

Inferring Probabilistic Reward Machines from Non-Markovian Reward Signals for Reinforcement Learning

Taylor Dohmen^{1*}, Noah Topper^{2*}, George Atia², Andre Beckus³,
Ashutosh Trivedi¹, Alvaro Velasquez³

¹ University of Colorado Boulder

² University of Central Florida

³ Air Force Research Laboratory

Abstract

The success of reinforcement learning in typical settings is predicated on Markovian assumptions on the reward signal by which an agent learns optimal policies. In recent years, the use of reward machines has relaxed this assumption by enabling a structured representation of non-Markovian rewards. In particular, such representations can be used to augment the state space of the underlying decision process, thereby facilitating non-Markovian reinforcement learning. However, these reward machines cannot capture the semantics of stochastic reward signals. In this paper, we make progress on this front by introducing probabilistic reward machines (PRMs) as a representation of non-Markovian stochastic rewards. We present an algorithm to learn PRMs from the underlying decision process and prove results around its correctness and convergence.

Introduction

Traditionally, reinforcement learning (RL) (Sutton and Barto 1998) techniques have relied on strong Markovian assumptions of the underlying decision process. In particular, the reward signal by which an agent learns through positive or negative reinforcement is generally defined as a function of the current state of the environment and the action of an agent. The history of states observed by the agent are not considered in such settings. However, many problem domains require taking into account this history. Examples include learning in settings with sparse rewards (Neider et al. 2021), rewards defined as regular expressions and formal logics (Camacho et al. 2019), and decision-making under partial observability (Icarte et al. 2019). In such settings, the non-Markovian reward signal may not be known, but can be learned from traces of behavior. Recently, the problem of learning a structured representation from the reward signal has received significant attention. These representations often take the form of finite automata called reward machines (Icarte et al. 2018, 2022). Reward machines can augment the state space of decision processes with non-Markovian rewards and thereby facilitate the use of classical techniques for Markov Decision Process (MDP) on these more general types of environments (Gaon and Brafman 2020; Xu et al.

2021). They also serve as a memory mechanism for reasoning over partially observable environments (Icarte et al. 2019), are useful for reward shaping to mitigate sparse reward signals (Camacho et al. 2019; Velasquez and Melcer 2020; Velasquez et al. 2021), and provide explanations of RL systems (Xu et al. 2021).

While much progress has been made on learning and leveraging reward machines for decision processes with non-Markovian rewards (Xu et al. 2020, 2021; Abadi and Brafman 2020; Gaon and Brafman 2020; Neider et al. 2021; Rens et al. 2021), the more general setting where rewards exhibit both non-Markovian and stochastic dynamics has not been addressed. In this paper, we make progress on this front by introducing *probabilistic reward machines* (PRMs).

We present an algorithm to learn PRM representations of non-Markovian stochastic reward signals by integrating a probabilistic variant of the L^* algorithm (Angluin 1987) for automaton inference with RL. Our algorithm uses a product construction to produce a standard MDP from the environmental decision process and the learned PRMs in order to perform RL, and we provide an accompanying proof that these products correctly simulate the original decision processes with non-Markovian stochastic rewards. To further motivate PRMs, we show that an exponential blowup in size is required to emulate the system under learning by embedding probabilities in the environment and using deterministic reward machines. We establish theoretical guarantees on convergence of the algorithm in the limit.

Related Work

The classical problem of grammatical inference seeks to learn a formal language, or some representation thereof, from a finite number of samples (Horning 1969; de la Higuera 2010). Approaches for such learning problems are generally categorized as either active or passive. Passive inference seeks to mine the underlying specification from a static dataset of observed traces. Active methods differ in that the system under learning can be queried to guide the learning process. The L^* algorithm for learning regular languages (Angluin 1987) is the quintessential example of active inference, and assumes the existence of a minimally adequate teacher capable of answering membership and equivalence queries. This method has been broadly adopted and generalized to learn interface automata (Aarts and Vaan-

*These authors contributed equally.

drager 2010), Mealy machines (Niese 2003), automaton representations of recurrent neural networks (Weiss, Goldberg, and Yahav 2018, 2019), and MDPs (Tappler et al. 2021).

Grammatical inference has become a popular problem in the context of RL and learning reward signals. While most literature assumes a reward machine is given (Icarte et al. 2018; Hahn et al. 2019), the problem of learning the machine from observations has only been considered recently. The work in (Xu et al. 2020; Neider et al. 2021) explores a solving approach based on satisfiability. In (Icarte et al. 2019), the problem of learning a reward machine is viewed through the lens of discrete optimization. Methods that merge Q-learning and L^* for learning reward machines are considered in (Gaon and Brafman 2020; Xu et al. 2021).

Probabilistic Reward Machines

We proceed with our treatment on PRMs with some preliminary definitions pertaining to the underlying decision process from which samples are obtained to estimate the PRM representation of the objective.

A *Transition-Markov Decision Process* (TMDP) maintains the classical assumption that state transitions are Markovian while generalizing the standard MDP model by allowing history-dependent stochastic rewards.

Definition 1 (TMDP). A transition-Markov decision process M is a tuple (X, x_I, A, P, AP, L, R) where X is a finite set of states, $x_I \in X$ is an initial state, A is a finite set of actions, $P : (X \times A \times X) \rightarrow [0, 1]$ is a probabilistic transition function, AP is a finite set of atomic propositions, $L : (X \times A \times X) \rightarrow 2^{AP}$ is a function labelling each transition with a subset of AP , and $R : (2^{AP})^* \rightarrow \text{Dist}(\mathbb{R})$ is a non-Markovian stochastic reward function, where $\text{Dist}(\mathbb{R})$ denotes the set of possible probability distributions over \mathbb{R} .

A run of a TMDP is a string $x_0 a_1 x_1 \cdots a_n x_n \in X \cdot (A \cdot X)^*$ such that $0 < \prod_{k=1}^n P(x_{k-1}, a_k, x_k)$. For a run $x_0 a_1 x_1 \cdots a_n x_n$, its corresponding label sequence is $\ell_1 \ell_2 \cdots \ell_n$, where $\ell_k = L(x_{k-1}, a_k, x_k)$, and the corresponding reward sequence is $r_1 r_2 \cdots r_n$ where $r_k \sim R(\ell_1 \cdots \ell_k)$. A trajectory is a run for which $x_0 = x_I$. A policy is a function $\pi : (X \cdot (A \cdot X)^* \times A) \rightarrow [0, 1]$ specifying a distribution over the set of actions following a run.

