

# Average Reward Reinforcement Learning for Omega-Regular and Mean-Payoff Objectives

MILAD KAZEMI, King’s College London, UK

MATEO PEREZ, University of Colorado Boulder, USA

FABIO SOMENZI, University of Colorado Boulder, USA

SADEGH SOUDJANI, University of Birmingham, UK, and MPI-SWS, Germany

ASHUTOSH TRIVEDI, University of Colorado Boulder, USA

ALVARO VELASQUEZ, University of Colorado Boulder, USA

Recent advances in reinforcement learning (RL) have renewed focus on the design of reward functions that shape agent behavior. Manually crafting such functions is often tedious and error-prone. A more principled alternative is to specify behavioral requirements using a formal, unambiguous language that can be automatically translated into a reward function. Omega-regular languages are a natural choice for this purpose, given their established role in formal verification and synthesis. However, existing approaches using omega-regular specifications typically rely on discounted reward RL in an episodic setting, where the environment is periodically reset to an initial state during learning. This setup is misaligned with the semantics of omega-regular specifications, which describe properties over infinite behavior traces. In such cases, the average reward criterion and the continuing setting—where the agent interacts with the environment over a single, uninterrupted lifetime—are more appropriate.

To address the challenges of infinite-horizon, continuing tasks, we restrict our focus to the subclass of omega-regular languages known as absolute liveness specifications. These specifications cannot be violated by any finite prefix of the agent’s behavior, aligning naturally with the continuing setting. We present the first model-free RL framework that translates absolute liveness specifications to average-reward objectives. In contrast to prior work, our approach enables learning in communicating Markov Decision Processes without episodic resetting. We further introduce a reward structure for lexicographic multi-objective optimization, where the goal is to maximize an external average-reward objective among the policies that also maximize the satisfaction probability of a given absolute liveness omega-regular specification. Our method guarantees convergence in unknown communicating MDPs and supports on-the-fly reductions that do not require full knowledge of the environment, thus enabling model-free RL. Empirical results across various benchmarks demonstrate that our average-reward approach in the continuing setting is more effective than competing methods based on discounting.

**JAIR Track:** Integration of Logical Constraints in Deep Learning

**JAIR Associate Editor:** Bettina Koenighofer

## JAIR Reference Format:

Milad Kazemi, Mateo Perez, Fabio Somenzi, Sadegh Soudjani, Ashutosh Trivedi, and Alvaro Velasquez. 2026. Average Reward Reinforcement Learning for Omega-Regular and Mean-Payoff Objectives. *Journal of Artificial Intelligence Research* 85, Article 25 (March 2026), 29 pages. DOI: [10.1613/jair.1.19233](https://doi.org/10.1613/jair.1.19233)

Authors’ Contact Information: Milad Kazemi, ORCID: [0000-0003-2716-2758](https://orcid.org/0000-0003-2716-2758), [milad.kazemi@kcl.ac.uk](mailto:milad.kazemi@kcl.ac.uk), King’s College London, London, UK; Mateo Perez, ORCID: [0000-0003-4220-3212](https://orcid.org/0000-0003-4220-3212), [Mateo.Perez@colorado.edu](mailto:Mateo.Perez@colorado.edu), University of Colorado Boulder, Colorado, Boulder, USA; Fabio Somenzi, ORCID: [0000-0002-2085-2003](https://orcid.org/0000-0002-2085-2003), [fabio@colorado.edu](mailto:fabio@colorado.edu), University of Colorado Boulder, Colorado, Boulder, USA; Sadegh Soudjani, ORCID: [0000-0003-1922-6678](https://orcid.org/0000-0003-1922-6678), [sadegh@mpi-sws.org](mailto:sadegh@mpi-sws.org), University of Birmingham, UK, and MPI-SWS, Kaiserslautern, Germany; Ashutosh Trivedi, ORCID: [0000-0001-9346-0126](https://orcid.org/0000-0001-9346-0126), [Ashutosh.Trivedi@colorado.edu](mailto:Ashutosh.Trivedi@colorado.edu), University of Colorado Boulder, Colorado, Boulder, USA; Alvaro Velasquez, ORCID: [0000-0001-6757-105X](https://orcid.org/0000-0001-6757-105X), [Alvaro.Velasquez@colorado.edu](mailto:Alvaro.Velasquez@colorado.edu), University of Colorado Boulder, Colorado, Boulder, USA.



This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

© 2026 Copyright held by the owner/author(s).

DOI: [10.1613/jair.1.19233](https://doi.org/10.1613/jair.1.19233)

## 1 Introduction

Reinforcement learning (RL) (Sutton and Barto 2018) is a foundational framework for sequential decision making under uncertainty, where agents learn to optimize their behavior through trial-and-error interaction with the environment. While RL has achieved superhuman performance in domains such as board games (Schrittwieser et al. 2020; Silver, Huang, et al. 2016; Silver, Hubert, et al. 2018; Silver, Schrittwieser, et al. 2017), robotics (Tang et al. 2025; Zhao et al. 2020), and resource optimization (Mao et al. 2016; Mirhoseini et al. 2021), much of this success has been in *episodic tasks*, i.e., settings where interactions naturally reset after a finite sequence, and learning proceeds by leveraging repeated experience across episodes. In contrast, many real-world applications—such as autonomous monitoring, industrial process control, or mission planning—are inherently continuing tasks, where the agent must learn over a single, unending interaction with its environment. For such tasks, short-term rewards often fail to capture the long-term behavioral goals of interest. Instead, it is more natural to specify objectives in terms of structured, temporally extended criteria. Formal languages, particularly  $\omega$ -regular languages recognized by automata over infinite traces (C. Baier and Katoen 2008), provide a rich and precise way to express such long-term goals. This has led to increasing interest in integrating formal methods with RL for the synthesis of correct-by-construction policies (Belta and Sadraddini 2019; Jagtap et al. 2020; Majumdar, Mallik, Schmuck, et al. 2021), enabling logic-based specifications to guide data-driven learning. This paper explores how *model-free average-reward RL*, a formulation especially suited for continuing tasks (Mahadevan 1996; Wan, Naik, et al. 2021), can be used to synthesize policies that satisfy  $\omega$ -regular objectives.

### 1.1 Why $\omega$ -Regular Objectives?

While many successful applications of RL rely on reward functions that depend only on the agent’s current state and action, it is often necessary—or more natural—to account for the agent’s history when defining learning objectives. This becomes especially important in settings with sparse rewards (Neider et al. 2021), partial observability (Toro Icarte et al. 2019), or temporally extended goals (Camacho, Icarte, et al. 2019). In such cases, the agent’s goal is better expressed as a sequence of desirable or undesirable behaviors rather than as immediate outcomes. Formal language structures—particularly automata over sequences of observations—have emerged as a powerful tool for specifying these non-Markovian objectives. Rooted in formal verification (M. Y. Vardi 1985), this approach enables the use of automata with a variety of acceptance conditions to capture complex reward structures and behavioral constraints. In some cases, even natural language objectives can be systematically translated into automata representations (Brunello et al. 2019).

This line of work has gained significant traction in the formal synthesis of control policies (Belta and Sadraddini 2019; Jagtap et al. 2020; Lavaei, Somenzi, et al. 2020; Majumdar, Mallik, Schmuck, et al. 2021), where developers define high-level goals in a formal language and synthesis algorithms compute correct-by-construction policies—eliminating the need for manual, error-prone reward engineering (Kress-Gazit et al. 2018). In this paper, we adopt this paradigm to synthesize policies using model-free RL for a class of non-Markovian objectives defined by  $\omega$ -regular languages (C. Baier and Katoen 2008). They capture infinite sequences of semantically meaningful observations and provide a rich formalism for specifying long-term behavioral goals.

To operationalize such specifications within RL, we introduce nondeterministic reward machines, which encode the reward structure of  $\omega$ -regular automata. By constructing a product between the agent’s environment—modeled as a Markov Decision Process (MDP)—and the automaton, we reduce the synthesis problem to a Markovian task that can be addressed using standard model-free algorithms. This builds on the broader use of reward machines (Icarte, T. Klassen, et al. 2018), which represent reward functions as automata over histories of semantically meaningful events. These machines provide a structured and interpretable way to incorporate temporal logic into RL and enable the transformation of non-Markovian objectives into equivalent Markovian

formulations over an augmented state space (Gaon and Brafman 2020; Xu, Wu, et al. 2021). This transformation allows existing RL techniques to be applied to long-horizon, history-dependent tasks without modification.

## 1.2 Limitations of Discounted Reward in Continuing Tasks.

Despite the widespread use of discounted reward formulations in RL, they are often ill-suited for continuing tasks. The discounted return, defined as the sum of future rewards scaled by a discount factor, ensures mathematical tractability by bounding the total return over infinite horizons. However, this formulation inherently prioritizes short-term gains over long-term performance, which misaligns with the objectives of many continuing systems. In practice, achieving a suitable approximation of long-run behavior requires setting the discount factor very close to one. Yet, doing so weakens the contraction properties of learning algorithms, leading to slow convergence and potential instability.

These issues are particularly pronounced in continuing environments, where no natural episode boundaries exist. While discounted RL has shown remarkable success in episodic domains (Mirhoseini et al. 2021; Mnih et al. 2015; Silver, Huang, et al. 2016), its solutions are sensitive to initial state distributions—an undesirable property in continuing settings where learning should be state-agnostic over the long run. Furthermore, discounted formulations in continuing tasks are often incompatible with function approximation techniques commonly used in large-scale RL. In (Naik et al. 2019), the authors argue that in continuing tasks, discounted RL is fundamentally incompatible with function approximation and does not constitute a proper optimization problem. Unlike tabular settings where any policy is representable, most practical RL problems have state or action spaces that are too large for tables. This necessitates function approximation to find the best representable policy. However, without a well-posed objective function, there is often no representable policy that is unambiguously superior to all others (Naik et al. 2019).

Discounted RL algorithms for continuing tasks often rely on increasing the discount factor close to one to approximate the Blackwell-optimal policy, and they will succeed particularly when action gaps are large (Bellemare et al. 2016). However, two major obstacles hinder the use of discounted-value algorithms as proxies for average-value algorithms in practice. First, most algorithms that learn discounted value functions become increasingly unstable as the discount factor approaches one. Second, estimating the required critical discount factor is difficult since it is tied to the environment’s unknown dynamics (Naik et al. 2019). Finally, the complexity of representing the optimal solution for long-horizon problems is studied in (Lehnert et al. 2018) by isolating the representation challenge from the learning challenge. The authors show that function approximation methods may fail if action gaps collapse, which can occur over long planning horizons. The approximator may lack the necessary “resolution” to distinguish the optimal action from sub-optimal ones.

Given the growing reliance on large neural architectures in modern RL, the limitations of discounted RL motivate the need for alternative formulations aligned with the demands of continuing tasks. These limitations motivate the need to explore average-reward RL as a more principled alternative in continuing settings.

## 1.3 Average-Reward RL

A natural alternative to discounting in continuing environments is to optimize the agent’s average reward over time. This formulation aligns more directly with long-term performance objectives, making it particularly well-suited for tasks without episodic resets. However, average-reward RL introduces unique challenges. Unlike discounted RL—where the discount factor ensures convergence through a built-in contraction—average-reward methods rely on structural properties of the underlying MDP. In particular, the convergence of model-free average-reward algorithms (Mahadevan 1996; Wan, Naik, et al. 2021; Wei et al. 2020) typically requires the MDP to be communicating, meaning every state is reachable from every other state under some policy. This communicating assumption is especially important—and potentially problematic—when synthesizing policies for  $\omega$ -regular

specifications. The synthesis process involves constructing a product MDP between the agent’s environment and the automaton representing the specification. While the original MDP may satisfy the communicating property, the product MDP may not, creating a barrier to applying average-reward RL directly in such settings.

#### 1.4 Absolute Liveness Properties.

Absolute liveness specifications form an important subclass of  $\omega$ -regular specifications and are prevalent throughout the temporal property hierarchy (Alpern and Schneider 1985; Manna and Pnueli 1990). These specifications are prefix-independent—their satisfaction is unaffected by the addition of finite prefixes to an infinite word—making them well-suited for continuing tasks. Moreover, when the underlying MDP is communicating, absolute liveness specifications are satisfied with probability zero or one.

A key observation is that many temporally extended objectives, including those expressed in Linear Temporal Logic (LTL), can be framed in terms of absolute liveness semantics. In particular, a satisfiable LTL formula  $\varphi$  is an absolute liveness property if it is expressively equivalent to  $F\varphi$ , where  $F\varphi$  denotes eventual satisfaction of  $\varphi$ . Thus, the eventual-satisfaction semantics of an arbitrary  $\omega$ -regular or LTL specification  $\varphi$  can be captured by the absolute liveness property  $F\varphi$ , motivating their use in long-horizon learning tasks.

#### 1.5 Contributions.

The main contribution of this paper is an average-reward model-free reinforcement learning algorithm for synthesizing policies that satisfy a given absolute liveness  $\omega$ -regular specification. Our approach preserves the communicating property in the product construction, thereby enabling the learning of optimal policies without requiring episodic resetting. Although we assume the underlying MDP is communicating, a naive synchronization of the MDP with the automaton does not, in general, preserve this property. To address this, we introduce a reward machine construction and an augmented specification that guarantee the synchronized product MDP remains communicating. This paper makes the following primary contributions:

- (1) We present average-reward model-free RL algorithms for:
  - (a) synthesizing policies that satisfy a given *absolute liveness*  $\omega$ -regular specification, and
  - (b) solving lexicographic multi-objective optimizations in which the goal is to maximize a mean-payoff objective among the set of satisfying policies.
- (2) We study the structure of  $\omega$ -regular specifications, including safety, liveness, absolute liveness, fairness, and stability. We introduce novel automata-theoretic characterizations of *absolute liveness* (Lemma 3.5) and *stable* specifications (Lemma 3.4), which may be of independent interest.
- (3) We analyze the convergence of our algorithms and demonstrate that naive synchronization of a communicating MDP with a specification automaton can break the communicating property. To address this, we construct *reward machines* that preserve communication, enabling optimal learning without episodic resets.
- (4) Our work is the *first to establish a formal reduction from absolute liveness  $\omega$ -regular objectives to average-reward RL* in both single- and multi-objective settings with convergence guarantees.
- (5) We generalize our prior results from communicating MDPs (Kazemi, Perez, et al. 2022) to the broader class of *weakly communicating MDPs*, thereby extending the applicability of our framework to environments where some states may not be mutually reachable, yet still support meaningful long-run average-reward optimization.
- (6) We implement our proposed reduction using Differential Q-learning and evaluate it on a suite of communicating MDPs with absolute liveness specifications. Unlike prior methods, our approach composes the product MDP on-the-fly without requiring episodic resetting. Empirical results show that our method reliably converges to optimal policies in the continuing setting, even when prior approaches fail due to

non-communicating product MDPs. In the episodic setting, our approach remains competitive in training time while requiring no resets, highlighting its robustness and practical effectiveness.

To the best of our knowledge, this is the first framework that provides a formal translation from absolute liveness  $\omega$ -regular objectives to average-reward RL with convergence guarantees. Preliminary results were reported at the AAMAS conference (Kazemi, Perez, et al. 2022); the current paper extends that work with complete proofs, generalizations (weakly communicating MDPs), and novel results on multi-objective synthesis.

## 1.6 Organization.

The paper is organized as follows. Section 2 includes the preliminaries, qualitative and quantitative objectives, RL for continuing tasks, and the problem statements. Section 3 provides our results on qualitative objectives suitable for continual learning. Section 4 establishes our theoretical results on average-reward RL for qualitative objectives. Section 5 presents our algorithm for average-reward RL with lexicographic objectives, where the primary goal is to satisfy a qualitative specification and the secondary goal is to optimize a quantitative average-reward objective. In Section 6, we test the performance of our approach on different case studies and compare it against prior techniques. Section 7 discusses related work before concluding in Section 8.

## 2 Preliminaries and Problem Statement

This section presents the foundational concepts and formalizes the problem setting. We begin with Markov Decision Processes (MDPs) as the underlying model of the environment in Subsection 2.1. In Subsection 2.2, we introduce probabilistic reward machines, which serve as quantitative objectives for RL-based policy synthesis. Subsection 2.3 covers  $\omega$ -regular specifications as qualitative objectives for policy synthesis. We then formulate the formal problem statements in the context of reinforcement learning for continuing tasks in Subsection 2.4, and conclude by discussing the challenges of average-reward RL in Subsection 2.5.

