 5.2. [Convergence of Linear Stochastic Approximation] If Condition 5.1(1-9) are satisfied, then

$$\lim_{k \rightarrow \infty} |\bar{\mathbf{G}}(\boldsymbol{\theta}_k) \mathbf{s}_k - \bar{\mathbf{h}}(\boldsymbol{\theta}_k)| = 0. \quad (14)$$

Proof. See Chapter 3 of (Konda, 2002). \blacksquare

Based on Theorem 5.2, we present the convergence of the critic in Alg. 1 in the following theorem.

Theorem 5.3. [Critic Convergence] For the LSTD actor-critic method described in Alg. 1 with some deterministic and non-increasing step-size sequences $\{\beta_k\}$, $\{\gamma_k\}$ satisfying:

$$\sum_k \beta_k = \infty, \quad \sum_k \beta_k^2 < \infty, \quad \lim_{k \rightarrow \infty} \frac{\beta_k}{\gamma_k} = 0, \quad (15)$$

the sequence \mathbf{s}_k is bounded, and

$$\begin{aligned} \lim_{k \rightarrow \infty} |\mathbf{b}_k - E_{\boldsymbol{\theta}}[g(x, u) \mathbf{z}]| &= 0, \\ \lim_{k \rightarrow \infty} |\mathbf{v}(\mathbf{A}_k) - E_{\boldsymbol{\theta}}[\mathbf{v}(\mathbf{z}((P_{\boldsymbol{\theta}} \boldsymbol{\psi}'_{\boldsymbol{\theta}})(x, u) - \boldsymbol{\psi}'_{\boldsymbol{\theta}}(x, u)))]| &= 0, \end{aligned}$$

where $E_{\boldsymbol{\theta}}[\cdot]$ is the expectation with respect to the stationary distribution of the Markov chain $\{x_k, u_k, \mathbf{z}_k\}$, and for any matrix \mathbf{A} , $\mathbf{v}(\mathbf{A})$ is a column vector that stacks all row vectors of \mathbf{A} (also written as column vectors).

Proof. Simple algebra suggests that the critic update in Alg. 1 can be written as the form of (13) with

$$\begin{aligned} \mathbf{s}_k &= \begin{bmatrix} \mathbf{b}_k \\ \mathbf{v}(\mathbf{A}_k) \\ 1 \end{bmatrix}, \quad \mathbf{y}_k = (x_k, u_k, \mathbf{z}_k), \\ \mathbf{h}_{\boldsymbol{\theta}}(\mathbf{y}) &= \begin{bmatrix} g(x, u) \mathbf{z} \\ \mathbf{v}(\mathbf{z}((P_{\boldsymbol{\theta}} \boldsymbol{\psi}'_{\boldsymbol{\theta}})(x, u) - \boldsymbol{\psi}'_{\boldsymbol{\theta}}(x, u))) \\ 1 \end{bmatrix}, \\ \mathbf{G}_{\boldsymbol{\theta}}(\mathbf{y}) &= [\mathbf{I}], \\ \boldsymbol{\Xi}_k &= \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & D_k \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}, \end{aligned} \quad (16)$$

where \mathbf{A}_k , \mathbf{b}_k , x_k , u_k , and \mathbf{z}_k are the iterates defined in Alg. 1, $\mathbf{y} = (x, u, \mathbf{z})$ denotes a value of the triplet \mathbf{y}_k , and $D_k = \mathbf{v}(\mathbf{z}_k(\boldsymbol{\psi}'_{\boldsymbol{\theta}_k}(x_{k+1}, u_{k+1}) - (P_{\boldsymbol{\theta}} \boldsymbol{\psi}_{\boldsymbol{\theta}})'(x_k, u_k)))$.

The step-sizes γ_k and β_k in Alg. 1 correspond exactly to the γ_k and β_k in Condition 5.1.(1) and 5.1.(2), respectively. For MDPs with finite state and action space, Condition 5.1.(1-2) reduces to (15).

A direct yet verbose way to prove the theorem is to verify Conditions 5.1.(3)-(9). However, a comparison with the convergence proof for the TD(λ) critic in (Konda and Tsitsiklis, 2003) gives a simpler proof. Let

$$\mathbf{F}_\theta(\mathbf{y}) = \mathbf{z}(\psi'_\theta(x, u) - (P_\theta\psi_\theta)'(x, u)).$$