Given a TMDP $M = (X, x_I, A, P, AP, L, R)$, we seek to construct a PRM that encodes the underlying reward R .

Definition 2 (PRM). A probabilistic reward machine H is a tuple $(AP, \Gamma, Y, y_I, \tau, \varrho)$ where AP is a set of atomic propositions and 2^{AP} is the input alphabet, $\Gamma \subseteq \mathbb{R}$ is a finite set of rewards, Y is a finite set of states, $y_I \in Y$ is a distinguished initial state, $\tau : (Y \times 2^{AP} \times Y) \rightarrow [0, 1]$ is a probabilistic transition function, and $\varrho : (Y \times 2^{AP} \times Y) \rightarrow \Gamma$ is a function mapping each transition to a reward from Γ .

A run of H is a sequence $y_0(\ell_1, \gamma_1)y_1 \cdots (\ell_n, \gamma_n)y_n$ where $\gamma_k = \varrho(y_{k-1}, \ell_k, y_k)$ and $\tau(y_{k-1}, \ell_k, y_k) > 0$, for all k . Let $\mathbf{H} : (2^{AP} \cdot \Gamma) \rightarrow [0, 1]^{|Y| \times |Y|}$ map each pair $\ell, \gamma \in 2^{AP} \times \Gamma$ to the transition matrix $\mathbf{H}(\ell\gamma)$ such that

$$\mathbf{H}(\ell\gamma)[i, j] = \begin{cases} \tau(y_i, \ell, y_j) & \text{if } \varrho(y_i, \ell, y_j) = \gamma, \\ 0 & \text{otherwise.} \end{cases}$$




Figure 1: A PRM defined over the office gridworld.

As usual, we extend the domain from $2^{AP} \cdot \Gamma$ to $(2^{AP} \cdot \Gamma)^*$ by setting $\mathbf{H}(\ell\gamma w) = \mathbf{H}(\ell\gamma)\mathbf{H}(w)$ and $\mathbf{H}(\varepsilon) = \mathbf{I}$, where \mathbf{I} is the identity matrix and ε the empty string. A distribution $\mathbb{P}_H : ((2^{AP})^* \times \Gamma^*) \rightarrow [0, 1]$ over pairs of inputs and reward sequences is induced:

$$\mathbb{P}_H(\ell_1 \cdots \ell_n, r_1 \cdots r_n) = \mathbf{y}_I \mathbf{H}(\ell_1 r_1 \cdots \ell_n r_n) \mathbf{1},$$

where \mathbf{y}_I is an initial distribution over Y (typically with an entry of 1 at the index for state y_I and zeros everywhere else) and $\mathbf{1}$ is the vector with 1 at each entry. The PRM H is said to encode the reward function R of a TMDP if, for any label sequence $\ell = \ell_1 \cdots \ell_n$ and reward sequence $r = r_1 \cdots r_n$, the following equality holds:

$$\mathbb{P}_H(\ell, r) = \prod_{k=1}^n R(\ell_1 \cdots \ell_k)(r_k).$$

Example 1. Consider the office gridworld environment found in (Icarte et al. 2018). The original task, which takes the form of a deterministic reward machine, requires that the agent acquire coffee at location c and delivers the coffee to office o , at which point the agent receives a reward of 1. If the agent steps on an $*$, they fail the task and observe a reward of 0. A variant of this task, in which the reward signal is probabilistic, may be obtained by introducing a 10% chance that the coffee machine malfunctions, producing weak coffee which is rejected upon delivery and receives a reward of 0. A graphical representation of a PRM for this task is shown in Figure 1. An edge $y \xrightarrow[\tau]{\ell p} y'$ indicates that y transitions to y' on label $\ell \in 2^{AP}$ with probability $p = \tau(y, \ell, y')$ and receives reward $r = \varrho(y, \ell, y')$.

The following property formalizes the notion of a PRM in which every state and label uniquely determines the reward.

Definition 3 (Reward-Determinism). A PRM is reward-deterministic if, for any $y, y_1, y_2 \in Y$ and $\ell \in 2^{AP}$ such that $\tau(y, \ell, y_1) > 0$ and $\tau(y, \ell, y_2) > 0$, it holds that either $y_1 = y_2$ or $\varrho(y, \ell, y_1) \neq \varrho(y, \ell, y_2)$.

An Excursus on L^*

In the seminal work of (Angluin 1987), an algorithm called L^* is proposed by which a deterministic finite automaton (DFA) for a given regular language \mathcal{L} can be learned through the use of a minimally adequate teacher that can answer membership and equivalence queries. When the learner executes a membership query, it presents a word $w \in (2^{\text{AP}})^*$ to the teacher and the teacher outputs whether $w \in \mathcal{L}$. If the learner executes an equivalence query, it presents a hypothesis DFA to the teacher who must then answer whether this automaton encodes the language to be learned. If not, the teacher generates a counterexample in the form of a word on which the two languages differ. The L^* algorithm keeps track of state words $S \subseteq (2^{\text{AP}})^*$, which are closed under prefix operations (e.g. if $ab \in S$, then $a \in S$) and test words $E \subseteq (2^{\text{AP}})^*$, which are closed under suffix operations (e.g. if $ab \in E$, then $b \in E$). Initially, we have $S = E = \{\varepsilon\}$, where ε is the empty string. As the algorithm proceeds, there are two critical properties that must be tracked and revolve around the notion of E -equivalence.

Definition 4 (E -Equivalence). *Given $w, w' \in (2^{\text{AP}})^*$ and set $E \subseteq (2^{\text{AP}})^*$, the words w and w' are E -equivalent with respect to the language \mathcal{L} , denoted $w \equiv_E w'$, if $w \in \mathcal{L} \iff w'e \in \mathcal{L}$ holds for every $e \in E$.*

Definition 5 (Consistency). *Given (S, E) , the consistency property holds if and only if for all $s, s' \in S$, the following implication holds: if $s \equiv_E s'$, then $s\ell \equiv_E s'\ell$ for all $\ell \in \Sigma$.*

Definition 6 (Closedness). *Given (S, E) , the closedness property holds if and only if for all $s \in S$ and $\ell \in 2^{\text{AP}}$, there exists some $s' \in S$ such that $s\ell \equiv_E s'$.*

For a closed and consistent (S, E) , a corresponding automaton can be derived by taking each E -equivalence class of S to be a state, with the empty string ε as the starting state. The transition function is defined using the closedness property. That is, whenever we have $s\ell \equiv_E s'\ell$ per the closedness property, then we also have the transition from s to s' upon observing the label ℓ in the DFA. Furthermore, it follows from the consistency property that s' is unique. The accepting states are those in the language of the teacher.