### 2.1 The Environment: Markov Decision Processes

Let  $\mathcal{D}(S)$  be the set of distributions over a given set  $S$ .

**Definition 2.1** (Markov decision process (MDP)). An MDP  $\mathcal{M}$  is a tuple  $(S, s_0, A, T, AP, L)$  where

- $S$  is a finite set of states, with  $s_0 \in S$  as the initial state;
- $A$  is a finite set of actions;
- $T: S \times A \rightarrow \mathcal{D}(S)$  is the probabilistic transition function;
- $AP$  is the set of *atomic propositions*; and
- $L: S \rightarrow 2^{AP}$  is the *labeling function*.

For any state  $s \in S$ , we let  $A(s)$  denote the set of actions that can be selected in state  $s$ . An MDP is a Markov chain if  $A(s)$  is a singleton for all  $s \in S$ .

*Runs.* For states  $s, s' \in S$  and  $a \in A(s)$ ,  $T(s, a)(s')$  equals the conditional probability  $p(s' | s, a)$ , which is the probability of jumping to state  $s'$  from state  $s$  under action  $a$ . A *run* of  $\mathcal{M}$  is an  $\omega$ -word  $\langle s_0, a_1, s_1, \dots \rangle \in S \times (A \times S)^\omega$  such that  $p(s_{i+1} | s_i, a_{i+1}) > 0$  for all  $i \geq 0$ . A *finite run* is a finite such sequence. For a run  $r = \langle s_0, a_1, s_1, \dots \rangle$  we define the corresponding *labeled run* as  $L(r) = \langle L(s_0), L(s_1), \dots \rangle \in (2^{AP})^\omega$ . We write  $Runs^{\mathcal{M}}(FRuns^{\mathcal{M}})$  for the set of (finite) runs of the MDP  $\mathcal{M}$  and  $Runs^{\mathcal{M}}(s)(FRuns^{\mathcal{M}}(s))$  for the set of (finite) runs of the MDP  $\mathcal{M}$  from  $s$ . We write  $last(r)$  for the last state of a finite run  $r$ . The superscript  $\mathcal{M}$  will be dropped when clear from context.

*Strategies or Policies.* A *strategy* in  $\mathcal{M}$  is a function  $\sigma: FRuns \rightarrow \mathcal{D}(A)$  such that  $\text{supp}(\sigma(r)) \subseteq A(last(r))$ , where  $\text{supp}(d)$  denotes the support of the distribution  $d$ . A *memory skeleton* is a tuple  $M_s = (M, \Sigma, m_0, \alpha_s)$  where  $M$  is a finite set of memory states,  $\Sigma$  is a finite alphabet,  $m_0 \in M$  is the initial state, and  $\alpha_s: M \times \Sigma \rightarrow M$  is the



memory update function. We define the extended memory update function  $\hat{\alpha}_s : M \times \Sigma^* \rightarrow M$  inductively in the usual way. A finite-memory strategy for  $\mathcal{M}$  over a memory skeleton  $M_s$  is a Mealy machine  $(M_s, \alpha_a)$  where  $\alpha_a : S \times M \rightarrow \mathcal{D}(A)$  is the *next-action function* that suggests the next action based on the MDP and memory state. The semantics of a finite-memory strategy  $(M, \alpha_a)$  is given as a strategy  $\sigma : \text{FRuns} \rightarrow \mathcal{D}(A)$  such that, for every  $r \in \text{FRuns}$ , we have that  $\sigma(r) = \alpha_a(\text{last}(r), \hat{\alpha}_s(m_0, L(r)))$ .

A strategy  $\sigma$  is *pure* if  $\sigma(r)$  is a point distribution for all runs  $r \in \text{FRuns}^M$  and a strategy is *mixed* (short for strictly mixed) if  $\text{supp}(\sigma(r)) = A(\text{last}(r))$  for all runs  $r \in \text{FRuns}^M$ . Let  $\text{Runs}_\sigma^M(s)$  denote the subset of runs  $\text{Runs}^M(s)$  from initial state  $s$  that follow strategy  $\sigma$ . Let  $\Pi_M$  be the set of all strategies. We say that  $\sigma$  is *stationary* if  $\text{last}(r) = \text{last}(r')$  implies  $\sigma(r) = \sigma(r')$  for all runs  $r, r' \in \text{FRuns}^M$ . A stationary strategy can be given as a function  $\sigma : S \rightarrow \mathcal{D}(A)$ . A strategy is *positional* if it is both pure and stationary.

*Probability Space.* An MDP  $\mathcal{M}$  under strategy  $\sigma$  results in a Markov chain  $\mathcal{M}_\sigma$ . If  $\sigma$  is a finite-memory strategy, then  $\mathcal{M}_\sigma$  is a finite-state Markov chain. The behavior of an MDP  $\mathcal{M}$  under a strategy  $\sigma$  and starting state  $s \in S$  is defined on the probability space  $(\text{Runs}_\sigma^M(s), \mathcal{F}_{\text{Runs}_\sigma^M(s)}, \text{Pr}_\sigma^M(s))$  over the set of infinite runs of  $\sigma$  with starting state  $s$ .  $\mathcal{F}_{\text{Runs}_\sigma^M(s)}$  denotes the sigma-algebra on these runs generated by the underlying MDP, and  $\text{Pr}_\sigma^M(s)$  is the probability distribution over these runs constructed inductively. Given a random variable  $f : \text{Runs}^M \rightarrow \mathbb{R}$ , we denote by  $\mathbb{E}_\sigma^M(s)\{f\}$  the expectation of  $f$  over the runs of  $\mathcal{M}$  starting from  $s$  that follow  $\sigma$ .

*Structural Properties of MDPs.* A *sub-MDP* of  $\mathcal{M}$  is an MDP  $\mathcal{M}' = (S', A', T', AP, L')$ , where  $S' \subset S$ ,  $A' \subseteq A$  is such that  $A'(s) \subseteq A(s)$  for all  $s \in S'$ , and  $T'$  and  $L'$  are  $T$  and  $L$  restricted to  $S'$  and  $A'$ . An *end-component* (De Alfaro 1998) of an MDP  $\mathcal{M}$  is a sub-MDP  $\mathcal{M}'$  that is closed under the transitions in  $T'$  and such that for every state pair  $s, s' \in S'$  there is a strategy that can reach  $s'$  from  $s$  with positive probability. A maximal end-component is an end-component that is maximal under set-inclusion. Every state  $s$  of an MDP  $\mathcal{M}$  belongs to at most one maximal end-component. A *bottom strongly connected component* (BSCC) of a Markov chain is any of its end-components.

An MDP  $\mathcal{M}$  is *communicating* if it is equal to its (only) maximal end-component. An MDP  $\mathcal{M}$  is *weakly communicating* if its state space can be decomposed into two sets: in the first set, each state is reachable from every other state in the set under some strategy; in the second set, all states are transient under all strategies, meaning that the probability of starting from  $s$  in this set and returning to  $s$  is less than one under any strategy.

## 2.2 Quantitative Objectives: Probabilistic Reward Machines

In the classical RL literature, the learning objective is specified using *Markovian* reward functions, i.e., a function  $\rho : S \times A \times S \rightarrow \mathbb{R}$  assigning utility to state-action pairs. A *rewardful* MDP is a tuple  $\mathcal{M} = (S, s_0, A, T, \rho)$ , where  $S, s_0, A$ , and  $T$  are defined as for MDPs, and  $\rho$  is a Markovian reward function. A rewardful MDP  $\mathcal{M}$  under a strategy  $\sigma$  determines a sequence of random rewards  $\rho(X_{i-1}, Y_i, X_i)_{i \geq 1}$ , where  $X_i$  and  $Y_i$  are the random variables denoting the  $i$ -th state and action, respectively. For  $\lambda \in [0, 1)$ , the *discounted reward* under  $\sigma$  is defined as:

$$\text{Disc}(\lambda)_\sigma^M(s) := \lim_{N \rightarrow \infty} \mathbb{E}_\sigma^M(s) \left\{ \sum_{i=1}^N \lambda^{i-1} \rho(X_{i-1}, Y_i, X_i) \right\},$$

while the *average reward* is defined as

$$\text{Avg}_\sigma^M(s) := \liminf_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_\sigma^M(s) \left\{ \sum_{i=1}^N \rho(X_{i-1}, Y_i, X_i) \right\}.$$

The average reward is defined using  $\liminf$  of the time-averaged expected reward, which is standard in average-reward maximization (Puterman 1994). For a finite MDP with bounded rewards, optimal stationary policies exist for which the average reward converges. Hence, for optimal policies,  $\liminf$ ,  $\limsup$ , and the limit coincide and induce the same optimal value.

For an objective  $Reward^M \in \{Disct(\lambda)^M, Avg^M\}$  and state  $s$ , we define the optimal reward  $Reward_*^M(s)$  as  $\sup_{\sigma \in \Pi_M} Reward_\sigma^M(s)$ . A strategy  $\sigma$  is optimal for  $Reward^M$  if  $Reward_\sigma^M(s) = Reward_*^M(s)$  for all  $s \in S$ . Optimal reward and strategies for these objectives can be computed in polynomial time (Puterman 1994).

Often, complex learning objectives cannot be expressed using Markovian reward functions. A recent trend is to express learning objectives using finite-state reward machines (see, e.g., (Icarte, T. Klassen, et al. 2018; Icarte, T. Q. Klassen, et al. 2022)), which provides a structured formalism to decompose high-level tasks and exploits the internal structure of the reward function. These reward machines often serve as an intermediate representations for formal languages, enabling the automatic translation of logical specifications into reward structures (Camacho, Icarte, et al. 2019). Recent works also include learning reward machines directly from agent interactions to handle non-Markovian rewards (Gaon and Brafman 2020) and partial observability (Toro Icarte et al. 2023), with extensions to multi-agent stochastic games (Hu et al. 2024) and inverse reinforcement learning (Shehab et al. 2025).

For the objectives we consider, we require a more expressive variant of reward machines that allows probabilistic and nondeterministic transitions, as well as spurious  $\epsilon$ -labeled transitions. We refer to these as probabilistic reward machines, with the understanding that any of the three transition types may be absent. For instance, the definition reduces to that of a nondeterministic reward machine as in (Kazemi, Perez, et al. 2022) when all enabled transitions are taken with probability one. The use of  $\epsilon$ -transitions serves only to streamline the presentation of our technical results.

**Definition 2.2** (Probabilistic Reward Machines). A *probabilistic reward machine* is a tuple  $\mathcal{R} = (\Sigma_\epsilon, U \times U^p, (u_0, u_0^p), \delta_r, T^p, \rho)$  where

- $\Sigma_\epsilon = (\Sigma \cup \{\epsilon\})$  with  $\Sigma$  being a finite alphabet and  $\epsilon$  indicating a silent transition,
- $U$  and  $U^p$  are two finite sets of states,
- $(u_0, u_0^p) \in U \times U^p$  is the starting state,
- $\delta_r : U \times \Sigma_\epsilon \rightarrow 2^U$  is the transition relation (that allows nondeterminism),
- $T^p : U^p \rightarrow \mathcal{D}(U^p)$  is the probabilistic transition function, and
- $\rho : U \times U^p \times \Sigma_\epsilon \times U \times U^p \rightarrow \mathbb{R}$  is the reward function.

To simplify the notation we let  $U^{\mathcal{R}} = U \times U^p$  and  $u_0^{\mathcal{R}} = (u_0, u_0^p)$ .

**Definition 2.3** (Product (rewardful) MDP). Given an MDP  $\mathcal{M} = (S, s_0, A, T, AP, L)$ , a reward machine  $\mathcal{R} = (\Sigma_\epsilon, U^{\mathcal{R}}, u_0^{\mathcal{R}}, \delta_r, T^p, \rho)$  with the alphabet  $\Sigma = 2^{AP}$ , and the labeling function  $L : S \rightarrow \Sigma$ , their product

$$\mathcal{M} \times \mathcal{R} = (S \times U^{\mathcal{R}}, s_0 \times u_0^{\mathcal{R}}, (A \times U^{\mathcal{R}}) \cup \{\epsilon\}, T^\times, \rho^\times)$$

is a rewardful MDP where  $T^\times : (S \times U^{\mathcal{R}}) \times ((A \times U^{\mathcal{R}}) \cup \{\epsilon\}) \rightarrow \mathcal{D}(S \times U^{\mathcal{R}})$  is such that

$$T^\times((s, (u, u^p)), \alpha)((s', (u', u^{p'}))) = \begin{cases} T(s, a)(s') \cdot T^p(u^p)(u^{p'}) & \text{if } \alpha = (a, (u', u^{p'})) \text{ and } (u, L(s), u') \in \delta_r \\ 1 & \text{if } \alpha = \epsilon, s = s', \text{ and } \delta(u, \epsilon, u') \in \delta_r \\ 0 & \text{otherwise.} \end{cases}$$

and  $\rho^\times : (S \times U^{\mathcal{R}}) \times ((A \times U^{\mathcal{R}}) \cup \{\epsilon\}) \times (S \times U^{\mathcal{R}}) \rightarrow \mathbb{R}$  is defined such that

$$\rho^\times((s, u^{\mathcal{R}}), \alpha, (s', u^{\mathcal{R}'})) = \begin{cases} \rho(u^{\mathcal{R}}, L(s), u^{\mathcal{R}'}) & \text{if } \alpha = (a, u^{\mathcal{R}'}) \\ \rho(u^{\mathcal{R}}, \epsilon, u^{\mathcal{R}'}) & \text{if } \alpha = \epsilon. \end{cases}$$

The Product MDP effectively runs the MDP and the reward machine in parallel. When an action  $a \in A$  is taken, a transition occurs in the MDP to update its state, and an appropriate transition occurs in the reward machine to update its state. When an  $\epsilon$ -action is taken, the state of the MDP remains the same, and one of the  $\epsilon$ -transitions of the reward machine is activated to update its state.

A deterministic reward machine is retrieved from the above definition by setting  $|\delta_r(u, a)| \leq 1$  and  $T^P(u^P) \in \{0, 1\}$  for all  $u \in U$ ,  $a \in \Sigma_e$ , and  $u^P \in U^P$ . For technical convenience, we assume that  $\mathcal{M} \times \mathcal{R}$  contains only states reachable from  $(s_0, u_0^{\mathcal{R}})$ . For both discounted and average objectives, the optimal strategies of  $\mathcal{M} \times \mathcal{R}$  are positional on  $\mathcal{M} \times \mathcal{R}$ . Moreover, these positional strategies characterize a finite memory strategy (with memory skeleton based on the states of  $\mathcal{R}$  and the next-action function based on the positional strategy) over  $\mathcal{M}$  maximizing the learning objective given by  $\mathcal{R}$ .

### 2.3 Qualitative Objectives: Omega-Regular Specifications

Formal specification languages, such as Büchi automata and logic, provide a rigorous and unambiguous mechanism to express learning objectives for continuing tasks. There is a growing trend (Bozkurt, Y. Wang, Zavlanos, et al. 2020; Hahn, Perez, et al. 2019; H. Hasanbeig et al. 2023; Kazemi and Soudjani 2020; Sadigh et al. 2014) toward expressing learning objectives in RL using linear temporal logic (LTL) and  $\omega$ -regular languages (which strictly generalize LTL). We will describe  $\omega$ -regular languages by *good-for-MDP* Büchi automata (Hahn, Perez, et al. 2020).

Linear Temporal Logic (LTL) (C. Baier and Katoen 2008) is a temporal logic that is often used to specify objectives in human-readable form. Given a set of atomic propositions  $AP$ , the LTL formulae over  $AP$  can be defined via the following grammar:

$$\varphi := a \mid \neg\varphi \mid \varphi \vee \psi \mid X\varphi \mid \varphi U \psi, \quad (2.1)$$

where  $a \in AP$ , using negation  $\neg$ , disjunction  $\vee$ , next  $X$ , and until  $U$  operators. Additional operators are introduced as abbreviations:  $\top \stackrel{\text{def}}{=} a \vee \neg a$ ;  $\perp \stackrel{\text{def}}{=} \neg\top$ ;  $\varphi \wedge \psi \stackrel{\text{def}}{=} \neg(\neg\varphi \vee \neg\psi)$ ;  $\varphi \rightarrow \psi \stackrel{\text{def}}{=} \neg\varphi \vee \psi$ ;  $F\varphi \stackrel{\text{def}}{=} \top U \varphi$ ; and  $G\varphi \stackrel{\text{def}}{=} \neg F\neg\varphi$ . We write  $w \models \varphi$  if  $\omega$ -word  $w$  over  $2^{AP}$  satisfies LTL formula  $\varphi$ . The satisfaction relation is defined inductively (C. Baier and Katoen 2008). We will provide details of classes of specifications including safety, liveness, and fairness in Subsection 3.