While proving the convergence of TD(λ) critic operating concurrently with the actor, (Konda and Tsitsiklis, 2003) showed that

$$\tilde{\mathbf{h}}_\theta(\mathbf{y}) = \begin{bmatrix} \tilde{h}_\theta^{(1)}(\mathbf{y}) \\ \tilde{h}_\theta^{(2)}(\mathbf{y}) \end{bmatrix} = \begin{bmatrix} Mg(x, u) \\ g(x, u)\mathbf{z} \end{bmatrix}, \quad (17)$$

$$\tilde{\mathbf{G}}_\theta(\mathbf{y}) = \begin{bmatrix} 1 & \mathbf{0} \\ \mathbf{z}/M & \mathbf{F}_\theta(\mathbf{y}) \end{bmatrix}, \quad (18)$$

and

$$\tilde{\mathbf{\Xi}}_k = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{v}(\mathbf{z}_k(\psi'_{\theta_k}(x_{k+1}, u_{k+1}) - (P_\theta\psi_\theta)'(x_k, u_k))) \end{bmatrix} \quad (19)$$

satisfy Condition 5.1.(3)-5.1.(8). M in (17) and (18) is an arbitrary (large) positive constant whose role is to facilitate the convergence proof. In our case, (16) can be rewritten as

$$\mathbf{h}_\theta(\mathbf{y}) = \begin{bmatrix} \tilde{\mathbf{h}}_\theta^{(2)}(\mathbf{y}) \\ -\mathbf{v}(\mathbf{F}_\theta(\mathbf{y})) \\ 1 \end{bmatrix}, \quad \mathbf{G}_\theta(\mathbf{y}) = [\mathbf{I}], \quad \mathbf{\Xi}_k = \begin{bmatrix} \tilde{\mathbf{\Xi}}_k \\ \mathbf{0} \end{bmatrix}. \quad (20)$$

Note that although $\mathbf{h}_\theta(\mathbf{y})$, $\mathbf{G}_\theta(\mathbf{y})$, and $\mathbf{\Xi}_k$ in (20) are very different from $\tilde{\mathbf{h}}_\theta(\mathbf{y})$, $\tilde{\mathbf{G}}_\theta(\mathbf{y})$, and $\tilde{\mathbf{\Xi}}_k$ in (17), (18) and (19), they involve the same quantities and both in a linear fashion. So, $\mathbf{h}_\theta(\cdot)$, $\mathbf{G}_\theta(\cdot)$ and $\mathbf{\Xi}_k$ also satisfy conditions 5.1.(3)-5.1.(8). Meanwhile, the step-size $\{\gamma_k\}$ satisfies condition 5.1.(1), and the step-size $\{\beta_k\}$ satisfies Eq. (15) (which is as explained above implies condition 5.1.(2)). Now, only condition 5.1.(9) remains to be checked. To that end, note that all diagonal elements of $\mathbf{G}_\theta(\mathbf{y})$ equal to one, so, $\mathbf{G}_\theta(\mathbf{y})$ is positive definite. This proves the convergence. Using the same correspondence and the result in (Konda and Tsitsiklis, 2003), one can further check that (14) also holds here. The theorem can be proved by substituting (16) to (14) and using the fact that $\bar{\mathbf{G}}(\theta_k) = E_\theta[\mathbf{G}_{\theta_k}(\mathbf{y})]$ and $\bar{\mathbf{h}}(\theta_k) = E_\theta[\mathbf{h}_\theta(\mathbf{y})]$. ■

Theorem 5.3 states that the critic step in Alg. 1 can converge to the optimal \mathbf{r} . In addition, note that λ in Step 6 of Alg 1 is a decay factor used in Sutton's eligibility trace, and has been applied in other approximate dynamic programming methods like TD(λ) (Sutton and Barto, 1998). For any sequence $\{c_k\}$, let $\liminf_{k \rightarrow \infty} c_k = \lim_{k \rightarrow \infty} \inf\{c_m : m \geq k\}$. Then the following theorem establishes the convergence for the actor step of Alg. 1.

Theorem 5.4. [Actor Convergence] For the LSTD actor-critic algorithm (Alg. 1) with some step-size sequence $\{\beta_k\}$ satisfying (15), for any $\epsilon > 0$, there exists some λ sufficiently close to 1, such that $\liminf_{k \rightarrow \infty} \|\nabla \bar{\alpha}(\theta_k)\| < \epsilon$ w.p.1. That is, θ_k visits an arbitrary neighborhood of a stationary point infinitely often.

Proof. Since critic convergence has been proved (cf. Theorem 5.3), the result follows by setting $\phi_\theta = \psi_\theta$ and following the proof in Sec. 6 of (Konda and Tsitsiklis, 2003). ■

References

C. Baier, M. Größer, M. Leucker, B. Bollig, and F. Ciesinski. Controller synthesis for probabilistic systems. In *Exploring New Frontiers of Theoretical Informatics*, pages 493–506. Springer, 2004.

C. Baier, J.-P. Katoen, et al. *Principles of model checking*, volume 26202649. MIT press Cambridge, 2008.

A. Barto, R. Sutton, and C. Anderson. Neuronlike adaptive elements that can solve difficult learning control problems. *IEEE Transactions on Systems, Man, & Cybernetics*, 1983.

D. Bertsekas and J. Tsitsiklis. Neuro-dynamic programming. 1996.

D. Bertsekas, D. Bertsekas, D. Bertsekas, and D. Bertsekas. Dynamic programming and optimal control. 1995.

A. Bhatia, L. E. Kavraki, and M. Y. Vardi. Sampling-based motion planning with temporal goals. In *IEEE International Conference on Robotics and Automation (ICRA) 2010*, pages 2689–2696. IEEE, 2010.