The L^* algorithm ensures that (S, E) remains closed and consistent until the target language is learned. This is done by constructing a DFA from (S, E) . If a counterexample is produced, it is used to modify (S, E) by adding the counterexample and all its prefixes to S . Membership queries are then used to determine the values of the new entries.

We now consider again the RL setting where we wish to learn the PRM representation of the objective as given by the reward signal, and mention some challenges that arise. Indeed, note that the seemingly innocuous membership queries must now be answered indirectly through the TMDP. This introduces two challenges to the adoption of the L^* algorithm within RL. First, there is no obvious way to perform membership queries. The construction of a meaningful w is not obvious on the part of the learner and answering this query is difficult on the part of the teacher. This is because we can only reason about these words indirectly through the TMDP, whereas traditional approaches reason directly over the target language.

The second challenge lies in performing equivalence queries, since the teacher cannot answer whether the currently learned PRM encodes the reward of the underlying TMDP. However, it is worth noting that counterexamples are obtained naturally through the RL process. These counterexamples are words that either have not been seen before or whose corresponding reward sequence differs from other encounters of the same word. The latter case arises due to the stochastic nature of the reward signal. Thus, the answers to equivalence queries are obtained implicitly through RL.

Explicating Probabilistic Machines from Non-Markovian Reward Signals

We now describe an algorithm that learns a PRM representation of the stochastic reward function of a TMDP.

Definition 7 (Sampling Observation Table). *For an underlying PRM $(\text{AP}, \Gamma, Y, y_I, \tau, \rho)$, an observation table is a tuple (S, E, T) , where $S \subseteq (2^{\text{AP}} \cdot \Gamma)^*$ is a prefix-closed set of samples, $E \subseteq (2^{\text{AP}} \cdot \Gamma)^* \cdot 2^{\text{AP}}$ is a suffix-closed set of experiments, and $T : S \cdot ((2^{\text{AP}} \cdot \Gamma) \cup \{\varepsilon\}) \cdot E \rightarrow (\Gamma \rightarrow \mathbb{N})$ is a function mapping label-reward words to reward frequencies.*

Defining appropriate analogs of closedness and consistency for this type of observation table relies on a notion of statistical difference, based on the Hoeffding bound.

Definition 8. *We say that the sequences $s, s' \in S \cdot 2^{\text{AP}}$ are statistically different, written $\text{diff}_f(s, s')$, with respect to a function $f : (2^{\text{AP}} \cdot \Gamma)^* \cdot 2^{\text{AP}} \rightarrow (\Gamma \rightarrow \mathbb{N})$ if either*

- $0 < N = \sum_{g \in \Gamma} f(s)(g)$ and $0 < N' = \sum_{g \in \Gamma} f(s')(g)$,
- or there exists $\gamma \in \Gamma$ such that we have the inequality $\sqrt{\frac{1}{2} \ln \frac{2}{C}} \left(\sqrt{\frac{1}{N}} + \sqrt{\frac{1}{N'}} \right) < \left| \frac{f(s)(\gamma)}{N} - \frac{f(s')(\gamma)}{N'} \right|$, where C is a data-dependent confidence level.¹

Two cells se and $s'e'$ of an observation table (S, E, T) are compatible, denoted $se \sim s'e'$, if $\text{diff}_T(se, s'e')$ is false.

We say that two rows s, s' are compatible, denoted $s \stackrel{E}{\sim} s'$ if $se \sim s'e$ holds for all $e \in E$. Based on these relations, S can be partitioned into *compatibility classes* which are groups of statistically similar samples. Unlike in traditional L^* , where each sample translates to a state in the derived automaton, each state of the derived PRM will correspond to a representative from a compatibility class. Each class is defined via compatibility to its representative. Note that the compatibility relation is not an equivalence relation; a given sample may be compatible with several representatives and two elements of a compatibility class may be incompatible. To resolve these ambiguities, define the function $\text{rank} : S \rightarrow \mathbb{N}$ as $\text{rank}(s) = \sum_{\ell \in 2^{\text{AP}}} \sum_{\gamma \in \Gamma} T(s\ell)(\gamma)$, to quantify the amount of information about each sample contained in the table. The unique representative of s is then defined as the compatible sample of maximal rank and is characterized by the function $\text{rep} : S \rightarrow S$, given as $\text{rep}(s) = \arg \max_{s' \in S} \{\text{rank}(s') : s \stackrel{E}{\sim} s'\}$. Denote by $\text{rep}(S)$ the set $\{\text{rep}(s) : s \in S\}$ of all representatives.

¹A good choice (Tappler et al. 2021) for C is $\frac{1}{M^3}$ where M is the number of samples taken thus far.

Algorithm 1: L^* -based Active Inference for PRM representation of Stochastic Non-Markovian Reward

Output: Observation table O and/or PRM H .

```
1 Initialize observation table  $O = (S, E, T)$ 
2 repeat
3   if exists counterexample  $\chi = \ell_1\gamma_1 \cdots \ell_n\gamma_n$  then
4      $Prefixes \leftarrow \{\ell_1\gamma_1 \cdots \ell_k\gamma_k : 1 \leq k \leq n\}$ 
5      $Queries \leftarrow Prefixes \cdot ((2^{AP} \cdot \Gamma) \cup \{\varepsilon\}) \cdot E$ 
6      $S \leftarrow S \cup Prefixes$ 
7   repeat
8     if  $O$  is not consistent then
9       Choose  $s, s' \in S, \ell \in 2^{AP}, \gamma \in \Gamma, e \in E$  such that  $s \stackrel{E}{\sim} s'$  and  $\text{diff}_T(sl\gamma e, s'l\gamma e)$ .
10       $E \leftarrow E \cup \{l\gamma e\}$ 
11       $Queries \leftarrow Queries \cup (S \cdot ((2^{AP} \cdot \Gamma) \cup \{\varepsilon\}) \cdot \{l\gamma e\})$ 
12     else if  $O$  is not closed then
13       Choose  $s \in S, \ell \in 2^{AP}, \gamma \in \Gamma$  such that  $T(sl)(\gamma) > 0$  and  $sl\gamma \stackrel{E}{\not\sim} r$  for all  $r \in \text{rep}(S)$ .
14        $S \leftarrow S \cup \{sl\gamma\}$ 
15        $Queries \leftarrow Queries \cup (\{sl\gamma\} \cdot ((2^{AP} \cdot \Gamma) \cup \{\varepsilon\}) \cdot E)$ 
16     for  $\zeta \in Queries$  do
17        $T \leftarrow \text{MEMBERSHIPQUERY}(\zeta, T)$ 
18      $Queries \leftarrow \emptyset$ 
19   until  $O$  is closed and consistent
20    $T, \chi \leftarrow \text{EQUIVALENCEQUERY}(O)$ 
21 until there exists no counterexample  $\chi$ 
```