Nondeterministic Büchi automata are finite state acceptors for all  $\omega$ -regular languages.

**Definition 2.4** (Büchi Automata). A (nondeterministic) *Büchi automaton* is a tuple  $\mathcal{A} = (\Sigma, Q, q_0, \delta, F)$ , where

- $\Sigma$  is a finite *alphabet*,
- $Q$  is a finite set of *states*,
- $q_0 \in Q$  is the *initial state*,
- $\delta: Q \times \Sigma \rightarrow 2^Q$  is the *transition function*, and
- $F \subset Q \times \Sigma \times Q$  is the set of *accepting transitions*.

A *run*  $r$  of  $\mathcal{A}$  on  $w \in \Sigma^\omega$  is an  $\omega$ -word  $\langle r_0, w_0, r_1, w_1, \dots \rangle \in (Q \times \Sigma)^\omega$  such that  $r_0 = q_0$  and, for  $i > 0$ ,  $r_i \in \delta(r_{i-1}, w_{i-1})$ . Each triple  $(r_{i-1}, w_{i-1}, r_i)$  is a *transition* of  $\mathcal{A}$ . We write  $\text{Inf}(r)$  for the set of transitions that appear infinitely often in the run  $r$ . A run  $r$  of  $\mathcal{A}$  is *accepting* if  $\text{Inf}(r) \cap F \neq \emptyset$ . The *language*  $\mathcal{L}(\mathcal{A})$  of  $\mathcal{A}$  is the subset of words in  $\Sigma^\omega$  with accepting runs in  $\mathcal{A}$ . A language is  *$\omega$ -regular* if it is accepted by a Büchi automaton.

*Good-for-MDP Büchi Automata.* Given an MDP  $\mathcal{M}$  and a specification  $\varphi$  represented with a Büchi automaton  $\mathcal{A} = (\Sigma, Q, q_0, \delta, F)$ , we want to compute an optimal strategy satisfying the objective. We define the satisfaction probability of  $\sigma$  from starting state  $s$  as

$$\text{PSem}_{\mathcal{A}}^{\mathcal{M}}(s, \sigma) = \Pr_{\sigma}^{\mathcal{M}}(s) \{r \in \text{Runs}_{\sigma}^{\mathcal{M}}(s) : L(r) \in \mathcal{L}(\mathcal{A})\}.$$

The optimal satisfaction probability  $\text{PSem}_{\mathcal{A}}^{\mathcal{M}}(s)$  for specification  $\mathcal{A}$  is defined as  $\sup_{\sigma \in \Pi_{\mathcal{M}}} \Pr_{\sigma}^{\mathcal{M}}(s, \sigma)$  and we say that  $\sigma \in \Pi_{\mathcal{M}}$  is an optimal strategy for  $\mathcal{A}$  if  $\text{PSem}_{\mathcal{A}}^{\mathcal{M}}(s, \sigma)(s) = \text{PSem}_{\mathcal{A}}^{\mathcal{M}}(s)$ .

**Definition 2.5** (Product MDP). Given an MDP  $\mathcal{M} = (S, s_0, A, T, AP, L)$  and automaton  $\mathcal{A} = (\Sigma, Q, q_0, \delta, F)$  with alphabet  $\Sigma = 2^{AP}$ , the *product*  $\mathcal{M} \times \mathcal{A} = (S \times Q, (s_0, q_0), A \times Q, T^{\times}, F^{\times})$  is an MDP with initial state  $(s_0, q_0)$  and



accepting transitions  $F^\times$  where  $T^\times : (S \times Q) \times (A \times Q) \rightarrow \mathcal{D}(S \times Q)$  is defined by

$$T^\times((s, q), (a, q'))((s', q')) = \begin{cases} T(s, a)(s') & \text{if } (q, L(s), q') \in \delta \\ 0 & \text{otherwise.} \end{cases}$$

The accepting transition set  $F^\times \subseteq (S \times Q) \times (A \times Q) \times (S \times Q)$  is defined by  $((s, q), (a, q'), (s', q')) \in F^\times$  if, and only if,  $(q, L(s), q') \in F$  and  $T(s, a)(s') > 0$ . A strategy  $\sigma^\times$  on the product defines a strategy  $\sigma$  on the MDP with the same value, and vice versa. Note that for a stationary  $\sigma^\times$ , the strategy  $\sigma$  may need memory. End-components and runs of the product MDP are defined just like for MDPs.

A run  $r$  of  $\mathcal{M} \times \mathcal{A}$  is accepting if  $\text{Inf}(r) \cap F^\times \neq \emptyset$ . The *syntactic satisfaction* probability is the probability of accepting runs:

$$\text{PSat}_{\mathcal{A}}^{\mathcal{M}}((s, q), \sigma^\times) = \Pr_{\sigma^\times}^{\mathcal{M} \times \mathcal{A}}(s, q) \{r \in \text{Runs}_{\sigma^\times}^{\mathcal{M} \times \mathcal{A}}(s, q) : \text{Inf}(r) \cap F^\times \neq \emptyset\}.$$

Similarly, we define  $\text{PSat}_{\mathcal{A}}^{\mathcal{M}}(s)$  as the optimal probability over the product, i.e.,  $\sup_{\sigma^\times} (\text{PSat}_{\mathcal{A}}^{\mathcal{M}}((s, q_0), \sigma^\times))$ . For a deterministic  $\mathcal{A}$  the equality  $\text{PSat}_{\mathcal{A}}^{\mathcal{M}}(s) = \text{PSem}_{\mathcal{A}}^{\mathcal{M}}(s)$  holds; however, this is not guaranteed for nondeterministic Büchi automata as the optimal resolution of nondeterministic choices may require access to future events. This motivates the definition of a good-for-MDP nondeterministic Büchi automata.

**Definition 2.6** (Good-for-MDP (Hahn, Perez, et al. 2020)). A Büchi automaton  $\mathcal{A}$  is *good for MDPs* (GFM), if  $\text{PSat}_{\mathcal{A}}^{\mathcal{M}}(s_0) = \text{PSem}_{\mathcal{A}}^{\mathcal{M}}(s_0)$  holds for all MDPs  $\mathcal{M}$  and starting states  $s_0$ .

Note that every  $\omega$ -regular objective can be expressed as a GFM automaton (Hahn, Perez, et al. 2020). A popular class of GFM automata is suitable limit-deterministic Büchi automata (Hahn, Li, et al. 2015; Sickert et al. 2016). There are other types of automata that are good-for-MDPs. For example, in parity automata, each transition is assigned an integer priority. A run of a parity automaton is accepting if the highest recurring priority is odd. In this paper, we only use GFM Büchi automata.

The satisfaction of an  $\omega$ -regular objective given as a GFM automaton  $\mathcal{A}$  by an MDP  $\mathcal{M}$  can be formulated in terms of the accepting maximal end-components of the product  $\mathcal{M} \times \mathcal{A}$ , i.e., the maximal end-components that contain an accepting transition from  $F^\times$ . The optimal satisfaction probabilities and strategies can be computed by first computing the accepting maximal end-components of  $\mathcal{M} \times \mathcal{A}$  and then maximizing the probability to reach states in such components. The optimal strategies are positional on  $\mathcal{M} \times \mathcal{A}$  and induce finite-memory strategies over  $\mathcal{M}$  that maximize the satisfaction probability of the learning objective given by  $\mathcal{A}$ .

## 2.4 Reinforcement Learning for Continuing Tasks and Problem Statements

RL is a sampling-based optimization approach where an agent learns to optimize its strategy by repeatedly interacting with the environment relying on the reinforcements (numerical reward signals) it receives for its actions. We focus on the model-free approach to RL, where the learner computes optimal strategies without explicitly estimating the transition probabilities and rewards. These approaches are asymptotically more space-efficient (Strehl et al. 2006) than model-based RL and have been shown to scale well (Mnih et al. 2015; Silver, Huang, et al. 2016). Some prominent model-free RL algorithms for discounted and average reward objectives include Q-learning and TD( $\lambda$ ) (Sutton and Barto 2018) and Differential Q-learning (Wan, Naik, et al. 2021).

In some applications, such as running a maze or playing tic-tac-toe, the interaction between the agent and the environment naturally breaks into finite length learning sequences, called episodes. Thus the agent optimizes its strategy by combining its experience over different episodes. We call such tasks *episodic*. On the other hand, for some applications—such as process control and reactive systems—this interaction continues ad-infinitum and the agent learns over a single lifetime. We call such tasks *continuing*. This paper develops a model-free RL algorithm for continuing tasks where the learning objective is an  $\omega$ -regular specification given as a GFM automaton. Prior

solutions (Bozkurt, Y. Wang, Zavlanos, et al. 2020; Hahn, Perez, et al. 2019; H. Hasanbeig et al. 2023; Sadigh et al. 2014) focused on episodic setting and have proposed a model-free reduction from  $\omega$ -regular objectives to discounted-reward objectives. Several researchers (Naik et al. 2019; Sutton and Barto 2018) have made the case for adopting average reward formulation for continuing tasks due to several limitations of discounted-reward RL in continuing tasks. This paper investigates a model-free reduction from  $\omega$ -regular objectives to average-reward objectives in model-free RL.

**Problem 1** ( $\omega$ -Regular to Average Reward Translation). Given an unknown communicating MDP  $\mathcal{M} = (S, s_0, A, T, AP, L)$  and a GFM automaton  $\mathcal{A} = (\Sigma, Q, q_0, \delta, F)$ , the *reward translation* problem is to design a reward machine  $\mathcal{R}$  such that an optimal positional strategy maximizing the average reward for  $\mathcal{M} \times \mathcal{R}$  provides a finite memory strategy maximizing the satisfaction probability of  $\mathcal{L}(\mathcal{A})$  in  $\mathcal{M}$ .

We also consider the lexicographic multi-objective optimization in which the goal is to maximize a mean-payoff objective over those policies that satisfy a given liveness specification.

**Problem 2** (Lexicographic  $\omega$ -regular and Average Reward). Given an unknown communicating MDP  $\mathcal{M} = (S, s_0, A, T, AP, L)$ , a GFM automaton  $\mathcal{A} = (\Sigma, Q, q_0, \delta, F)$ , and a reward function  $\rho(s, s')$ , the problem is to design an algorithm such that the resulting policy maximizes the average reward

$$\text{Avg}_\sigma^{\mathcal{M}}(s) := \liminf_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_\sigma^{\mathcal{M}}(s) \left\{ \sum_{i=1}^N \rho(s_{i-1}, a_i, s_i) \right\},$$

among the policies that maximize the satisfaction probability of  $\mathcal{A}$  in  $\mathcal{M}$ .

Next we give an account of challenges with average-reward RL, identify suitable classes of qualitative temporal specifications in Section 3, and then solve Problems 1-2 respectively in Sections 4-5 under appropriate assumptions on the structure of the underlying MDP.

## 2.5 Average-Reward Reinforcement Learning

Given an MDP  $\mathcal{M}$ , reward machine  $\mathcal{R}$ , and an objective (discounted or average reward), an optimal strategy can be computed in polynomial time using linear programming (Puterman 1994). Similarly, graph-theoretic techniques to find maximal end-components can be combined with linear programming to compute optimal strategies for  $\omega$ -regular objectives (Hahn, Li, et al. 2015). However, such techniques are not applicable when the transitions or reward structure of the MDP is unknown. The existing average-reward RL algorithms such as Differential Q-learning provide convergence guarantees under the assumption that the MDP  $\mathcal{M}$  is communicating (Wan, Naik, et al. 2021). Thus, for the reward translation in Problem 1 to be effective, we need the product  $\mathcal{M} \times \mathcal{A}$  to be communicating. Unfortunately, even when  $\mathcal{M}$  is communicating,  $\mathcal{M} \times \mathcal{A}$  may violate the communicating requirement. We tackle this issue in Section 4.

The communicating assumption requires that, for every pair of states  $s$  and  $s'$ , there is a policy under which the MDP can reach  $s'$  from  $s$  in a finite number of steps with positive probability. Under the communicating assumption, there exists a unique optimal reward rate  $r^*$  that is independent of the start state. The differential Q-learning algorithm constructs a table of estimates  $Q_t : S \times U \rightarrow \mathbb{R}$  at each time step  $t$ . Let us denote the state-action pair of the MDP at time  $t$  by  $(s_t, u_t)$ . Then, differential Q-learning updates the table as follows:

$$\begin{aligned} Q_{t+1}(s, u) &:= Q_t(s, u) && \text{for all } (s, u) \neq (s_t, u_t) \\ Q_{t+1}(s, u) &:= Q_t(s_t, u_t) + \alpha_t \delta_t, && \text{for } (s, u) = (s_t, u_t), \end{aligned}$$

where  $\alpha_t$  is the step-size at step  $t$ , and  $\delta_t$  is the temporal difference error defined as

$$\delta_t := \rho_{t+1} - \bar{r}_t + \max_u Q_t(s_{t+1}, u) - Q_t(s_t, u_t) \text{ and } \bar{r}_{t+1} := \bar{r}_t + \eta \cdot \alpha_t \cdot \delta_t,$$

where  $\rho_{t+1}$  is the reward received for the transition  $(s_t, u_t, s_{t+1})$ ,  $\bar{r}_t$  is a scalar estimate of the optimal reward rate  $r^*$ , and  $\eta$  is a positive constant. It is shown in (Wan, Naik, et al. 2021) that  $\bar{r}_t$  converges to  $r^*$  under the following assumptions:

- (a) the MDP is communicating;
- (b) the associated Bellman equation has a unique solution up to a constant;
- (c) the step sizes  $\alpha_t$  decrease appropriately with  $t$ ;
- (d) all state–action pairs are updated infinitely often; and
- (e) the ratio of the update frequency of the most-updated state–action pair to the least-updated one is finite.

### 3 Specifications Suitable for Continual Learning

In this section, we study and characterize temporal logic specifications that express various classes of tasks. We then identify those suitable for continual reinforcement learning and provide solutions to Problems 1–2 in subsequent sections. Two important classes of temporal logic specifications are *safety* and *liveness* properties (Alpern and Schneider 1985; Manna and Pnueli 1990). Loosely speaking, safety specifications ensure that *something bad never happens*, while liveness specifications guarantee that *something good eventually happens*. Below, we briefly recall the formal definition of safety and then focus on liveness.

**Definition 3.1** (Safety and Liveness). Let  $w$  be a sequence over alphabet  $\Sigma$  and let  $w(i)$  represent the prefix of  $w$  of length  $i + 1$ ; i.e.,  $x(i) = (x_0, x_1, \dots, x_i)$ .

- A language  $\mathcal{L} \subseteq \Sigma^\omega$  is a safety specification if, and only if, for all  $\omega$ -sequences  $w \in \Sigma^\omega$  not in  $\mathcal{L}$ , some prefix of  $w$  cannot be extended to an  $\omega$ -sequence in the language.
- A language  $\mathcal{L} \subseteq \Sigma^\omega$  is a liveness specification if, and only if, we can extend any finite sequence in  $\Sigma^*$  to an  $\omega$ -sequence in the language. In other words, the set  $\{w(i) : w \in \mathcal{L}\}$  is  $\Sigma^*$ . For instance  $\psi_1 = a \vee F b$  is a liveness specification.

**Definition 3.2** (Absolute Liveness). A language  $\mathcal{L} \subseteq \Sigma^\omega$  is an *absolute liveness* specification if, and only if,  $\mathcal{L}$  is non-empty and, for every finite sequence  $\bar{w} \in \Sigma^*$  and every  $w \in \mathcal{L}$ ,  $\bar{w}w \in \mathcal{L}$  holds. That is, appending an arbitrary finite prefix to an accepted word produces an accepted word. This implies that an LTL specification  $\varphi$  is absolute liveness specification if  $\varphi$  is satisfiable and  $\varphi$  is equivalent to  $F\varphi$ . One can observe that  $\psi_1$  is not an absolute liveness specification since adding the prefix  $\neg a$  to a trace that does not satisfy  $F b$  yields a trace not in the language. However,  $\psi_2 = F a$  is an absolute liveness specification.