J. A. Boyan. Least-squares temporal difference learning. In *ICML*, pages 49–56. Citeseer, 1999.

E. M. Clarke, O. Grumberg, and D. Peled. *Model checking*. MIT press, 1999.

T. H. Cormen, C. E. Leiserson, R. L. Rivest, C. Stein, et al. *Introduction to algorithms*, volume 2. MIT press Cambridge, 2001.

C. Courcoubetis and M. Yannakakis. Markov decision processes and regular events. In *Automata, Languages and Programming*, pages 336–349. Springer, 1990.

L. De Alfaro. *Formal verification of probabilistic systems*. PhD thesis, Citeseer, 1997.

X. C. Ding, S. L. Smith, C. Belta, and D. Rus. LTL control in uncertain environments with probabilistic satisfaction guarantees. In *18th IFAC World Congress*, number 1, pages 3515–3520, 2011.

X. C. Ding, J. Wang, M. Lahijanian, I. C. Paschalidis, and C. A. Belta. Temporal logic motion control using actor-critic methods. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 4687–4692. IEEE, 2012.

R. M. Estanjini, K. Li, and I. C. Paschalidis. A least squares temporal difference actor-critic algorithm with applications to warehouse management. *Naval Research Logistics (NRL)*, 59(3-4):197–211, 2012.

J. Fu and U. Topcu. Probably approximately correct mdp learning and control with temporal logic constraints. *Proceedings of Robotics: Science and Systems*, 2014.

E. Gradel, W. Thomas, and T. Wilke. A guide to current research, ser. lecture notes in computer science., 2002.

S. Karaman and E. Frazzoli. Sampling-based motion planning with deterministic μ -calculus specifications. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, pages 2222–2229. IEEE, 2009.

V. R. Konda and J. N. Tsitsiklis. On actor-critic algorithms. *SIAM journal on Control and Optimization*, 42(4):1143–1166, 2003.

V. V. G. Konda. *Actor-critic algorithms*. PhD thesis, Massachusetts Institute of Technology, 2002.

H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas. Where’s Waldo? Sensor-based temporal logic motion planning. In *Robotics and Automation, 2007 IEEE International Conference on*, pages 3116–3121. IEEE, 2007.

M. Lahijanian, J. Wasniewski, S. B. Andersson, and C. Belta. Motion planning and control from temporal logic specifications with probabilistic satisfaction guarantees. In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, pages 3227–3232. IEEE, 2010.

J. Liu, N. Ozay, U. Topcu, and R. M. Murray. Synthesis of reactive switching protocols from temporal logic specifications. *Automatic Control, IEEE Transactions on*, 58(7):1771–1785, 2013.

S. G. Loizou and K. J. Kyriakopoulos. Automatic synthesis of multi-agent motion tasks based on LTL specifications. In *Decision and Control, 2004. CDC. 43rd IEEE Conference on*, volume 1, pages 153–158. IEEE, 2004.

R. Luna, M. Lahijanian, M. Moll, and L. E. Kavraki. Optimal and efficient stochastic motion planning in partially-known environments. 2014.

R. Moazzez Estanjini, X. C. Ding, M. Lahijanian, J. Wang, C. A. Belta, and I. C. Paschalidis. Least squares temporal difference actor-critic methods with applications to robot motion control. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 704–709. IEEE, 2011a.

R. Moazzez Estanjini, X. C. Ding, M. Lahijanian, J. Wang, C. A. Belta, and I. C. Paschalidis. Least squares temporal difference actor-critic methods with applications to robot motion control. In *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, pages 704–709. IEEE, 2011b.

M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*, volume 414. John Wiley & Sons, 1994.

M. M. Quottrup, T. Bak, and R. Zamanabadi. Multi-robot planning: A timed automata approach. In *Robotics and Automation, 2004 IEEE International Conference on*, volume 5, pages 4417–4422. IEEE, 2004.

J. J. Rutten, P. Panangaden, and F. Van Breugel. *Mathematical techniques for analyzing concurrent and probabilistic systems*. Number 23. American Mathematical Soc., 2004.

D. Sadigh, E. S. Kim, S. Coogan, S. S. Sastry, and S. A. Seshia. A learning based approach to control synthesis of markov decision processes for linear temporal logic specifications. *Proceedings of the 53rd IEEE Conference on Decision and Control (CDC)*, 2014.

J. Si. *Handbook of learning and approximate dynamic programming*, volume 2. Wiley-IEEE Press, 2004.

R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT Press, 1998.

M. Y. Vardi. Probabilistic linear-time model checking: An overview of the automata-theoretic approach. In *Formal Methods for Real-Time and Probabilistic Systems*, pages 265–276. Springer, 1999.

E. M. Wolff, U. Topcu, and R. M. Murray. Robust control of uncertain markov decision processes with temporal logic specifications. In *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pages 3372–3379. IEEE, 2012.

T. Wongpiromsarn, U. Topcu, and R. M. Murray. Receding horizon temporal logic planning for dynamical systems. In *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, pages 5997–6004. IEEE, 2009.