Definition 9 (Sampling Closedness). *An observation table is closed if, for all $sl\gamma \in S \cdot 2^{AP} \cdot \Gamma$ with $T(sl)(\gamma) > 0$, there exists $r \in \text{rep}(S)$ such that $sl\gamma \stackrel{E}{\sim} r$ holds.*

Definition 10 (Sampling Consistency). *An observation table is consistent, if, for all compatible pairs in the set $\{(s, s') \in S \times S : s \stackrel{E}{\sim} s'\}$ and all $\ell, \gamma \in 2^{AP} \times \Gamma$, either $T(sl)(\gamma) = 0$ or $T(s'l)(\gamma) = 0$, or $sl\gamma \stackrel{E}{\sim} s'l\gamma$.*

With the above definitions established, we are now in position to describe the high-level flow of the learning procedure, which is specified explicitly in algorithm 1. While the details of the sub-procedures are significantly different from those involved in traditional L^* , the overall structure of the process is quite similar. The algorithm starts with an empty observation table and enters the loop at line 2. Since the table is empty to begin with and there does not exist a counterexample, the conditional at line 3 and the inner loop are bypassed and an equivalence query is executed at line 20. Deferring the details for now, the equivalence query builds a hypothesis PRM from the observation table and challenges the teacher to find a counterexample showing a discrepancy between the hypothesis and the true reward function. If no counterexample is found, the halting condition at line 21 is met and the algorithm terminates.

If a counterexample is found, it is returned and stored in the variable χ and we return to the beginning of the loop. In this case, the conditional block at line 3 is entered and the counterexample χ and all of its prefixes are added to the set of samples S of the observation table. Additionally, we would like to gain more information about the counterexam-

ple, so we save it and its extensions in the variable $Queries$.

Next, the loop at line 7 is entered, in which the table is checked for consistency. If the table is inconsistent (lines 8-11), a witness $l\gamma e$ is found, added to the set E of experiments, and the set of all samples extended by this witness is added to the set $Queries$. If the table is consistent, we proceed (lines 12-15) by checking if it is closed. If it is not closed, we similarly find a witness, this time $sl\gamma$, add it to the sample set S and add its extensions by the set of all experiments to the set $Queries$. For each string in $Queries$ a membership query is executed to gather more data into the table (lines 16-17). The variable $Queries$ is then emptied (line 17), and this sub-process (lines 7-19) is iterated until the observation table is both closed and consistent. At this point, a hypothesis PRM can be constructed again, and the entire process repeats until the equivalence query fails to produce a counterexample to the current hypothesis.

The reward-determinism property (c.f. definition 3) is inspired by the label-determinism property used for learning MDPs in (Tappler et al. 2021). Algorithm 1 always learns a reward-deterministic PRM, even if the true reward is specified as a non-reward-deterministic PRM.

Equivalence Queries

In order to construct a hypothesis PRM, we need a statistically significant amount of data about system trajectories to be stored in the observation table. We assume that a threshold parameter $N \in \mathbb{N}$ specifying the minimum number of samples necessary to estimate a transition probability is provided. We use \perp to denote a sink state to which transitions are diverted when the data is insufficient for estimating tran-

Algorithm 2: Equivalence Query

Data: N_{stop}

```

1 Function EQUIVALENCEQUERY( $O$ ):
2   Make hypothesis  $H$  from table  $O$ , initialize  $Q_H$ .
3    $flag \leftarrow false$ 
4   for  $1 \leq i \leq N_{stop}$  do
5      $\lambda, Q_H, flag \leftarrow QUERY(Q_H, H, equiv)$ 
6     for prefix  $\ell_1 r_1 \dots \ell_k r_k$  of  $\lambda$  do
7        $\lambda_k \leftarrow \ell_1 r_1 \dots \ell_k$ 
8        $T(\lambda_k)(r_k) \leftarrow T(\lambda_k)(r_k) + 1$ 
9     if  $flag$  then return  $T, \lambda$ 
10  return  $T, None$ 

```

sition probabilities otherwise.

Definition 11 (Hypothesis Reward Machine). *Given a closed and consistent observation table (S, E, T) , we construct a PRM $H = (AP, \Gamma, Y, y_I, \tau, \varrho)$ as follows:*

- $Y = (\text{rep}(S) \times \Gamma) \cup \{y_I, \perp\}$ with $y_I = (\varepsilon, 0)$,
- $\tau(y, \ell, y') = \begin{cases} 0 & \text{if } y = \perp \text{ and } y' \neq \perp, \\ 1 & \text{if } y = y' = \perp, \\ 1 & \text{if } \sum_{\Gamma} T(y\ell)(g) < N, \\ & \text{and } y' = \perp \\ & \text{if } \sum_{\Gamma} T(y\ell)(g) \geq N, \\ \frac{T(y\ell)(\gamma)}{\sum_{\Gamma} T(y\ell)(g)} & y' = (\text{rep}(r\ell\gamma), \gamma), \\ & \text{and } y = (r, \eta), \end{cases}$
- $\varrho(y, \ell, y') = \begin{cases} 0 & \text{if } y = \perp \text{ or } y' = \perp, \\ \gamma & \text{if } y' = (\text{rep}(r\ell\gamma), \gamma) \text{ and } y = (r, \eta). \end{cases}$

Equivalence queries are one of two interfaces between the learner and the teacher. In our setting, the teacher does not have complete knowledge of the system under learning, and so the equivalence query amounts to repeatedly executing the system and the hypothesis in tandem, trying to find inconsistencies between the two. Algorithm 2 specifies the procedure, which takes a closed and consistent observation table as its single parameter. It begins on line 2 by constructing a hypothesis PRM according to definition 11 and initializing the corresponding Q-table Q_H . The variable $flag$ is initially set to `false` and will be set to true if the teacher finds a counterexample, i.e. a label-reward string impossible under the hypothesis PRM, including those leading to the failure state \perp . Additionally, each call to the teacher (line 5) returns a label-reward string λ and an updated Q-table Q_H . Regardless of whether λ is a counterexample, its data is added to the observation table (lines 6-8). Next (lines 9-10), we check if $flag$ is set to `true`, indicating λ is a counterexample, and if so, the updated table is returned along with λ , ending the loop early. Otherwise, the loop repeats, and halts after N_{stop} iterations without encountering a counterexample. In this case, the updated table is returned and the “counterexample” returned is `None`.