**Definition 3.3** (Stable Specification). A language  $\mathcal{L} \subseteq \Sigma^\omega$  is a *stable* specification if, and only if, it is closed under suffixes; that is, if  $x$  is in  $\mathcal{L}$ , then every suffix of  $x$  is also in  $\mathcal{L}$ . Since deleting the first element from  $a(\neg a)^\omega$  results in a sequence that does not satisfy  $\psi_2$ ,  $\psi_2$  is not stable. However,  $\psi_3 = G a \vee G F b$  is stable. Moreover, a language is a *fairness* specification if and only if it is both a stable specification and absolute liveness specification, e.g.,  $\psi_4 = G F a$ . One can easily conclude that none of the specifications  $\psi_1$ ,  $\psi_2$ , and  $\psi_3$  are fairness specifications.

We give a solution for the translation in Problem 1 for *absolute liveness* specifications (Sistla 1994). Adding a prefix to a trace for average reward objectives should not change the average value associated with the trace. This is aligned with the satisfaction of absolute liveness specifications. Moreover, since absolute liveness specifications cannot be rejected for any finite word, they preserve the continual nature of the learning procedure. In the following, we focus on the characteristics of absolute liveness specifications.

The set of absolute liveness specifications is closed under union and intersection. The proofs mentioned below rely on the property  $\varphi \equiv F\varphi$  in LTL notation. The more general  $\omega$ -regular proofs use  $\varphi \equiv \Sigma^*\varphi$ . In addition, closure under intersection requires an “asterisk:” the intersection is an absolute liveness specification unless it is empty. Suppose  $\varphi_1$  is equivalent to  $F\varphi_1$  and  $\varphi_2$  is equivalent to  $F\varphi_2$ . Then,

$$\varphi_1 \vee \varphi_2 \equiv (F\varphi_1) \vee (F\varphi_2) \equiv F(\varphi_1 \vee \varphi_2) .$$

Moreover, since for every LTL specification  $\varphi$ ,  $F\varphi$  is equivalent to  $F F\varphi$ ,

$$\varphi_1 \wedge \varphi_2 \equiv (F F\varphi_1) \wedge (F F\varphi_2) \equiv F((F\varphi_1) \wedge (F\varphi_2)) \equiv F(\varphi_1 \wedge \varphi_2) .$$

Lemma 2.1 of (Sistla 1994), which says that the complement of a stable specification different from  $\Sigma^\omega$  is an absolute liveness specification and vice versa, can be used to show that the stable specifications are also closed under union and intersection. The case when either stable specification is  $\top$  is easily proved. Because of these closure properties,  $\bigvee_{1 \leq k \leq n} (F G a_k) \wedge (G F b_k)$  (where each  $a_k$  and  $b_k$  is an atomic proposition) is an absolute liveness specification of Rabin index  $n$ . This means that there is no upper bound to the number of priorities required by a parity automaton accepting an absolute liveness specification (Wagner 1979). The simplest absolute liveness specification is  $\top$ , which requires one priority. The automata that accept absolute liveness and stable specifications have special properties discussed next.

**Lemma 3.4.** *Let  $\mathcal{A}$  be a deterministic Büchi automaton with initial state  $s_0$ . The automaton  $\mathcal{A}$  accepts an absolute liveness specification if and only if, for every reachable state  $s$  of  $\mathcal{A}$ , the language accepted from  $s$  contains the language accepted from  $s_0$ .*

PROOF. Let  $u$  be a word that takes the automaton  $\mathcal{A}$  from the initial state  $s_0$  to a reachable state  $s$ , and let  $w$  be a word accepted from  $s_0$ . Suppose first that  $\mathcal{A}$  accepts an absolute liveness specification. Since  $w$  is accepted from  $s_0$ , the concatenated word  $uw$  is also accepted from  $s_0$ . Because  $\mathcal{A}$  is deterministic, the unique run of  $uw$  starting from  $s_0$  reaches  $s$  after reading  $u$  and then continues with  $w$ . It follows that  $w$  is accepted from  $s$ . Conversely, suppose that for every reachable state  $s$  of  $\mathcal{A}$ , the language accepted from  $s$  contains the language accepted from  $s_0$ . Let  $u$  be a word that takes  $\mathcal{A}$  from  $s_0$  to  $s$ , and let  $w$  be a word accepted from  $s_0$ . By assumption,  $w$  is accepted from  $s$ , which implies that the concatenated word  $uw$  is accepted from  $s_0$ . Therefore,  $\mathcal{A}$  accepts an absolute liveness specification.  $\square$

The “only-if” direction of Lemma 3.4 does not extend to nondeterministic automata. This failure is witnessed by the two-state automaton shown in Fig. 1(a), which recognizes the absolute liveness specification  $F G a$ . In this automaton, the unique accepting state recognizes precisely the language  $G a$ . An analogous separation result holds for stable specifications.

**Lemma 3.5.** *Let  $\mathcal{A}$  be a Büchi automaton with initial state  $s_0$ . The automaton  $\mathcal{A}$  accepts a stable specification if and only if, for every reachable state  $s$  of  $\mathcal{A}$ , the language accepted from  $s$  is contained in the language accepted from  $s_0$ .*

PROOF. Let  $u$  be a word that takes the automaton  $\mathcal{A}$  from the initial state  $s_0$  to a reachable state  $s$ , and let  $w$  be a word accepted from  $s$ . Suppose first that  $\mathcal{A}$  accepts a stable specification. Then the concatenated word  $uw$  is accepted from  $s_0$ . Since stable specifications are closed under taking suffixes, it follows that the suffix  $w$  of  $uw$  is also accepted from  $s_0$ . Conversely, suppose that for every reachable state  $s$  of  $\mathcal{A}$ , the language accepted from  $s$  is contained in the language accepted from  $s_0$ . Let  $u$  take  $\mathcal{A}$  from  $s_0$  to  $s$  along an accepting run of  $uw$ . Since  $w$  is accepted from  $s$ , the assumption implies that  $w$  is accepted from  $s_0$ . Hence,  $uw$  is accepted from  $s_0$ , and therefore  $\mathcal{A}$  accepts a stable specification.  $\square$

Note that Lemma 3.5 does not require determinism. Combined with Lemma 3.4, it proves that a deterministic automaton accepts a fairness specification if, and only if, all its states are language-equivalent. This implies that a fairness specification accepted by a deterministic (Büchi) automaton is accepted by a strongly connected deterministic (Büchi) automaton. Any reachable sink SCC of a deterministic automaton for a fairness specification is, in itself, an automaton for the specification. The above lemmas show that checking whether a specification is stable or an absolute liveness specification is reducible to checking language containment for deterministic Büchi automata. To solve Problem 1, we make the following assumption.



(a) A nondeterministic Büchi automaton for the absolute liveness specification  $FG a$ , used as a counterexample to the “only-if” direction of Lemma 3.4.

(b) Automaton for the absolute liveness specification  $(FG a) \vee (FG \neg a)$  used in Example 4.1.

Fig. 1. Nondeterministic Büchi automata illustrating absolute liveness specifications.

**Assumption 1.** Given an MDP  $\mathcal{M}$  and Büchi automaton  $\mathcal{A}$ , we assume that: 1)  $\mathcal{M}$  is communicating; 2)  $\mathcal{A}$  is a GFM automaton; and 3)  $\mathcal{A}$  accepts an absolute liveness specification.

We also study relaxing the communicating assumption on  $\mathcal{M}$  to weakly communicating, which requires strengthening the specification to fairness in order to retain correctness and convergence guarantees.

#### 4 Average-Reward RL for Qualitative Objectives

This section provides a solution for Problem 1. Let us fix a communicating MDP  $\mathcal{M} = (S, s_0, A, T, AP, L)$  and an absolute liveness GFM property  $\mathcal{A} = (\Sigma, Q, q_0, \delta, F)$  for the rest of this section. Our goal is to construct a reward machine  $\mathcal{R}$  such that we can use an off-the-shelf average reward RL on  $\mathcal{M} \times \mathcal{R}$  to compute an optimal strategy of  $\mathcal{M}$  against  $\mathcal{A}$ . Since the optimal strategies are not positional on  $\mathcal{M}$  but rather positional on  $\mathcal{M} \times \mathcal{A}$ , it is natural to assume that the reward machine  $\mathcal{R}$  takes the structure of  $\mathcal{A}$  with a reward function providing positive reinforcement with every accepting transition. Unfortunately, even for absolute liveness GFM automata  $\mathcal{A}$ , the product  $\mathcal{M} \times \mathcal{A}$  with a communicating MDP  $\mathcal{M}$  may not be communicating.

**Example 4.1.** Assume a communicating MDP  $\mathcal{M}$  with at least one state labeled  $a$  or  $\neg a$ , and the absolute liveness property  $\varphi = (FG a) \vee (FG \neg a)$  with its automaton shown in Fig. 1(b). Observe that any run that visits one of the two accepting states cannot visit the other one. Hence, the product does not satisfy the communicating property.

*Reward Machine Construction.* Let  $\mathcal{A} = (\Sigma, Q, q_0, \delta, F)$  be an absolute liveness GFM Büchi automaton. Consider  $\mathcal{R}_{\mathcal{A}} = (\Sigma_{\epsilon}, Q, q_0, \delta', \rho)$  where  $\delta'(q, a) = \delta(q, a)$  for all  $a \in \Sigma$  and  $\epsilon$  transitions reset to the starting state, i.e.  $\delta'(q, \epsilon) = q_0$ . Note that by adding the reset ( $\epsilon$ ) action from every state of  $\mathcal{R}$  to its initial state, the graph structure of  $\mathcal{M}$  is strongly connected. The reward function  $\rho : Q \times (\Sigma \cup \{\epsilon\}) \times Q \rightarrow \mathbb{R}$  is such that

$$\rho(q, a, q') = \begin{cases} c & \text{if } a = \epsilon \\ 1 & \text{if } (q, a, q') \in F \\ 0 & \text{otherwise.} \end{cases}$$

*Remark 4.2.* The above construction that adds resets transitions to the starting state could also be seen from the lens of constructing a GFM parity automaton. This would be performed as follows. First, transform the original GFM Büchi automaton into a nondeterministic parity automaton by assigning parity equal to two for accepting transitions and equal to one for other transitions. Then, add reset transitions with parity equal to three. Therefore, reset transitions are rejecting at the highest priority. This parity automaton preserves the language of the original Büchi automaton: since an accepting run involves only finitely many resets, any accepted word effectively consists of a finite prefix followed by a suffix accepted by the original automaton. Because absolute liveness properties are closed under prefix addition, the resulting word remains in the language. Furthermore,



adding transitions yields an automaton that simulates the original one. Consequently, by invoking (Hahn, Perez, et al. 2020, Lemma 1), we conclude that the resulting parity automaton remains GFM.

**Lemma 4.3** (Preservation of Communication). *For a communicating MDP  $\mathcal{M}$  and reward machine  $\mathcal{R}_{\mathcal{A}}$  for an absolute liveness GFM automaton  $\mathcal{A}$ , we have that the product  $\mathcal{M} \times \mathcal{R}_{\mathcal{A}}$  is communicating.*

PROOF. To show that  $\mathcal{M} \times \mathcal{R}_{\mathcal{A}}$  is communicating, we need to show that for arbitrary states  $(s, q), (s', q') \in S \times Q$  reachable from the initial state  $(s_0, q_0)$ , there is a strategy that can reach  $(s', q')$  from  $(s, q)$  with positive probability. Note that since  $\mathcal{M}$  is communicating, it is possible to reach  $(s_0, q'')$  from  $(s, q)$  for some  $q''$  of  $\mathcal{R}_{\mathcal{A}}$  using a strategy to reach  $s_0$  from  $s$  in  $\mathcal{M}$ . We can then use a reset ( $\epsilon$ ) action in  $\mathcal{R}_{\mathcal{A}}$  to reach the state  $(s_0, q_0)$ . Since  $(s', q')$  is reachable from the initial state  $(s_0, q_0)$ , we have a strategy to reach  $(s', q')$  from  $(s, q)$  with positive probability.  $\square$

**Lemma 4.4** (Average and Probability). *There exists a constant  $c^* < 0$  such that for all  $c < c^*$ , positional strategies that maximize the average reward on  $\mathcal{M} \times \mathcal{R}_{\mathcal{A}}$  will maximize the satisfaction probability of  $\mathcal{A}$ .*

PROOF. The proof is in three parts.

- (1) First observe that if  $c < 0$ , then for any average-reward optimal strategy in  $\mathcal{M} \times \mathcal{R}_{\mathcal{A}}$ , the expected average reward is non-negative. This is so because all other actions except  $\epsilon$  actions provide non-negative rewards. Hence, any strategy that takes  $\epsilon$  actions only finitely often, results in a non-negative average reward.
- (2) Let  $\Pi^*$  be the set of positional strategies in  $\mathcal{M} \times \mathcal{R}_{\mathcal{A}}$  such that the  $\epsilon$  actions are taken only finitely often, i.e. no BSCC of the corresponding Markov chain contains an  $\epsilon$  transition. Let  $\Pi^\epsilon$  be the set of remaining positional strategies, i.e., the set of positional strategies that visit an  $\epsilon$  transition infinitely often. Let  $0 < p_{\min} < 1$  be a lower bound on the expected long-run frequency of the  $\epsilon$  transitions among all strategies in  $\Pi^\epsilon$ . Let  $c^* = -1/p_{\min}$ . Observe that for every policy  $\sigma' \in \Pi^\epsilon$ , the expected average reward is negative and cannot be an optimal strategy in  $\mathcal{M} \times \mathcal{R}_{\mathcal{A}}$ . To see that, let  $0 < p \leq 1$  be the long-run frequency of the  $\epsilon$  transitions for  $\sigma$  and let  $0 \leq q < 1$  be the long-run frequency of visiting accepting transitions for  $\sigma$ . The average reward for  $\sigma$  is

$$\begin{aligned} \text{Avg}_\sigma^{\mathcal{M} \times \mathcal{R}_{\mathcal{A}}}(s_0, q_0) &= p \cdot c + q \cdot 1 + (1 - p - q) \cdot 0 \\ &\leq p \cdot c + q \cdot 1 + (1 - p - q) \cdot 1 = p \cdot c + (1 - p) \\ &\leq p \cdot c^* + (1 - p) = -p/p_{\min} + (1 - p) \leq -1 + (1 - p) \leq -p. \end{aligned}$$

Since every optimal policy must have a non-negative average reward, no policy in  $\Pi^\epsilon$  is optimal for  $c < c^*$ .

- (3) Now consider an optimal policy  $\sigma^*$  in  $\Pi^*$ . We show that this policy also optimizes the probability of satisfaction of  $\mathcal{A}$ . There are two cases to consider.
  - (a) If the expected average reward of  $\sigma^*$  is 0, then under no strategy it is possible to reach an accepting transitions (positive reward transition) in  $\mathcal{M} \times \mathcal{R}_{\mathcal{A}}$ . Hence, every policy is optimal in  $\mathcal{M}$  against  $\mathcal{A}$ , and so is  $\sigma^*$ .
  - (b) If the expected average reward of  $\sigma^*$  is positive, then notice that for every BSCC of the Markov chain of  $\mathcal{M} \times \mathcal{R}_{\mathcal{A}}$  under  $\sigma^*$ , the average reward is the same. This is so because otherwise, there is a positional policy that reaches the BSCC with the optimal average from all the other BSCCs with lower averages, contradicting the optimality of  $\sigma^*$ . Since for an optimal policy  $\sigma^*$ , every BSCC provides the same positive average, every BSCC must contain an accepting transition. Hence, every run of the MDP  $\mathcal{M}$  under  $\sigma^*$  will eventually dwell in an accepting component and in the process will see a finitely many  $\epsilon$  (reset) transitions. For any such given run  $r$ , consider the the suffix  $r'$  of the run after the last  $\epsilon$  transition is taken and let  $r = wr'$  for some finite run  $w$ . Since  $L(r')$  is an accepting word in  $\mathcal{A}$ , and since  $\mathcal{A}$  is an absolute liveness property any arbitrary prefix  $w'$  to this run  $r'$  is also accepting. This implies that the original run  $r$  is also accepting for  $\mathcal{A}$ . It follows that for such a strategy  $\sigma^*$ , the probability of satisfaction of  $\mathcal{A}$  is 1, making  $\sigma^*$  an optimal policy for  $\mathcal{M}$  against  $\mathcal{A}$ .  $\square$




Fig. 2. Automaton of FG  $a$ , dashed lines represent resets.




Fig. 3. MDP, each transition represents an action.

Since our translation from absolute liveness  $\omega$ -regular objective to reward machines is model-free, the following theorem is immediate.

**Theorem 4.5** (Convergence of Model-free RL). *Differential Q-learning algorithm for maximizing average reward objective on  $\mathcal{M} \times \mathcal{R}_{\mathcal{A}}$  will converge to a strategy maximizing the probability of satisfaction of  $\mathcal{A}$  for a suitable value of  $c$ . Moreover, the product construction  $\mathcal{M} \times \mathcal{R}_{\mathcal{A}}$  can be done on-the-fly and it is model-free.*

As an example, consider the property FG  $a$  and an MDP with two states and all transitions between states are available as deterministic actions (Fig. 2 and Fig. 3). Only one of the states is labeled  $a$ . An infinite memory strategy could see  $a$  for one step, reset, then see two  $a$ s, reset, then see three  $a$ s and so forth. This strategy will produce the same average value as the positional strategy which sees  $a$  forever without resetting. However, the infinite memory strategy will fail the property while the positional one will not.

*Shaping Rewards via Hard Resets.* For a Büchi automaton  $\mathcal{A}$ , we say that its state  $q \in Q$  is *coaccessible* if there exists a path starting from that state to an accepting transition. If a state is not coaccessible then any run of the product  $\mathcal{M} \times \mathcal{A}$  that ends in such a state will never be accepting, and hence one can safely redirect all of its outgoing (even better, incoming) transitions to the initial state with reward  $c$  (a hard reset). Such hard resets will promote speedy learning by reducing the time spent in such states during unsuccessful explorations, and at the same time adding these resets does not make a non-accepting run accepting or vice versa. Lemma 4.3, Lemma 4.4, and Theorem 4.5 continue to hold with such hard resets. Introducing hard resets is a reward shaping procedure in that it is a reward transformation (Ng et al. 1999) under which optimal strategies remain invariant.

*Extension to weakly communicating MDPs.* To relax further the communicating assumption on the MDP, we can apply our method to weakly communicating MDPs but the class of specifications will be more restricted. The set of weakly-communicating MDPs is the most general set of MDPs such that there currently exists a learning algorithm that can, using a single stream of experience, guarantee to identify a policy that achieves the optimal average reward rate in the MDP (Bartlett and Tewari 2009; Wan and Sutton 2022). We show how to apply our method in the following for weakly communicating MDPs and fairness specifications.

**Lemma 4.6** (Preservation of Weak Communication). *For a weakly communicating MDP  $\mathcal{M}$  and reward machine  $\mathcal{R}_{\mathcal{A}}$  for a fairness GFM automaton  $\mathcal{A}$ , the product  $\mathcal{M} \times \mathcal{R}_{\mathcal{A}}$  is weakly communicating.*




Fig. 4. Example showing that the lexicographically optimal policy may require infinite memory. Actions are denoted as arrows; the reward is 0 unless otherwise indicated. The specification of interest is  $\text{GF}a$ .

**PROOF.** The proof directly follows the proof of Lemma 4.3. Since we have a weakly-communicating MDP we can partition the states of the MDP into two subsets: the first set is all the transient states under stationary policies, and the second set are the set of states where every pair of states are reachable from each other under a stationary policy. Since the sub-MDP resulting from the second set is communicating, and a fairness property is also an absolute liveness property, then based on Lemma 4.3, the sub-product MDP is communicating. On the other hand, all of the states in the first set are transient, and the fairness property is closed under the addition and deletion of prefixes then eventually we reach a state of the communicating sub-product MDP which leads to the weakly communicating property of the product MDP.  $\square$

## 5 Average-Reward RL for Lexicographic Objectives

This section provides a solution for Problem 2. We first note that an optimal policy for the Problem 2 may require infinite memory. Consider the example in Figure 4. The  $\omega$ -regular specification  $\text{GF}a$  can be satisfied surely by staying in the state  $a$ . Likewise, the maximum possible average reward of 1 can be obtained by staying in the state labeled  $\neg a$ . Surprisingly, there is an infinite memory policy that achieves both. This infinite memory policy visits the state labeled  $\neg a$   $k$  times, then visits the state labeled  $a$ , where  $k$  increases towards infinity forever. Any finite memory policy that visits  $a$  infinitely often with probability 1 must visit the state labeled  $a$  with positive expected frequency in steady state, say  $f > 0$ , and thus achieves a suboptimal external average reward of at most  $(1 - f) < 1$ .

However, one can find finite memory policies that achieve an average reward within  $\varepsilon$  of the optimal average reward for any  $\varepsilon > 0$  while maximizing the probability of satisfaction of the  $\omega$ -regular specification. Namely, consider the policy which visits the state labeled  $\neg a$   $k$  times, then visits the state labeled  $a$  for a fixed  $k$ . This policy obtains an average reward of  $\frac{k}{k+2}$  while satisfying the  $\omega$ -regular specification with probability 1. For any  $\varepsilon > 0$ , one can select  $k \geq \lceil \frac{2}{\varepsilon} - 2 \rceil$  to ensure that  $\frac{k}{k+2} \geq 1 - \varepsilon$ .

In the following, we show that for any  $\varepsilon > 0$ , there always exist a finite-memory policy that is  $\varepsilon$ -optimal for Problem 2 and show how this policy can be constructed. Our approach to solve Problem 2 reduces the task of producing a lexicographically optimal strategy to one of solving a single average reward problem on an MDP, which can then be solved by any off-the-shelf RL algorithm for average reward on communicating MDPs, e.g. Differential Q-learning. The resulting policy is  $\varepsilon$ -optimal and finite memory. We then perform a transformation of this learned policy to yield an infinite memory optimal policy.

We define our average reward reduction via a reward machine. This reward machine augments the usual product with an additional bit. When this bit is zero, the agent receives reward based on the external reward function. With probability  $\beta$  on every step, this bit flips to one. When this bit is one, the agent is incentivized to find an accepting edge and is given a punishing reward of  $\rho(s, s') + c_1$  where  $c_1$  is a constant such that  $c_1 + \max_{s, s'} \rho(s, s') < \min_{s, s'} \rho(s, s')$  on every step. When the agent finds an accepting edge, the bit flips back to zero. To ensure communication, we add two types of resets, one which resets the automaton state back to its initial state ( $\varepsilon_1$ ), and one which sets the extra bit to zero ( $\varepsilon_2$ ). Both of these incur some large penalty of  $c_2$ .




Fig. 5. Picture of the construction of the probabilistic reward machine in (5.1). **Left:** Two layers corresponding to  $b \in \{0, 1\}$ . **Right:** Probabilistic changes of the additional bit  $b$ . The  $\epsilon$ -transitions are excluded for a better pictorial presentation.

We refer to the case when the bit is zero as “the first layer” and the case when the bit is one as “the second layer”. Intuitively, the agent spends most of its time collecting the external average reward, but is occasionally called upon with probability  $\beta$  to prove that it can see an accepting edge. This construction is shown pictorially in Figure 5. We now define the reward machine formally. Let  $\mathcal{A} = (\Sigma, Q, q_0, \delta, F)$  be a GFM for an absolute liveness specification and let  $\rho(s, s')$  be a given reward function. We construct a probabilistic reward machine  $\mathcal{R}_{\mathcal{A}, \rho} = (\Sigma', Q \times \{0, 1\}, (q_0, 0), \delta', \rho')$  where  $\Sigma' = (\Sigma \times S \times S \times Q) \cup \{\epsilon_1, \epsilon_2\}$ ,

$$\delta'((q, b), (l, s, s', q'))((q', b')) = \begin{cases} \beta & b = 0, b' = 1, q' \in \delta(q, l), l \notin \{\epsilon_1, \epsilon_2\} \\ 1 - \beta & b = 0, b' = 0, q' \in \delta(q, l), l \notin \{\epsilon_1, \epsilon_2\} \\ 1 & b = 1, b' = 1, q' \in \delta(q, l), (q, l, q') \notin F, l \notin \{\epsilon_1, \epsilon_2\} \\ 1 & b = 1, b' = 0, q' \in \delta(q, l), (q, l, q') \in F, l \notin \{\epsilon_1, \epsilon_2\} \\ 1 & q \neq q_0, q' = q_0, b' = b, l = \epsilon_1 \\ 1 & q' = q, b = 1, b' = 0, l = \epsilon_2 \\ 0 & \text{otherwise} \end{cases} \quad (5.1)$$

and

$$\rho'((q, b), (l, s, s', q'))((q', b')) = \begin{cases} \rho(s, s') & \text{if } b = 0 \text{ and } l \notin \{\epsilon_1, \epsilon_2\} \\ c_1 + \rho(s, s') & \text{if } b = 1 \text{ and } l \notin \{\epsilon_1, \epsilon_2\} \\ c_2 & \text{if } l \in \{\epsilon_1, \epsilon_2\} \end{cases}$$

where  $c_1 < \min_{s, s'} \rho(s, s') - \max_{s, s'} \rho(s, s')$ ,  $c_2 < 0$  and  $0 < \beta < 1$ .

We now show a few results before proceeding to the main theorems.

**Lemma 5.1** (Preservation of Communication). *For a communicating MDP  $\mathcal{M}$  and reward machine  $\mathcal{R}_{\mathcal{A}, \rho}$  defined above, the resulting product  $\mathcal{M} \times \mathcal{R}_{\mathcal{A}, \rho}$  is communicating.*

**PROOF.** The claim holds due to the additions of  $\epsilon$ -transitions  $\{\epsilon_1, \epsilon_2\}$  that resets both the automaton ( $\epsilon_1$ ) and the extra bit  $b$  to zero ( $\epsilon_2$ ).  $\square$

To develop our formal proof in Theorem 5.3, we make use of the following observation.

**Lemma 5.2** (Binary Probability of Satisfaction). *The probability of satisfaction of an absolute liveness specification is either 0 or 1 in a communicating MDP.*

PROOF. To show this result, we just need to show that if the probability of satisfaction of the specification is positive, then the probability of satisfaction must be 1. This follows from the fact that if the probability of satisfaction is positive, then there is a winning end-component, and this end-component can be reached with probability 1 due to the communicating nature of the product.  $\square$

**Theorem 5.3** (Lexicographic Reward Machine  $\varepsilon$ -optimality). *Let  $v^*$  be the external average reward obtained under a lexicographically optimal strategy for Problem 2. For every  $\varepsilon > 0$ , there exists a threshold  $\beta^* > 0$  such that for all  $0 < \beta < \beta^*$  there exists a  $c^* < 0$  such that for all  $c < c^*$  all positional policies that maximize the average reward on  $\mathcal{M} \times \mathcal{R}_{\mathcal{A},\rho}$  maximize the probability of satisfaction of  $\mathcal{A}$  on  $\mathcal{M}$  and achieves an average reward  $v$  such that  $v \geq v^* - \varepsilon$ . We call such a policy  $\varepsilon$ -optimal. Additionally, these policies are finite memory policies on  $\mathcal{M}$ .*

PROOF. We first note that since the memory required in the resulting policies is defined by  $\mathcal{R}_{\mathcal{A},\rho}$  with fixed parameters, the resulting policies are finite memory. From Lemma 5.2, we can proceed by considering two cases: the absolute liveness  $\omega$ -regular objective is satisfied with probability 0 or 1.

- (1) We first consider the case where the absolute liveness  $\omega$ -regular objective is satisfied with probability 0 under any policy, i.e., there is no policy that can satisfy the specification with positive probability. We will show that any  $0 < \beta^* < 1$  and  $c^* < 0$  works. Since the lexicographically optimal policy only needs to maximize the external average reward, there is an optimal policy that is positional on  $\mathcal{M}$  that achieves an external average reward  $v^*$ . Consider an optimal policy  $\sigma$  on the extended product  $\mathcal{M} \times \mathcal{R}_{\mathcal{A},\rho}$ . In steady state, this policy either takes resets from the second layer to the first infinitely often, or it takes no reset transitions. If the policy takes no resets, then it obtains a value of  $R_\sigma = b + v \leq b + v^*$  where the upper bound  $b + v^*$  is obtained by applying a positional policy optimal for the external reward to the second layer. Since  $\sigma$  is optimal, it must obtain a value of  $b + v^*$ , which obtains an external average reward of  $v^*$  on  $\mathcal{M}$ . If this policy takes resets infinitely often, then it takes the reset transition immediately upon reaching the second layer. This is due to the penalty  $b$ : if a trajectory of length  $T$  in the second layer that ends in a reset obtains a value of  $T(b + v) + c$ , then taking the reset first, and followed by the same trajectory in the first layer obtains a value of  $Tv + c \geq T(b + v) + c$ , where  $v$  is the average reward collected along this trajectory. Thus, this policy obtains a value of  $R_\sigma = (1 - \beta)v_1 + \beta c$  in the extended product, which is maximized when  $v_1 = v^*$ . Since the resets are internal to the agent, this policy collects an external average reward of  $v^*$  on  $\mathcal{M}$ . We have shown that under any  $0 < \beta^* < 1$  and  $c^* < 0$  when the absolute liveness  $\omega$ -regular objective is satisfied with probability 0, an optimal policy on the extended product collects an external average reward of  $v^* \geq v^* - \varepsilon$  for every  $\varepsilon > 0$ .
- (2) We now consider the case where the absolute liveness  $\omega$ -regular objective is satisfied with probability 1 under some policy. Let  $G \subseteq (\mathcal{M} \times \mathcal{A}) \times (\mathcal{M} \times \mathcal{A})$  be the set of transitions such that for every  $t \in G$ , there exists a policy that satisfies  $\mathcal{A}$  with probability 1 and visits  $t$  infinitely often. In other words,  $G$  is the maximal set of transitions that can be visited infinitely often without diminishing the probability of satisfying  $\mathcal{A}$ . Consider a lexicographically optimal strategy. This strategy cannot visit any transition outside of  $G$  infinitely often by definition and achieves an external average reward of  $v^*$ . This implies that there exists a set of positional policies restricted to  $G$  that achieves an average reward of at least  $v^*$ , since there exists optimal positional policies for average reward on MDPs. We now extend these policies to the extended product by applying this policy in the first layer and applying a policy which visits an accepting edge in the second layer without taking resets infinitely often and leaves  $G$  finitely many times, with probability 1. This second layer policy exists due to the definition of  $G$ . We call this set of policies  $\sigma^*$ , and abuse notation by treating  $\sigma^*$  as if it is a single policy for the rest of the proof. We will select  $0 < \beta^* < 1$



and  $c^* < 0$  such that  $\sigma^*$  is optimal in the extended product, and is  $\varepsilon$ -optimal on  $\mathcal{M}$ . Consider all policies  $\Pi_1$  that are positional on the extended product, satisfy the specification with probability 1, and visit the first layer infinitely often with probability 1. We say that the average reward that such a policy  $\sigma \in \Pi_1$  obtains on the extended product for a particular  $\beta$  is  $R_\sigma(\beta)$ . Let  $v_\sigma$  be the external average reward that  $\sigma$  obtains when  $\beta = 0$ . Then, we have that  $(1 - \beta)v_\sigma + \beta b f_\sigma(\beta) \leq R_\sigma(\beta) \leq v_\sigma$  where  $f_\sigma(\beta)$  is an upper bound on the expected time it takes to return to a state in steady state after a transition from the first to the second layer occurs. The first inequality follows by breaking up the average reward into uninterrupted reward collected in the first layer in steady state, and reward collected while waiting to return to the state that was left by transitioning to the second layer in steady state. The second inequality follows from the fact that the reward in the second layer is always less than the first. Note that  $f_\sigma(\beta)$  is monotonically increasing in  $\beta$ . Consider a policy  $\sigma \in \Pi_1$  such that  $v_\sigma < v_{\sigma^*} = v^*$ . Then we have that for  $\beta \leq \min\{\frac{1}{2}, \frac{v^* - v_\sigma}{v^* - b f_{\sigma^*}(\frac{1}{2})}\}$ ,