Membership Queries

Membership queries are the second interface between learner and teacher, and are designed to gather data on

Algorithm 3: Membership Query

Data: N_{query}

```

1 Function MEMBERSHIPQUERY( $\zeta, T$ ):
2   Construct query machine  $H_{\zeta}$  and initialize  $Q_{\zeta}$ .
3   for  $1 \leq i \leq N_{query}$  do
4      $\lambda, Q_{\zeta} \leftarrow QUERY(Q_{\zeta}, H_{\zeta}, member)$ 
5     for prefix  $\ell_1 r_1 \dots \ell_k r_k$  of  $\lambda$  do
6        $\lambda_k \leftarrow \ell_1 r_1 \dots \ell_k$ 
7        $T(\lambda_k)(r_k) \leftarrow T(\lambda_k)(r_k) + 1$ 
8   return  $T$ 

```

strings that are witnesses to the observation table’s openness or inconsistency. Following (Xu et al. 2021), we perform a membership query for each such witness string using a pseudo-“reward machine”, defined below.

Definition 12 (Membership Query Machine). *Given a query string $\zeta = \ell_1 \gamma_1 \dots \gamma_{n-1} \ell_n$, the corresponding membership query machine is $H_{\zeta} = (AP, \Gamma, Y, y_I, \tau, \varrho)$ where*

- $Y = \{y_0, y_1, \dots, y_n\}$ with $y_I = y_0$,
- $\tau(y_k, \ell, y') = \begin{cases} 1 & \text{if } \ell = \ell_{k+1} \text{ and } y' = y_{k+1} \\ & \text{or if } \ell \neq \ell_{k+1} \text{ and } y' = y_k, \\ 0 & \text{otherwise,} \end{cases}$
- $\varrho(y_k, \ell, y') = \begin{cases} 1 & \text{if } \ell = \ell_{k+1} \text{ and } y' = y_{k+1}, \\ 0 & \text{otherwise.} \end{cases}$

This reward machine incentivizes exploration along the query trace. Algorithm 3 gives the membership query specification, which takes as parameters a query string ζ and the map T from the observation table. First (line 2), a membership PRM H_{ζ} is constructed for ζ , according to definition 12 and a corresponding Q-table Q_{ζ} is initialized. Next, the learner passes the membership PRM and the Q-table to the teacher (line 4) and is given back a label-reward string λ and the updated Q-table. As in equivalence queries, the data contained in λ is added to the observation table (lines 5-7), and the process repeats N_{query} times, where N_{query} is a contextual parameter, known a-priori. Finally, the modified observation table is returned (line 8).

RL-Based Teacher

The function `QUERY` used in both equivalence and membership queries belongs to the teacher and encapsulates the RL portion of the algorithm. It leverages the fact that the product of a TMDP and a PRM results in a standard MDP.

Definition 13. *Let $M = (X, x_I, A, P, AP, L, R)$ be a TMDP and let $H = (AP, \Gamma, Y, y_I, \tau, \varrho)$ be a PRM. The product $M \times H$ is an MDP $(X', x'_I, A, P', AP, L', R')$ with a Markovian reward function $R' : 2^{AP} \rightarrow \Gamma$ such that:*

- $X' = X \times Y$ with $x'_I = (x_I, y_I)$,
- $P'((x, y), a, (x', y')) = P(x, a, x') \cdot \tau(y, L(x, a, x'), y')$,
- $L'((x, y), a, (x', y')) = L(x, a, x')$, and
- $R'(L'((x, y), a, (x', y'))) = \varrho(y, L(x, a, x'), y')$.

Any convergent RL algorithm working on the product of the TMDP and membership reward machine, will converge to a policy generating the desired membership query

Algorithm 4: Generic RL-driven Query

Data: $N_{ep}, M = (X, x_I, A, P, AP, L, R), \epsilon, \beta, \alpha$

```
1 Function QUERY( $Q, H, type$ ):
2    $x, y, \lambda, flag \leftarrow x_I, y_I, \epsilon, false$ 
3   for  $1 \leq i \leq N_{ep}$  do
4      $a \leftarrow \text{EPSILONGREEDYACTION}(Q, x, y, \epsilon)$ 
5      $x', r \leftarrow \text{STEP}(M, x, a)$ 
6      $\lambda \leftarrow \lambda \cdot L(x, a, x') \cdot r$ 
7      $y', \gamma \leftarrow \text{STEP}(H, y, L(x, a, x'))$ 
8     if  $type = \text{member}$  then
9        $Q(y, x, a) \leftarrow (1 - \alpha)Q(y, x, a) + \alpha(\gamma + \beta \max_{a' \in A} Q(y', x', a'))$ 
10    else if  $type = \text{equiv}$  then
11       $Q(y, x, a) \leftarrow (1 - \alpha)Q(y, x, a) + \alpha(r + \beta \max_{a' \in A} Q(y', x', a'))$ 
12      if  $\lambda$  is a counterexample then
13         $flag \leftarrow true$ 
14     $x, y \leftarrow x', y'$ 
15  if  $type = \text{member}$  then
16    return  $\lambda, Q$ 
17  else if  $type = \text{equiv}$  then
18    return  $\lambda, Q, flag$ 
```

with optimal probability. The RL algorithm optimizes over the product of the TMDP and the membership reward machine, and continues to observe the reward signals from the PRM. In doing so, the policy that maximizes these rewards is guided towards answering the underlying membership query by finding a trajectory of states in the decision process whose induced trace of labels answers the membership query. Similarly, any convergent RL algorithm working on the product of the TMDP and hypothesis reward machine, will converge to an optimal policy with respect to the rewards consistent with the hypothesis reward machine. If the hypothesis correctly encodes the true reward function, we get an optimal policy for the environment.

The RL-driven query is given in algorithm 4, which takes a Q-table, a reward machine, and $type$ indicating whether the query is for membership or equivalence. Furthermore, the procedure assumes access to contextual parameters ϵ for choosing the next action (line 4), α and β as the learning rate and discount factor, respectively, used in the Q-table updates (lines 9 and 11), and N_{ep} as the episode length of each RL execution. Note that the teacher interacts with the environmental TMDP and the given PRM individually, but by doing so in tandem, it simulates the execution of their product. It treats both M and H as black-boxes in the sense that it may only observe the current state and outputs of each after executing them with particular inputs using the STEP function. It is also assumed that the environment M resets to its initial state after the QUERY function returns.

Correctness, Convergence, & Compactness

In this section we establish theoretical properties related to PRMs and the learning procedure presented above. In particular we show that the product construction of definition 13 preserves TMDP semantics, prove that algorithm 1 converges in the limit, and establish that TMDPs with PRM

encoded reward functions are more compact than TMDPs with deterministic reward machines when representing semantically equivalent systems.

Definition 14 (TMDP Semantics). *The semantics of a TMDP $M = (X, x_I, A, P, AP, L, R)$ is a function $\llbracket M \rrbracket : A^* \rightarrow \text{Dist}((2^{\text{AP}})^* \times \mathbb{R})$ such that*

$$\llbracket M \rrbracket(w)(\ell, r) = \begin{cases} 1 & \text{if } w = \epsilon \text{ and } \ell = \epsilon, \text{ and } r = 0 \\ \perp & \text{if } w, \ell \text{ is not observable} \\ p \cdot q & \text{if } w, \ell \text{ is observable with prob. } p \\ & \text{and } q = R(\ell)(r) \end{cases}$$

Define the language $\mathcal{L}(M)$ of the TMDP M as the set $\{(w, \ell) \in A^* \times (2^{\text{AP}})^* : \exists r \in \mathbb{R}. \llbracket M \rrbracket(w)(\ell, r) \neq \perp\}$.

Definition 15 (MDP Semantics). *The semantics of an MDP $M = (X, x_I, A, P, AP, L, R)$ is a function $\llbracket M \rrbracket : (A^* \times (2^{\text{AP}})^*) \rightarrow \text{Dist}(\mathbb{R})$ such that*

$$\llbracket M \rrbracket(w, \ell)(r) = \begin{cases} 1 & \text{if } w = \epsilon \text{ and } \ell = \epsilon \text{ and } r = 0 \\ \perp & \text{if } w, \ell \text{ is not observable} \\ p & \text{if } w, \ell \text{ is observable with prob. } p \\ & \text{and } \ell = \ell' l \text{ and } r = R(l) \end{cases}$$

Define the language $\mathcal{L}(M)$ of the MDP M as the set $\{(w, \ell) \in A^* \times (2^{\text{AP}})^* : \exists r \in \mathbb{R}. \llbracket M \rrbracket(w, \ell)(\gamma) \neq \perp\}$.

Two MDPs M, M' are equivalent iff, for all $w \in A^*$ and $\ell \in (2^{\text{AP}})^*$, it holds that $\llbracket M \rrbracket(w, \ell) = \llbracket M' \rrbracket(w, \ell)$.

Theorem 1. *If M is a TMDP with a reward encoded by a PRM H , then, for all $w \in A^*$, $\ell \in (2^{\text{AP}})^*$, and $\gamma \in \Gamma$:*

$$\llbracket M \rrbracket(w)(\ell, \gamma) = \llbracket M \times H \rrbracket(w, \ell)(\gamma).$$

Proof. Suppose that $M = (X, x_I, A, P, AP, L, R)$ is a TMDP, $H = (AP, \Gamma, Y, y_I, \tau, \varrho)$ is a PRM encoding R , and

$M \times H = (X', x'_I, A, P', AP, L', R')$ is their product constructed according to definition 13. If $\llbracket M \rrbracket(w)(\ell, \gamma) \neq \perp$, then, by definition 14, there exists $p, q \in [0, 1]$ such that $\llbracket M \rrbracket(w)(\ell, \gamma) = pq$ and w, ℓ is observable with probability p and $R(\ell)(\gamma) = q$. If H encodes the reward function of M , then it follows from the definition of PRM semantics that $\mathbb{P}_H(\ell, \Gamma^{|\ell|-1}\gamma) = R(\ell)(\gamma)$, where $\Gamma^{|\ell|-1}$ denotes all reward sequences length $|\ell| - 1$. Notice that $\mathbb{P}_H(s, \Gamma^n) = \sum_{r \in \Gamma^n} \mathbb{P}_H(s, r) = 1$ for any n and any $s \in (2^{AP})^*$. Suppose now that $\ell = \ell'l$ and $w = w'a$. Assuming H is reward-deterministic, there must be states $y, y' \in Y$ such that $\mathbb{P}_H(\ell, \Gamma^{|\ell|-1}\gamma) = R(\ell)(\gamma) = \tau(y, l, y') = q$ and $\varrho(y, l, y') = \gamma$. Moreover, the probability p of observing w, ℓ in $M \times H$, depends only on the transition probability function P of M whenever the reward probabilities are not in consideration. Therefore we have that $\Pr[(w, \ell) \in \mathcal{L}(M)]$ is equivalent to $\Pr[(w', \ell') \in \mathcal{L}(M \times H)] \cdot P(x, a, x')$ for some $x, x' \in X$. Since we have established that $R(\ell)(\gamma) = \tau(y, l, y')$, we may conclude that

$$\begin{aligned} \llbracket M \rrbracket(w)(\ell, \gamma) &= \Pr[(w, \ell) \in \mathcal{L}(M)] \cdot R(\ell)(\gamma) \\ &= \Pr[(w', \ell') \in \mathcal{L}(M \times H)] \cdot P(x, a, x') \cdot \tau(y, l, y') \\ &= \Pr[(w, \ell) \in \mathcal{L}(M \times H)] \\ &= \llbracket M \times H \rrbracket(w, \ell)(\gamma). \quad \square \end{aligned}$$

For an MDP M and a positive integer n , define $\mathcal{L}^{\leq n}(M)$ as the set $\bigcup_{k \leq n} \{(w, \ell) \in A^k \times (2^{AP})^k : (w, \ell) \in \mathcal{L}(M)\}$.

Lemma 1 ((Tappler et al. 2021)). *Two MDPs M and M' with n and at most n states, respectively, are equivalent iff $\llbracket M \rrbracket(w, \ell) = \llbracket M' \rrbracket(w, \ell)$, for all $(w, \ell) \in \mathcal{L}^{\leq n^2+1}(M)$.*

The following result follows from the convergence of L^* for label-deterministic MDPs, shown in (Tappler et al. 2021). In that work, the authors show that under uniformly randomized testing strategies, the sampling-based L^* algorithm converges almost surely in the limit to the MDP under learning. Our learning algorithm can be seen as a variant of sampling-based L^* in which sampling is implemented via reinforcement learning. By resetting the learning rate parameter to very low values at the start of each query and gradually increasing it after a sufficiently long initial period of exploration results in period of essentially random exploration and thus simulates a uniformly random testing strategy. This allows us to apply the convergence arguments from (Tappler et al. 2021) to our setting to assert that our learning procedure converges almost surely in the limit to the PRM (or an equivalent PRM) encoding the reward function of the TMDP under learning. In particular, lemma 1 provides a lower bound on the reinforcement learning episode lengths necessary to achieve convergence in the limit.