$$R_\sigma(\beta) \leq v_\sigma \leq (1 - \beta)v_{\sigma^*} + \beta b f_{\sigma^*}(\beta) \leq (1 - \beta)v_{\sigma^*} + \beta b f_{\sigma^*}(\frac{1}{2}) \leq R_{\sigma^*}(\beta)$$

where the third inequality follows from the monotonicity of  $f_{\sigma^*}(\beta)$ . Since there are only finitely many policies in  $\Pi_1$ , one can find a  $0 < \beta_{\text{thresh}} < 1$  that satisfies the above inequality for all policies in  $\Pi_1$ . Thus, for all  $\beta < \beta_{\text{thresh}}$ ,  $\sigma^*$  is an optimal average reward policy on the extended product amongst all the policies in  $\Pi_1$ . Note that the external average reward collected on  $\mathcal{M}$  by any policy  $\sigma$  for a fixed  $\beta$ , denoted by  $\bar{R}_\sigma(\beta)$ , is larger than the average reward collected on the extended product  $R_\sigma(\beta)$  because all of the rewards on the extended product are less than or equal to the external rewards. By selecting  $\beta^* = \min\{\beta_{\text{thresh}}, \frac{\varepsilon}{v^* - b f_{\sigma^*}(\frac{1}{2})}\} > 0$ , we have that

$$v^* - \varepsilon \leq (1 - \beta)v_{\sigma^*} + \beta b f_{\sigma^*}(\frac{1}{2}) \leq R_{\sigma^*}(\beta) \leq \bar{R}_{\sigma^*}(\beta),$$

so  $\sigma^*$  is  $\varepsilon$ -optimal for all  $0 < \beta < \beta^*$ . We will now select  $c^* < 0$  such that there is an optimal policy in  $\Pi_1$ . We denote all positional policies by  $\Pi$ . Note that all policies in  $\Pi \setminus \Pi_1$  either remain in the second layer forever, obtaining an average reward of  $b < \min_{s,s'} \rho(s, s')$  which is less than the average reward for all policies  $\Pi_1$ , or take reset transitions infinitely often. Policies which select reset transitions infinitely often obtain an average reward of at most  $(1 - \beta) \max_{s,s'} \rho(s, s') + \beta c$ . One can then select  $c^* \leq \frac{b}{\beta} - \frac{1 - \beta}{\beta} \max_{s,s'} \rho(s, s') < 0$  to ensure that  $(1 - \beta) \max_{s,s'} \rho(s, s') + \beta c \leq b$ . Since all policies in  $\Pi_1$  obtain an average reward on the extended product greater than  $b$ , there is an optimal policy  $\sigma^*$  in  $\Pi_1$  for  $c < c^*$  with the previously fixed  $\beta$ . This  $\sigma^*$  is an optimal policy on the extended product under these parameter selections and is  $\varepsilon$ -optimal when applied to  $\mathcal{M}$ . □

**Theorem 5.4** (Lexicographic Reward Machine Optimality). *Let  $\beta(i) : \mathbb{N} \rightarrow (0, 1)$  be a sequence such that  $\lim_{i \rightarrow \infty} \beta(i) = 0$ . There exists a threshold  $\beta^* > 0$  such that for all  $0 < \beta < \beta^*$  there exists a  $c^* < 0$  such that for all  $c < c^*$  all positional policies that maximize the average reward on  $\mathcal{M} \times \mathcal{R}_{\mathcal{A}, \rho}$  have the following property: if we transform this fixed, finite memory policy by setting  $\beta = \beta(i)$  on timestep  $i$ , then this transformed policy is optimal for Problem 2 and infinite memory.*

**PROOF.** We first note that this policy is infinite memory: it must keep track of  $\beta(i)$ , which requires an increasing number bits to represent exactly as  $\beta(i)$  decreases. Thus, as  $i$  goes to infinity, the amount of memory required grows in an unbounded manner. For the rest of the proof, we follow the results and proof of Theorem 5.3. Let  $\beta^* = \beta_{\text{thresh}}$  as defined in the proof of Theorem 5.3, and let

$$c^* \leq \frac{b}{\beta} - \frac{1 - \beta}{\beta} \max_{s,s'} \rho(s, s') < 0.$$

We have that there is a set of optimal positional policies  $\sigma^*$  on the extended product that satisfy  $\mathcal{A}$  with the maximum probability and accumulate an external average reward  $\bar{R}_{\sigma^*}(\beta) \geq (1 - \beta)v^* + \beta b f_{\sigma^*}(\frac{1}{2})$  where  $v^*$  is the external average reward obtained under a lexicographically optimal strategy. We now substitute  $\beta$  with  $\beta(i)$ . We have that

$$\lim_{i \rightarrow \infty} \bar{R}_{\sigma^*}(\beta(i)) \geq \lim_{i \rightarrow \infty} (1 - \beta(i))v^* + \beta(i) b f_{\sigma^*}(\frac{1}{2}) = v^*,$$

so  $\sigma^*$  is a lexicographically optimal policy.  $\square$

These results can be extended to weakly communicating MDPs in straightforward fashion.

**Lemma 5.5** (Preservation of Communication). *For a weakly-communicating MDP  $\mathcal{M}$  and reward machine  $\mathcal{R}_{\mathcal{A},\rho}$  over a fairness specification, the resulting product  $\mathcal{M} \times \mathcal{R}_{\mathcal{A},\rho}$  is weakly-communicating.*

**PROOF.** The proof follows directly from Lemma 4.6 and Lemma 5.1. As stated for Lemma 4.6 we can partition the states of the MDP into two sets and the sub-product MDP resulted from the second set with the fairness specification is communicating based on Lemma 5.1. Moreover, since all of the states of the first set, see the proof of Lemma 4.6, are transient the resulting product MDP is weakly communicating.  $\square$

## 6 Experimental Results

We implemented the reduction<sup>1</sup> with hard resets presented in Section 4. As described, we do not build the product MDP explicitly, and instead compose it on-the-fly by keeping track of the MDP and automaton states independently. We use Differential Q-learning (Wan, Naik, et al. 2021) to learn optimal, positional average reward strategies. For our experiments, we have collected a set of communicating MDPs with absolute liveness specifications.<sup>2</sup> All experiments were conducted on 64-bit machine equipped with a 12th Gen Intel Core i7 processor, 32 GB of RAM, and an NVIDIA A2000 GPU.

We compare with two previous approaches for translating absolute liveness  $\omega$ -regular languages to rewards: the method of (Hahn, Perez, et al. 2019) with Q-learning and the method of (Bozkurt, Y. Wang, Zavlanos, et al. 2020) with Q-learning. The method of (Hahn, Perez, et al. 2019) translates a GFM Büchi automaton into a reachability problem through a suitable parameter  $\zeta$ . This reachability problem can be solved with discounted RL by rewarding reaching the target state and using a large enough discount factor. The method of (Bozkurt, Y. Wang, Zavlanos, et al. 2020) uses a state dependent discount factor  $\gamma_B$  and a GFM Büchi automaton. By using a suitable  $\gamma_B$  and large enough discount factor, one can learn optimal strategies for the absolute liveness  $\omega$ -regular objective.

*RQ1. How do previous approaches perform in the continuing setting?* The methods of (Bozkurt, Y. Wang, Zavlanos, et al. 2020; Hahn, Perez, et al. 2019) may produce product MDPs that are not communicating (see Example 4.1). This means that a single continuing run of the MDP may not explore all relevant states and actions. Thus, these methods are not guaranteed to converge without episodic resetting. We studied if this behavior affects these prior methods in practice. As a baseline, we include our proposed approach. Instead of tuning hyperparameters for each method, where hyperparameters that lead to convergence may not exist, we take a sampling approach. We select a wide distribution over hyperparameters for each method and sample 200 hyperparameter combinations for each method and example. We then train for 10 million steps on each combination. The selected hyperparameter distribution is  $\alpha \sim \mathcal{D}(0.01, 0.5)$ ,  $\varepsilon \sim \mathcal{D}(0.01, 1.0)$ ,  $c \sim \mathcal{D}(1, 200)$ ,  $\eta \sim \mathcal{D}(0.01, 0.5)$ ,  $\zeta \sim \mathcal{D}(0.5, 0.995)$ ,  $\gamma_B \sim \mathcal{D}(0.5, 0.995)$ , and  $\gamma \sim \mathcal{D}(0.99, 0.99999)$  where  $\mathcal{D}(a, b)$  is a log-uniform distribution from  $a$  to  $b$ . The end points of these distributions and the training amount are selected by finding hyperparameters which lead to convergence in the episodic setting for these methods.

<sup>1</sup>The implementation is available at <https://plv.colorado.edu/mungojerrie/>.

<sup>2</sup>Case studies are available at <https://doi.org/10.5281/zenodo.15489871>.




Fig. 6. Comparison of the distributions of probability of satisfaction of learned policies across sampled hyperparameters in the continuing setting. For each distribution, the mean is shown as a circle, and the maximum and minimum are shown as vertical bars. We compare our proposed reduction, the reduction of Hahn et al. (Hahn, Perez, et al. 2019) with Q-learning, and the reduction of Bozkurt et al. (Bozkurt, Y. Wang, Zavlanos, et al. 2020) with Q-learning. Episodic resetting is not used.

Figure 6 shows the resulting distribution over runs. A distribution entirely at 0 (1) indicates that all sampled runs produced strategies that satisfy the specification with probability 0 (1). For many examples, prior approaches have no successful hyperparameter combinations, with distributions centered entirely at 0. However, our proposed approach always have some hyperparameters that lead to optimal, probability 1, strategies, as indicated by the tails of the distributions touching the probability 1 region of the plot.

*RQ2. How does our method compare to previous approaches when we allow episodic setting?* By allowing episodic resetting, we can now find hyperparameters for methods of (Bozkurt, Y. Wang, Zavlanos, et al. 2020; Hahn, Perez, et al. 2019) that lead to convergence. We tuned all hyperparameters by hand to minimize training time, while verifying with a model checker that the produced policies are optimal. Table 1 shows learning times, as well as hyperparameters for our reduction. We report the number of states reachable in the MDP and the product, learning times averaged over 5 runs, the reset penalty  $c$ , the  $\epsilon$ -greedy exploration rate  $\epsilon$ , the learning rates  $\alpha$  and  $\eta$  of the Differential Q-learning, as well as the number of training steps. Note that we do not do any episodic resetting when training with our reduction. This means that the RL agent must learn to recover from mistakes during training, while previous approaches are periodically reset to a good initial state. *Our reduction using Differential Q-learning is competitive with previous approaches while not being reliant on episodic resetting.*

*RQ3. How our approach works in multi-objective settings?* As shown in Table 2, we successfully satisfied the absolute liveness property with probability one across all case studies. We report the number of states reachable in the MDP and the product, the learning times averaged over five runs, the average reward over five runs, the reset penalty  $c$  and  $\beta$ , the  $\epsilon$ -greedy exploration rate  $\epsilon$ , the learning rates  $\alpha$  and  $\eta$  of Differential Q-learning, as well as the number of training steps. Note that, also in the multi-objective setting, we do not perform any episodic resetting when training with our reduction. We achieved average rewards that are almost optimal in all cases (the optimal value for the first case study is 1 and for the rest is 0).

Table 1. Learning results and comparison. Hyperparameters used for our reduction are shown. Blank entries indicate that default values are used. The default parameters are  $c = -1$ ,  $\varepsilon = 0.1$ ,  $\alpha = 0.1$ , and  $\eta = 0.1$ . Times are in seconds. Superscript <sup>†</sup> indicates results from Q-learning with reduction from (Hahn, Perez, et al. 2019), while superscript <sup>‡</sup> indicates Q-learning with reduction from (Bozkurt, Y. Wang, Zavlanos, et al. 2020). Results for <sup>†</sup> and <sup>‡</sup> require episodic resetting. All hyperparameters are tuned by hand.

Name	states	prod.	time	time <sup>†</sup>	time <sup>‡</sup>	$c$	$\varepsilon$	$\alpha$	$\eta$	train-steps
adverse	202	507	8.51	7.09	12.56	-150		0.2		10M
frozenSmall	16	64	0.99	20.23	9.88					500k
frozenLarge	64	256	4.07	3.88	8.79			0.02	0.02	3M
windy	123	366	1.40	1.81	2.61		0.95	0.5	0.05	1M
windyStoch	130	390	2.97	3.91	2.53			0.5		2M
grid5x5	25	100	0.62	1.12	1.02			0.5		200k
ishift	4	29	0.03	0.01	0.02					10k
doublegrid	1296	5183	16.43	3.45	3.09	-2	0.5	0.05	0.01	12M
busyRingMC2	72	288	0.03	0.03	0.03				0.01	10k
busyRingMC4	2592	15426	6.08	3.94	2.33				0.01	1.5M

Table 2. Learning results for multi-objective case studies. The default parameters are  $\text{tol} = 0$ ,  $\beta = 0.05$ ,  $\text{ep-n} = 1$ ,  $\text{ep-l} = 1000000$ ,  $\alpha = 0.01$ ,  $\eta = 0.01$ ,  $\varepsilon = 0.1$ . Times are in seconds. All hyperparameters are tuned by hand.

Name	states	prod.	time	Prob.	rew.	$c$
infmem	2	8	0.758	1	0.999	100
vessels	24	576	1.115	1	-0.100	50
grid4x4c4	16	512	1.959	1	0.004	50
grid4x4c2	16	256	1.930	1	0.002	50
grid5x5c4	25	800	1.977	1	-0.001	60
grid5x5c2	25	400	1.872	1	-0.001	60
grid4x7	28	896	1.741	1	0.000	100

## 7 Related Work

*Learning for Formal Specifications over Finite Traces.* The development and use of formal reward structures for RL have witnessed increased interest in recent years. For episodic RL, logics have been developed over finite traces of the agent’s behavior, including  $\text{LTL}_f$  and Linear Dynamic Logic ( $\text{LDL}_f$ ) (Camacho, J. A. Baier, et al. 2018; De Giacomo and M. Vardi 2015). These logics have equivalent automaton and reward machine representations that have catalyzed a series of efforts on defining novel reward shaping functions to accelerate the convergence of RL algorithms subject to formal specifications (Camacho, Icarte, et al. 2019; De Giacomo, Iocchi, et al. 2019; Icarte, T. Klassen, et al. 2018). These methods leverage the graph structure of the automaton to provide an artificial reward signal to the agent. More recently, dynamic reward shaping using  $\text{LTL}_f$  has been introduced as a means to both learn the transition values of a given reward machine and leverage these values for reward shaping and transfer learning (Velasquez et al. 2021). There has also been work on learning or synthesizing the entire structure of such reward machines from agent interactions with the environment by leveraging techniques from satisfiability and active grammatical inference (Gaon and Brafman 2020; Neider et al. 2021; Toro Icarte et al. 2019; Xu, Gavran, et al. 2020; Xu, Wu, et al. 2021). Other data-driven methods have also been developed recently for satisfying specifications over finite traces. Data-driven distributionally robust policy synthesis approaches are presented in (Kordabad et al. 2025; Schön, Zhong, et al. 2024). Data-driven construction of abstractions with correctness guarantees has also been studied (Kazemi, Majumdar, et al. 2024; Nazeri et al. 2025; Schön, Naseer, et al. 2024), which enables policy synthesis against temporal specifications. Data-driven computation of resilience with respect to finite temporal logic specifications is studied in (Saoud et al. 2024).

*Formal Specifications over Infinite Traces.* For the infinite-trace settings, LTL has been extensively used to verify specifications and synthesize policies formally using the mathematical model of a system (C. Baier and Katoen 2008; Blom et al. 2006; Haesaert, Nilsson, et al. 2021; Haesaert and Soudjani 2020; Kazemi, Perez, et al. 2024; Kazemi and Soudjani 2020; Lavaei, Somenzi, et al. 2020; Lavaei, Soudjani, et al. 2022; Majumdar, Mallik, and Soudjani 2020; Majumdar, Sathiyarayanan, et al. 2024; Schön, Huijgevoort, et al. 2022; Yoon et al. 2019). Considering the generality of the available results in terms of structure of the underlying MDP, most of the research focuses on discounted reward structures. Despite the simplicity of discounted Markov decision problems, the discounted reward structure (unlike average reward) prioritizes the transient response of the system. However, application of the average reward objective because of the restriction over the structure of the MDP is limited. The work (Ding et al. 2014) proposes a policy iteration algorithm for satisfying specifications of the form  $G F \phi \wedge \psi$  for a communicating MDP almost surely. The work (Ashok et al. 2017) proposes a value iteration algorithm for solving the average reward problem for multichain MDPs, where the algorithm first computes the optimal value for each of strongly connected components and then weighted reachability to find the optimal policy. The work (G. K. Atia et al. 2021) provides a linear program for policy synthesis of multichain MDPs with steady-state constraints. The work (Majumdar, Salamati, et al. 2025) provides a regret-free learning algorithm for policy synthesis that is based on identifying the structure of the underlying MDP using data with a certain confidence.