Theorem 2. *Fix a TMDP M , and suppose that there exists a PRM H , with k states, encoding the reward function of M . Let H_n be the hypothesis PRM passed to the n^{th} equivalence query, and suppose that the episode length is at least $k^2 + 1$.*

1. *If every possible $\ell \in (2^{AP})^*$ occurs in a pair $(w, \ell) \in \mathcal{L}(M)$, then $H_n \rightarrow H$ almost surely as $n \rightarrow \infty$.*
2. *Otherwise, $M \times H_n \rightarrow M \times H$ almost surely as $n \rightarrow \infty$.*

Theorem 2 tells us that algorithm 1 almost surely converges to a PRM H encoding the reward function of the TMDP M under learning, under the assumption that every possible label sequence is attainable from M . If this assumption does not hold, then we cannot claim that the true PRM encoding the reward function has been learned, since we cannot sample the behavior of the system on unobservable traces. On the other hand, we can say that the learned PRM is equivalent to H , *modulo the language of M* . Given sufficient interactions, the algorithm will almost surely converge to a PRM that is equivalent to H on those label sequences in the language of M . Therefore, optimal strategies for $M \times H$ can be learned from the product of M and the learned PRM.

The following result justifies the use of PRMs by establishing that PRMs can provide a more compact representation of systems than might otherwise be modeled by embedding probabilities capturing the stochasticity of a stochastic non-Markovian reward function directly into the environment and using a deterministic reward machine to encapsulate the non-Markovian aspects of the reward function.

Theorem 3. *Suppose that M is a TMDP with reward function encoded by a reward-deterministic PRM H . If M has m states and H has n states, then there exists a TMDP M' with mn states with a deterministic reward function encoded by deterministic reward machine H' with $O(2^n)$ states such that the product MDPs $M \times H$ and $M' \times H'$ are equivalent.*

Proof. Suppose that $M = (X, x_I, A, P, AP, L, R)$ is a TMDP, $H = (AP, \Gamma, Y, y_I, \tau, \varrho)$ is a PRM encoding R . We proceed to construct a deterministic PRM as follows:

1. We construct an NFA N from H such that rewards are moved into the states and a transition is included in N for every transition of positive probability in H .
2. A DFA D is produced using a subset construction on N .
3. A deterministic reward machine H' is obtained by interchanging deterministic transitions in D with transitions of probability 1 and bringing rewards back out onto the transitions while still keeping them in the states.

Note that reward-determinism of H ensures that the determinization step can be carried out without introducing sequences of rewards in H' that were unobtainable in H . The three steps above can be composed into a single step to obtain a deterministic PRM $H' = (AP, \Gamma, Y', y'_I, \tau', \varrho')$ where

- $Y' = 2^{Y \times \Gamma}$ with $y'_I = \{(y_I, 0)\}$,
- $\tau'(Y_1, l, Y_2) = 1$ iff $\tau(y, l, y') > 0$ and $\varrho(y, l, y') = \gamma'$, for all $(y, \gamma) \in Y_1$ and $(y', \gamma') \in Y_2$
- $\varrho'(Y_1, l, Y_2) = \gamma$ iff $Y_2 \subseteq \{(y, \gamma) : y \in Y\}$.

Now, we construct, as a modified product of H and M , a TMDP $M' = (X', x'_I, A, P', AP, L', R')$ which simulates the transition probabilities of H and has a deterministic reward function encoded by H' :

- $X' = X \times Y$ with $x'_I = (x_I, y_I)$,
- $P'((x, y), a, (x', y')) = P(x, a, x') \cdot \tau(y, L(x, a, x'), y')$,
- $L'((x, y), a, (x', y')) = L(x, a, x')$,
- $R'(\ell)(\gamma) = \mathbb{P}_{H'}(\ell, \Gamma^{|\ell|-1}\gamma)$.

Taking the product of M' and H' yields an MDP $M' \times H' = (X'', x''_I, A, P'', AP, L'', R'')$ such that

- $X'' = X' \times Y'$ with $x''_I = (x'_I, y'_I)$,
- $P''((x'_1, y'_1), a, (x'_2, y'_2)) = P'(x'_1, a, x'_2) \cdot \tau'(y'_1, L'(x'_1, a, x'_2), y'_2)$,
- $L''((x'_1, y'_1), a, (x'_2, y'_2)) = L'(x'_1, a, x'_2)$,
- $R''(L''((x'_1, y'_1), a, (x'_2, y'_2))) = \rho'(y'_1, L'(x'_1, a, x'_2), y'_2)$.

What remains to be shown is that the MDPs $M \times H$ and $M' \times H'$ are equivalent. We proceed by induction on (w, ℓ) .

Base case: $(w, \ell) = (\varepsilon, \varepsilon)$.

It follows immediately from definition 15 that $\llbracket M \times H \rrbracket(\varepsilon, \varepsilon)(\gamma) = \llbracket M' \times H' \rrbracket(\varepsilon, \varepsilon)(\gamma)$, for any $\gamma \in \Gamma$.

Inductive case: If $\llbracket M \times H \rrbracket(w', \ell') = \llbracket M' \times H' \rrbracket(w', \ell')$, for $w' \in A^*$ and $\ell' \in (2^{AP})^*$, then it follows, for all $a \in A$ and $l \in 2^{AP}$, that $\llbracket M \times H \rrbracket(w'a, \ell'l) = \llbracket M' \times H' \rrbracket(w'a, \ell'l)$.