*RL for Formal Specifications over Infinite Traces.* In the last few years, researchers have started developing data-driven policy synthesis techniques in order to satisfy temporal specifications. There is a large body of literature in safe reinforcement learning (RL) (see e.g. (Efroni et al. 2020; Garcia and Fernández 2015; Recht 2018)). The problem of learning a policy to maximize the satisfaction probability of a temporal specification using discounted RL is studied recently (Bozkurt, Y. Wang, Zavlanos, et al. 2020; Brázdil et al. 2014; Fu and Topcu 2014; H. Hasanbeig et al. 2023; M. Hasanbeig et al. 2019; Oura et al. 2020; Sadigh et al. 2014). The work (Hahn, Perez, et al. 2019) uses a parameterized augmented MDP to provide an RL-based policy synthesis for finite MDPs with unknown transition probabilities. It shows that the optimal policy obtained by RL for the reachability probability on the augmented MDP gives a policy for the MDP with a suitable convergence guarantee

In (Bozkurt, Y. Wang, Zavlanos, et al. 2020) authors provide a path-dependent discounting mechanism for the RL algorithm based on a limit-deterministic Büchi automaton (LDBA) representation of the underlying  $\omega$ -regular specification, and prove convergence of their approach on finite MDPs when the discounting factor goes to one. An LDBA is also leveraged in (H. Hasanbeig et al. 2023; M. Hasanbeig et al. 2019; Oura et al. 2020) for discounted-reward model-free RL in both continuous- and discrete-state MDPs. The LDBA is used to define a reward function that incentivizes the agent to visit all accepting components of the automaton. The paper (Falah et al. 2023) proposes a translation from  $\omega$ -regular properties into discounted reward objectives for model-free RL in continuous-time MDPs utilizing similar automata-based reward functions.

These works use episodic discounted RL with discount factor close to one to solve the policy synthesis problem. There are two issues with the foregoing approaches. First, because of the episodic nature of the algorithms they are not applicable in continuing settings. Second, because of high discount factors in practice these algorithms are difficult to converge. On the other hand, recent work on reward shaping for average reward RL has been explored based on safety specifications to be satisfied by the synthesized policy (Y. Jiang et al. 2021). In contrast to the solution proposed in this paper, the preceding approach requires knowledge of the graph structure of the underlying MDP and does not account for absolute liveness specifications.

*Average-Reward RL.* There is a rich history of studies in average reward RL (Dewanto et al. 2020; Mahadevan 1996; Tadepalli and Ok 1998). Lack of stopping criteria for multichain MDPs affect the generality of model-free RL algorithms. Therefore, all model-free RL algorithms put some restrictions on the structure of MDP (e.g. ergodicity (Abbasi-Yadkori et al. 2019; Wei et al. 2020) or communicating property). The closest line of work to this work is to use average reward objective for safe RL. The work (Singh et al. 2020) proposes a model-based RL algorithm for



maximizing average reward objective with safety constraint for communicating MDPs. It is worth noting that in multichain setting, the state-of-the-art learning algorithms use model-based RL algorithms. The work (Kretínský et al. 2020) studies satisfaction of  $\omega$ -regular specifications using data-driven approaches. The authors introduce an algorithm where the optimality of the policy is conditioned to not leaving the corresponding maximal end component which leads to a sub-optimal solution. The authors provide PAC analysis for the algorithm as well. Finally, there is a growing area of research on robust average-reward RL, whereby the transition probability function is defined over an uncertainty set, e.g., a KL divergence set (Y. Wang, Velasquez, G. Atia, et al. 2024; Y. Wang, Velasquez, G. K. Atia, et al. 2023). The problem of finding an optimal policy in that context is to maximize the minimum expected reward within that uncertainty set of transition probabilities.

*Multi-objective Probabilistic Verification.* Chatterjee, Majumadar, and Henzinger (Chatterjee et al. 2006) considered MDPs with multiple discounted reward objectives. In the presence of multiple objectives, the trade-offs between different objectives can be characterized as Pareto curves. The authors of (Chatterjee et al. 2006) showed that every Pareto optimal point can be achieved by a memoryless strategy and the Pareto curve can be approximated in polynomial time. Moreover, the problem of checking the existence of a strategy that realizes a value vector can be decided in polynomial time. These multi-objective optimization problems were studied in the context of multiple long-run average objectives by Chatterjee (Chatterjee 2007). He showed that the Pareto curve can be approximated in polynomial time in the size of the MDP for irreducible MDPs and in polynomial space in the size of the MDP for general MDPs. Additionally, the problem of checking the existence of a strategy that guarantees values for different objectives to be equal to a given vector is in polynomial time for irreducible MDPs and in NP for general MDPs. Etessami et al. (Etessami et al. 2007) were the first to study the multi-objective model-checking problem for MDPs with  $\omega$ -regular objectives. Given probability intervals for the satisfaction of various properties, they developed a polynomial-time (in the size of the MDP) algorithm to decide the existence of such a strategy. They also showed that, in general, such strategies may require both randomization and memory. That paper also studies the approximation of the Pareto curve with respect to a set of  $\omega$ -regular properties in time polynomial in the size of the MDP. Forejt et al. (Forejt et al. 2011) studied quantitative multi-objective optimization over MDPs that combines  $\omega$ -regular and quantitative objectives. Those algorithms are implemented in the probabilistic model checker PRISM (Kwiatkowska et al. 2011). Multi-objective optimization on interval MDPs is studied by Monir et al. (Monir et al. 2024) to provide a Lyapunov-based policy synthesis approach.

*Multi-objective Reinforcement Learning.* There has been substantial work on lexicographic objectives in RL, including lexicographic discounted objectives (Skalse et al. 2022), lexicographic  $\omega$ -regular objectives (Hahn, Perez, et al. 2021), and a combination of safety and discounted objectives (Bozkurt, Y. Wang, and Pajic 2021). Hahn et al. (Hahn, Perez, et al. 2023) considered the general class of  $\omega$ -regular objectives with discounted rewards in model-free RL. However, to the best of our knowledge, this paper is the first work that considers lexicographic  $\omega$ -regular objectives with average objectives.

*Average-Reward RL for Formal Specifications.* Despite significant progress in data-driven approaches for satisfying  $\omega$ -regular specifications, there remains a gap in the use of average-reward, model-free RL algorithms for satisfying temporal logic specifications. While the paper (Alur et al. 2022) proves non-existence of robust, optimality-preserving reductions from general LTL specifications to discounted cumulative rewards, it identifies the reduction to average rewards as an open problem. This open problem has been addressed in the paper (Le et al. 2024) that shows that one can translate any specification in the general class of  $\omega$ -regular properties into average objectives and then approximate them successively with discounted objectives while preserving optimality in the limit in the episodic setting. Our preliminary work presented at AAMAS 2022 (Kazemi, Perez, et al. 2022) also aimed to address this gap without relying on successive discounted RL in episodic setting but by proposing a model-free average-reward RL algorithm tailored to a subclass of LTL specifications known as absolute liveness

specifications. We argue that this subclass captures a broad range of practically relevant specifications and is particularly well-suited to the average-reward RL setting.

This manuscript builds on (Kazemi, Perez, et al. 2022) and extends the results in the following directions: **(a)** we generalize our results from communicating MDPs to the broader class of weakly communicating MDPs; **(b)** we incorporate a lexicographic multi-objective optimization framework, wherein the goal is to maximize a mean-payoff objective among policies that satisfy a given liveness specification; **(c)** we provide novel automata-theoretic characterizations of absolute liveness and stable specifications; and **(d)** we conduct an experimental evaluation demonstrating the effectiveness of the proposed lexicographic multi-objective approach.

## 8 Conclusion

This work addressed the problem of synthesizing policies that satisfy a given absolute liveness  $\omega$ -regular specification while maximizing a mean-payoff objective in the continuing setting. Our first key contribution is a *model-free translation* from  $\omega$ -regular specifications to an average-reward objective, enabling the use of off-the-shelf average-reward reinforcement learning algorithms. In contrast to existing approaches that rely on discounted, episodic learning, which require environment resets and may be infeasible in many real-world settings, our approach learns optimal policies in a single, life-long episode without resetting.

Our second contribution is a solution to a *lexicographic multi-objective optimization problem*, where the goal is to maximize a mean-payoff objective among the set of policies that satisfy a given liveness specification. We provide convergence guarantees under the assumption that the underlying Markov Decision Process is communicating. Importantly, our solutions are *model-free* and do not require access to the environment’s transition structure or its graph representation. This removes the common assumption in prior work that synthesis requires the computation of end components in the product MDP.

We implemented our approach using Differential Q-learning and evaluated it on a range of case studies. The experimental results demonstrate that our method reliably converges to optimal strategies under the stated assumptions and outperforms existing approaches in the continuing setting. These results support the important and often overlooked hypothesis that *average-reward RL is better suited to continuing tasks than discounted RL*. As future work, we plan to explore the use of *function approximation*, with the aim of enabling average-reward RL to achieve the same level of success for continuing tasks as discounted RL has achieved in episodic settings.

## Acknowledgments

The research of S. Soudjani is supported by the following grants: EIC 101070802 and ERC 101089047. This research was also supported in part by the National Science Foundation (NSF) through CAREER Award CCF-2146563. Ashutosh Trivedi holds the position of Royal Society Wolfson Visiting Fellow and gratefully acknowledges the support of the Wolfson Foundation and the Royal Society for this fellowship.

## References

- Y. Abbasi-Yadkori, P. Bartlett, K. Bhatia, N. Lazic, C. Szepesvari, and G. Weisz. 2019. “Politex: Regret bounds for policy iteration using expert prediction.” In: *International Conference on Machine Learning*. PMLR, 3692–3702.
- B. Alpern and F. B. Schneider. 1985. “Defining liveness.” *Information processing letters*, 21, 4, 181–185.
- R. Alur, S. Bansal, O. Bastani, and K. Jothimurugan. 2022. “A framework for transforming specifications in reinforcement learning.” In: *Principles of systems design: essays dedicated to Thomas A. Henzinger on the occasion of his 60th birthday*. Springer, 604–624.
- P. Ashok, K. Chatterjee, P. Daca, J. Křetínský, and T. Meggendorfer. 2017. “Value iteration for long-run average reward in Markov decision processes.” In: *International Conference on Computer Aided Verification*. Springer, 201–221.
- G. K. Atia, A. Beckus, I. Alkhoury, and A. Velasquez. 2021. “Steady-state planning in expected reward multichain MDPs.” *Journal of Artificial Intelligence Research*, 72, 1029–1082.
- C. Baier and J.-P. Katoen. 2008. *Principles of Model Checking*. MIT Press, Cambridge, MA.

- P. L. Bartlett and A. Tewari. 2009. “REGAL: a regularization based algorithm for reinforcement learning in weakly communicating MDPs.” In: *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, 35–42.
- M. G. Bellemare, G. Ostrovski, A. Guez, P. Thomas, and R. Munos. 2016. “Increasing the action gap: New operators for reinforcement learning.” In: *Proceedings of the AAAI conference on artificial intelligence* 1. Vol. 30.
- C. Belta and S. Sadraddini. 2019. “Formal methods for control synthesis: An optimization perspective.” *Annual Review of Control, Robotics, and Autonomous Systems*, 2, 115–140.
- H. A. Blom, J. Lygeros, M. Everdij, S. Loizou, and K. Kyriakopoulos. 2006. *Stochastic hybrid systems: theory and safety critical applications*. Vol. 337. Springer, Heidelberg.
- A. K. Bozkurt, Y. Wang, and M. Pajic. 2021. “Model-free learning of safe yet effective controllers.” In: *2021 60th IEEE Conference on Decision and Control (CDC), Austin, TX, USA, December 14-17, 2021*. IEEE, 6560–6565. doi:10.1109/CDC45484.2021.9683634.
- A. K. Bozkurt, Y. Wang, M. M. Zavlanos, and M. Pajic. 2020. “Control synthesis from linear temporal logic specifications using model-free reinforcement learning.” In: *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 10349–10355.
- T. Brázdil, K. Chatterjee, M. Chmelik, V. Forejt, J. Křetínský, M. Kwiatkowska, D. Parker, and M. Ujma. 2014. “Verification of Markov decision processes using learning algorithms.” In: *Automated Technology for Verification and Analysis (ATVA)*. Springer, 98–114.
- A. Brunello, A. Montanari, and M. Reynolds. 2019. “Synthesis of LTL formulas from natural language texts: State of the art and research directions.” In: *26th International Symposium on Temporal Representation and Reasoning (TIME 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- A. Camacho, J. A. Baier, C. Muise, and S. A. McIlraith. 2018. “Finite LTL synthesis as planning.” In: *Twenty-Eighth International Conference on Automated Planning and Scheduling*.
- A. Camacho, R. T. Icarte, T. Q. Klassen, R. A. Valenzano, and S. A. McIlraith. 2019. “LTL and beyond: Formal languages for reward function specification in reinforcement learning.” In: *IJCAI*. Vol. 19, 6065–6073.
- K. Chatterjee. 2007. “Markov decision processes with multiple long-run average objectives.” In: *FSTTCS 2007: Foundations of Software Technology and Theoretical Computer Science*. Ed. by V. Arvind and S. Prasad. Springer Berlin Heidelberg, Berlin, Heidelberg, 473–484.
- K. Chatterjee, R. Majumdar, and T. A. Henzinger. 2006. “Markov decision processes with multiple objectives.” In: *Annual symposium on theoretical aspects of computer science*. Springer, 325–336.
- L. De Alfaro. 1998. “Formal verification of probabilistic systems.” Ph.D. Dissertation. Stanford University.
- G. De Giacomo, L. Iocchi, M. Favorito, and F. Patrizi. 2019. “Foundations for restraining bolts: Reinforcement learning with LTLf/LDLf restraining specifications.” In: *Proceedings of the International Conference on Automated Planning and Scheduling*. Vol. 29(1), 128–136.
- G. De Giacomo and M. Vardi. 2015. “Synthesis for LTL and LDL on finite traces.” In: *Twenty-Fourth International Joint Conference on Artificial Intelligence*.
- V. Dewanto, G. Dunn, A. Eshragh, M. Gallagher, and F. Roosta. 2020. “Average-reward model-free reinforcement learning: a systematic review and literature mapping.” *arXiv preprint arXiv:2010.08920*.
- X. Ding, S. L. Smith, C. Belta, and D. Rus. 2014. “Optimal control of Markov decision processes with linear temporal logic constraints.” *IEEE Transactions on Automatic Control*, 59, 5, 1244–1257.
- Y. Efroni, S. Mannor, and M. Pirodda. 2020. “Exploration-exploitation in constrained MDPs.” *arXiv preprint arXiv:2003.02189*.
- K. Etessami, M. Kwiatkowska, M. Y. Vardi, and M. Yannakakis. 2007. “Multi-objective model checking of Markov decision processes.” In: *Tools and Algorithms for the Construction and Analysis of Systems*. Ed. by O. Grumberg and M. Huth. Springer Berlin Heidelberg, Berlin, Heidelberg, 50–65.
- A. Falah, S. Guha, and A. Trivedi. 2023. “Reinforcement learning for omega-regular specifications on continuous-time MDP.” In: *Proceedings of the International Conference on Automated Planning and Scheduling*. Vol. 33, 578–586.
- V. Forejt, M. Kwiatkowska, G. Norman, D. Parker, and H. Qu. 2011. “Quantitative multi-objective verification for probabilistic systems.” In: *Tools and Algorithms for the Construction and Analysis of Systems*. Ed. by P. A. Abdulla and K. R. M. Leino. Springer Berlin Heidelberg, Berlin, Heidelberg, 112–127.
- J. Fu and U. Topcu. 2014. “Probably approximately correct MDP learning and control With temporal logic constraints.” In: *Proceedings of Robotics: Science and Systems*.
- M. Gaon and R. Brafman. 2020. “Reinforcement learning with non-Markovian rewards.” In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 34(04), 3980–3987.
- J. Garcia and F. Fernández. 2015. “A comprehensive survey on safe reinforcement learning.” *Journal of Machine Learning Research*, 16, 1, 1437–1480.
- S. Haesaert, P. Nilsson, and S. Soudjani. 2021. “Formal multi-objective synthesis of continuous-state MDPs.” *IEEE Control Systems Letters*, 5, 5, 1765–1770.
- S. Haesaert and S. Soudjani. 2020. “Robust dynamic programming for temporal logic control of stochastic systems.” *IEEE Transactions on Automatic Control*, 66, 6, 2496–2511.
- E. M. Hahn, G. Li, S. Schewe, A. Turrini, and L. Zhang. 2015. “Lazy probabilistic model checking without determinisation.” In: *International Conference on Concurrency Theory (CONCUR)*, 354–367.

- E. M. Hahn, M. Perez, S. Schewe, F. Somenzi, A. Trivedi, and D. Wojtczak. 2020. “Good-for-MDPs automata for probabilistic analysis and reinforcement learning.” In: *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. Springer, 306–323.
- E. M. Hahn, M. Perez, S. Schewe, F. Somenzi, A. Trivedi, and D. Wojtczak. 2021. “Model-free reinforcement learning for lexicographic omega-regular objectives.” In: *International symposium on formal methods*. Springer, 142–159.
- E. M. Hahn, M. Perez, S. Schewe, F. Somenzi, A. Trivedi, and D. Wojtczak. 2019. “Omega-regular objectives in model-free reinforcement learning.” In: *International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*. Springer, 395–412.
- E. M. Hahn, M. Perez, S. Schewe, F. Somenzi, A. Trivedi, and D. Wojtczak. 2023. “Omega-regular reward machines.” In: *ECAI 2023*. IOS Press, 972–979.
- H. Hasanbeig, D. Kroening, and A. Abate. 2023. “Certified reinforcement learning with logic guidance.” *Artificial Intelligence*, 322, 103949.
- M. Hasanbeig, Y. Kantaros, A. Abate, D. Kroening, G. J. Pappas, and I. Lee. 2019. “Reinforcement learning for temporal logic control synthesis with probabilistic satisfaction guarantees.” In: *IEEE Conference on Decision and Control (CDC)*. IEEE, 5338–5343.
- J. Hu, Y. Paliwal, H. Kim, Y. Wang, and Z. Xu. 2024. “Reinforcement learning with predefined and inferred reward machines in stochastic games.” *Neurocomputing*, 599, 128170.
- R. T. Icarte, T. Klassen, R. Valenzano, and S. McIlraith. 2018. “Using reward machines for high-level task specification and decomposition in reinforcement learning.” In: *International Conference on Machine Learning*. PMLR, 2107–2116.
- R. T. Icarte, T. Q. Klassen, R. Valenzano, and S. A. McIlraith. 2022. “Reward machines: Exploiting reward function structure in reinforcement learning.” *Journal of Artificial Intelligence Research*, 73, 173–208.
- P. Jagtap, S. Soudjani, and M. Zamani. 2020. “Formal synthesis of stochastic systems via control barrier certificates.” *IEEE Transactions on Automatic Control*.
- Y. Jiang, S. Bharadwaj, B. Wu, R. Shah, U. Topcu, and P. Stone. 2021. “Temporal-logic-based reward shaping for continuing reinforcement learning tasks.” *Good Systems-Published Research*.
- M. Kazemi, R. Majumdar, M. Salamati, S. Soudjani, and B. Wooding. 2024. “Data-driven abstraction-based control synthesis.” *Nonlinear Analysis: Hybrid Systems*, 52, 101467.
- M. Kazemi, M. Perez, F. Somenzi, S. Soudjani, A. Trivedi, and A. Velasquez. 2024. “Assume-guarantee reinforcement learning.” In: *Proceedings of the AAAI Conference on Artificial Intelligence* 19. Vol. 38, 21223–21231.
- M. Kazemi, M. Perez, F. Somenzi, S. Soudjani, A. Trivedi, and A. Velasquez. 2022. “Translating omega-regular specifications to average objectives for model-free reinforcement learning.” In: *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 732–741.
- M. Kazemi and S. Soudjani. 2020. “Formal policy synthesis for continuous-state systems via reinforcement learning.” In: *International Conference on Integrated Formal Methods*. Springer, 3–21.
- A. B. Kordabad, E. E. Vlahakis, L. Lindemann, S. Gros, D. V. Dimarogonas, and S. Soudjani. 2025. “Data-driven distributionally robust control for interacting agents under logical constraints.” *arXiv preprint arXiv:2503.09816*.
- H. Kress-Gazit, M. Lahijanian, and V. Raman. 2018. “Synthesis for robots: Guarantees and feedback for robot behavior.” *Annual Review of Control, Robotics, and Autonomous Systems*, 1, 211–236.
- J. Kretínský, F. Michel, L. Michel, and G. Perez. 2020. “Finite-memory near-optimal learning for Markov decision processes with long-run average reward.” In: *Conference on Uncertainty in Artificial Intelligence*. PMLR, 1149–1158.
- M. Kwiatkowska, G. Norman, and D. Parker. July 2011. “PRISM 4.0: Verification of probabilistic real-time systems.” In: *Computer Aided Verification (CAV)*. LNCS 6806. (July 2011), 585–591.
- A. Lavaei, F. Somenzi, S. Soudjani, A. Trivedi, and M. Zamani. 2020. “Formal controller synthesis for continuous-space MDPs via model-free reinforcement learning.” In: *International Conference on Cyber-Physical Systems (ICCPS)*, 98–107.
- A. Lavaei, S. Soudjani, A. Abate, and M. Zamani. 2022. “Automated verification and synthesis of stochastic hybrid systems: A survey.” *Automatica*, 146, 110617.
- X. B. Le, D. Wagner, L. Witzman, A. Rabinovich, and L. Ong. 2024. “Reinforcement learning with LTL and omega-regular objectives via optimality-preserving translation to average rewards.” *Advances in Neural Information Processing Systems*, 37, 117109–117132.
- L. Lehnert, R. Laroche, and H. van Seijen. 2018. “On value function representation of long horizon problems.” In: *Proceedings of the AAAI Conference on Artificial Intelligence* 1. Vol. 32.
- S. Mahadevan. 1996. “Average reward reinforcement learning: Foundations, algorithms, and empirical results.” *Machine learning*, 22, 1, 159–195.
- R. Majumdar, K. Mallik, A.-K. Schmuck, and S. Soudjani. 2021. “Symbolic qualitative control for stochastic systems via finite parity games.” *IFAC-PapersOnLine*, 54, 5, 127–132.
- R. Majumdar, K. Mallik, and S. Soudjani. 2020. “Symbolic controller synthesis for Büchi specifications on stochastic systems.” In: *Hybrid Systems: Computation and Control (HSCC)* Article 14. ACM, Sydney, New South Wales, Australia, 11 pages. ISBN: 9781450370189.
- R. Majumdar, M. Salamati, and S. Soudjani. 2025. “Regret-free reinforcement learning for LTL specifications.” *International Conference on Machine Learning (ICML)*.

- R. Majumdar, V. Sathiyarayanan, and S. Soudjani. 2024. "Necessary and sufficient certificates for almost sure reachability." *IEEE Control Systems Letters*.
- Z. Manna and A. Pnueli. 1990. "A hierarchy of temporal properties." In: *Proceedings of the ninth annual ACM symposium on Principles of distributed computing*, 377–410.
- H. Mao, M. Alizadeh, I. Menache, and S. Kandula. 2016. "Resource management with deep reinforcement learning." In: *Proceedings of the 15th ACM workshop on hot topics in networks*, 50–56.
- A. Mirhoseini et al. 2021. "A graph placement methodology for fast chip design." *Nature*, 594, 7862, 207–212.
- V. Mnih et al. Feb. 2015. "Human-level control through reinforcement learning." *Nature*, 518, (Feb. 2015), 529–533.
- N. Monir, O. Schön, and S. Soudjani. 2024. "Lyapunov-based policy synthesis for multi-objective interval MDPs." *IFAC-PapersOnLine*, 58, 11, 99–106.
- A. Naik, R. Shariff, N. Yasui, and R. Sutton. 2019. "Discounted reinforcement learning is not an optimization problem." In: *Optimization Foundations of Reinforcement Learning Workshop at NeurIPS 2019, Vancouver, Canada*.
- M. Nazeri, T. Badings, S. Soudjani, and A. Abate. 2025. "Data-driven yet formal policy synthesis for stochastic nonlinear dynamical systems." *Learning for Decision and Control (LADC)*.
- D. Neider, J.-R. Gaglione, I. Gavran, U. Topcu, B. Wu, and Z. Xu. 2021. "Advice-Guided Reinforcement Learning in a non-Markovian Environment." In: *Proceedings of the AAAI Conference on Artificial Intelligence*.
- A. Y. Ng, D. Harada, and S. Russell. 1999. "Policy invariance under reward transformations: Theory and application to reward shaping." In: *ICML*. Vol. 99, 278–287.
- R. Oura, A. Sakakibara, and T. Ushio. July 2020. "Reinforcement learning of control policy for linear temporal logic specifications using limit-deterministic Büchi automata." *IEEE Control Systems Letters*, 4, 3, (July 2020), 761–766.
- M. L. Puterman. 1994. *Markov decision processes: Discrete stochastic dynamic programming*. (1st ed.). John Wiley & Sons, Inc., New York, NY, USA. ISBN: 0471619779.
- B. Recht. 2018. "A tour of reinforcement learning: The view from continuous control." *Annual Review of Control, Robotics, and Autonomous Systems*, 253–279.
- D. Sadigh, E. S. Kim, S. Coogan, S. S. Sastry, and S. A. Seshia. 2014. "A learning based approach to control synthesis of Markov decision processes for linear temporal logic specifications." In: *Conference on Decision and Control*, 1091–1096.
- A. Saoud, P. Jagtap, and S. Soudjani. 2024. "Temporal logic resilience for dynamical systems." *IEEE Transactions on Automatic Control (under review)*.
- O. Schön, B. van Huijgevoort, S. Haesaert, and S. Soudjani. 2022. "Correct-by-design control of parametric stochastic systems." *61st IEEE Conference on Decision and Control (CDC)*.
- O. Schön, S. Naseer, B. Wooding, and S. Soudjani. 2024. "Data-driven abstractions via binary-tree Gaussian processes for formal verification." *IFAC-PapersOnLine*, 58, 11, 115–122.
- O. Schön, Z. Zhong, and S. Soudjani. 2024. "Data-driven distributionally robust safety verification using barrier certificates and conditional mean embeddings." In: *2024 American Control Conference (ACC)*. IEEE, 3417–3423.
- J. Schrittwieser et al. 2020. "Mastering atari, Go, chess and shogi by planning with a learned model." *Nature*, 588, 7839, 604–609.
- M. L. Shehab, A. Aspel, and N. Ozay. 2025. "Learning reward machines from partially observed policies." *arXiv preprint arXiv:2502.03762*.
- S. Sickert, J. Esparza, S. Jaax, and J. Křetínský. 2016. "Limit-deterministic Büchi automata for linear temporal logic." In: *International Conference on Computer Aided Verification (CAV)*. Springer, 312–332.
- D. Silver, A. Huang, et al. 2016. "Mastering the game of Go with deep neural networks and tree search." *nature*, 529, 7587, 484–489.
- D. Silver, T. Hubert, et al. 2018. "A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play." *Science*, 362, 6419, 1140–1144.
- D. Silver, J. Schrittwieser, et al. 2017. "Mastering the game of Go without human knowledge." *Nature*, 550, 7676, 354.
- R. Singh, A. Gupta, and N. B. Shroff. 2020. "Learning in Markov decision processes under constraints." *arXiv preprint arXiv:2002.12435*.
- A. P. Sistla. 1994. "Safety, liveness and fairness in temporal logic." *Formal Aspects of Computing*, 6, 5, 495–511.
- J. Skalse, L. Hammond, C. Griffin, and A. Abate. July 2022. "Lexicographic Multi-Objective Reinforcement Learning." In: *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-2022*. (July 2022), 3430–3436.
- A. L. Strehl, L. Li, E. Wiewiora, J. Langford, and M. L. Littman. 2006. "PAC model-free reinforcement learning." In: *Proceedings of the 23rd international conference on Machine learning*, 881–888.
- R. S. Sutton and A. G. Barto. 2018. *Reinforcement learning: An introduction*. MIT press, Cambridge, MA.
- P. Tadepalli and D. Ok. 1998. "Model-based average reward reinforcement learning." *Artificial intelligence*, 100, 1-2, 177–224.
- C. Tang, B. Abbatematteo, J. Hu, R. Chandra, R. Martín-Martín, and P. Stone. 2025. "Deep reinforcement learning for robotics: A survey of real-world successes." In: *Proceedings of the AAAI Conference on Artificial Intelligence 27*. Vol. 39, 28694–28698.
- R. Toro Icarte, T. Q. Klassen, R. Valenzano, M. P. Castro, E. Waldie, and S. A. McIlraith. 2023. "Learning reward machines: A study in partially observable reinforcement learning." *Artificial Intelligence*, 323, 103989. doi:<https://doi.org/10.1016/j.artint.2023.103989>.

- R. Toro Icarte, E. Waldie, T. Klassen, R. Valenzano, M. Castro, and S. McIlraith. 2019. “Learning reward machines for partially observable reinforcement learning.” *Advances in Neural Information Processing Systems*, 32, 15523–15534.
- M. Y. Vardi. 1985. “Automatic verification of probabilistic concurrent finite state programs.” In: *26th Annual Symposium on Foundations of Computer Science (SFCS 1985)*. IEEE, 327–338.
- A. Velasquez, B. Bissey, L. Barak, A. Beckus, I. Alkhoury, D. Melcer, and G. Atia. 2021. “Dynamic automaton-guided reward shaping for Monte Carlo tree search.” In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 35(13), 12015–12023.
- K. Wagner. Nov. 1979. “On  $\omega$ -regular sets.” *Information and Control*, 43, 2, (Nov. 1979), 123–177.
- Y. Wan, A. Naik, and R. S. Sutton. 2021. “Learning and planning in average-reward Markov decision processes.” In: *International Conference on Machine Learning*. PMLR, 10653–10662.
- Y. Wan and R. S. Sutton. 2022. “On convergence of average-reward off-policy control algorithms in weakly communicating MDPs.” In: *OPT 2022: Optimization for Machine Learning (NeurIPS 2022 Workshop)*.
- Y. Wang, A. Velasquez, G. Atia, A. Prater-Bennette, and S. Zou. 2024. “Robust average-reward reinforcement learning.” *Journal of Artificial Intelligence Research*, 80, 719–803.
- Y. Wang, A. Velasquez, G. K. Atia, A. Prater-Bennette, and S. Zou. 2023. “Model-free robust average-reward reinforcement learning.” In: *International Conference on Machine Learning*. PMLR, 36431–36469.
- C.-Y. Wei, M. J. Jahromi, H. Luo, H. Sharma, and R. Jain. 2020. “Model-free reinforcement learning in infinite-horizon average-reward Markov decision processes.” In: *International conference on machine learning*. PMLR, 10170–10180.
- Z. Xu, I. Gavran, Y. Ahmad, R. Majumdar, D. Neider, U. Topcu, and B. Wu. 2020. “Joint inference of reward machines and policies for reinforcement learning.” In: *Proceedings of the International Conference on Automated Planning and Scheduling*. Vol. 30, 590–598.
- Z. Xu, B. Wu, A. Ojha, D. Neider, and U. Topcu. 2021. “Active finite reward automaton inference and reinforcement learning using queries and counterexamples.” In: *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*. Springer, 115–135.
- H. Yoon, Y. Chou, X. Chen, E. Frew, and S. Sankaranarayanan. 2019. “Predictive runtime monitoring for linear stochastic systems and applications to geofence enforcement for UAVs.” In: *International Conference on Runtime Verification*. Springer, 349–367.
- W. Zhao, J. P. Queralta, and T. Westerlund. 2020. “Sim-to-real transfer in deep reinforcement learning for robotics: a survey.” In: *2020 IEEE symposium series on computational intelligence (SSCI)*. IEEE, 737–744.

Received 23 May 2025; accepted 21 January 2026