Let $w \in A^*$, $\ell \in (2^{AP})^*$, and $\gamma \in \Gamma$, and suppose that $w = w'a$ and $\ell = \ell'l$. Then, the following derivation holds for some $x, x' \in X$, $y, y' \in Y$ and $Y_1, Y_2 \subseteq Y$.

$$\begin{aligned}
& \llbracket M' \times H' \rrbracket(w, \ell)(\gamma) \\
&= \Pr[(w, \ell) \in \mathcal{L}(M' \times H') \text{ and } \gamma = R''(l)] \\
&= \Pr[(w', \ell') \in \mathcal{L}(M' \times H')] \\
&\quad \cdot P''(((x, y), Y_1), a, ((x', y'), Y_2)) \\
&= \Pr[(w', \ell') \in \mathcal{L}(M' \times H')] \cdot P(x, a, x') \\
&\quad \cdot \tau(y, l, y') \cdot \tau'(Y_1, l, Y_2) \\
&= \Pr[(w', \ell') \in \mathcal{L}(M' \times H')] \cdot P(x, a, x') \cdot \tau(y, l, y') \\
&= \llbracket M' \times H' \rrbracket(w', \ell')(\gamma) \cdot P(x, a, x') \cdot \tau(y, l, y') \\
&= \llbracket M \times H \rrbracket(w', \ell')(\gamma) \cdot P(x, a, x') \cdot \tau(y, l, y') \\
&= \llbracket M \times H \rrbracket(w, \ell)(\gamma) \quad \square
\end{aligned}$$

Empirical Evaluation

While the key contribution of this paper is theoretical, we implemented the proposed algorithms in Python to evaluate their effectiveness via a number of small-sized benchmarks. Recall the office gridworld scenario and the corresponding PRM shown in fig. 1. Notice that this PRM is neither minimal nor reward-deterministic, yet our algorithm was able to learn the equivalent PRM, shown in fig. 2, which is both minimal and reward-deterministic. We experimented with other small-sized environments and PRMs, and showed that our algorithm can consistently learn PRMs in less than 1 minute on a typical personal laptop.

As expected, the run-time of the algorithm increases exponentially with the number of states in the concept PRM; however our experiments showed that the algorithm scales gracefully (linearly) with the number of transitions. When varying the size of the environment, however, we noticed that certain kinds of environments seemed much more challenging than others. Upon further inspection, this coincided with the distinction that some of the PRMs operate on the




Figure 2: The PRM learned from the office gridworld.

predicates over the set of actions while others on the predicates over the set of states (observations). The office gridworld example, for instance, is labelled by observations c or o , while the actions the agent correspond to the direction of the movement (n, e, w, s). Of course, the key difference when the PRMs operate on predicates over actions is that every label is observable from every state and hence the MDP can generate every word, resulting in efficient membership queries. Interestingly, increasing the size of the environment where labels are observable from every state had minimal effect on the performance of the learning. On the other hand, if increasing the size resulted in very low probabilities of observing certain labels, the algorithm struggled to learn correct PRMs. This difference is a direct consequence of our definition of statistical difference. The employed notion of difference only distinguishes traces that have some data recorded in the table; it does not differentiate traces if either has no data at all. Thus, the algorithm cannot tell if the lack of data is because traces have yet to be sampled or if traces are unobservable. This issue can be alleviated by introducing oracle queries capable of answering questions regarding the observation language of the environment.

Conclusion

We introduced *probabilistic reward machines* as an extension of reward machines to accommodate non-Markovian and stochastic dynamics. We developed a learning algorithm to obtain PRMs from underlying decision processes via a combination of model-free reinforcement learning and a probabilistic variant of the L^* active inference method. While this work establishes a theoretical foundation for the use of reward machines to learn non-Markovian stochastic rewards, further research is needed to assess the practicality of our approach. We believe that certain tweaks, such as allowing oracle queries on the environmental observation language or using alternative notions of statistical difference and compatibility, will improve performance in practice.

Acknowledgments

This research was supported in part by the Air Force Research Laboratory through the Information Directorate's Information Institute Contract Number FA8750-20-3-1003, FA8750-20-3-1004, SA10032021060389, and SA2020071003-V0237, and the Air Force Office of Scientific Research through Award 20RICOR012, and the National Science Foundation through CAREER Award CCF-1552497 and Award CCF-2106339.

References

- Aarts, F.; and Vaandrager, F. W. 2010. Learning I/O Automata. In *Concurrency Theory, 21th International Conference, CONCUR*, volume 6269 of *Lecture Notes in Computer Science*, 71–85. Springer.
- Abadi, E.; and Brafman, R. I. 2020. Learning and Solving Regular Decision Processes. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI*, 1948–1954. ijcai.org.
- Angluin, D. 1987. Learning Regular Sets from Queries and Counterexamples. *Inf. Comput.*, 75(2): 87–106.
- Camacho, A.; Icarte, R. T.; Klassen, T. Q.; Valenzano, R. A.; and McIlraith, S. A. 2019. LTL and Beyond: Formal Languages for Reward Function Specification in Reinforcement Learning. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI*, 6065–6073. ijcai.org.
- de la Higuera, C. 2010. *Grammatical Inference: Learning Automata and Grammars*. Cambridge University Press.
- Gaon, M.; and Brafman, R. I. 2020. Reinforcement Learning with Non-Markovian Rewards. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI*, 3980–3987. AAAI Press.
- Hahn, E. M.; Perez, M.; Schewe, S.; Somenzi, F.; Trivedi, A.; and Wojtczak, D. 2019. Omega-Regular Objectives in Model-Free Reinforcement Learning. In *Tools and Algorithms for the Construction and Analysis of Systems - 25th International Conference, TACAS*, volume 11427 of *Lecture Notes in Computer Science*, 395–412. Springer.
- Horning, J. J. 1969. A study of grammatical inference. Technical report, Stanford Univ Calif Dept Of Computer Science.
- Icarte, R. T.; Klassen, T. Q.; Valenzano, R. A.; and McIlraith, S. A. 2018. Using Reward Machines for High-Level Task Specification and Decomposition in Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning, ICML*, volume 80 of *Proceedings of Machine Learning Research*, 2112–2121. PMLR.
- Icarte, R. T.; Klassen, T. Q.; Valenzano, R. A.; and McIlraith, S. A. 2022. Reward Machines: Exploiting Reward Function Structure in Reinforcement Learning. *J. Artif. Intell. Res.*, 73: 173–208.
- Icarte, R. T.; Waldie, E.; Klassen, T. Q.; Valenzano, R. A.; Castro, M. P.; and McIlraith, S. A. 2019. Learning Reward Machines for Partially Observable Reinforcement Learning. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems, NeurIPS*, 15497–15508.
- Neider, D.; Gaglione, J.; Gavran, I.; Topcu, U.; Wu, B.; and Xu, Z. 2021. Advice-Guided Reinforcement Learning in a non-Markovian Environment. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI*, 9073–9080. AAAI Press.
- Niese, O. 2003. *An integrated approach to testing complex systems*. Ph.D. thesis, Technical University of Dortmund.
- Rens, G.; Raskin, J.; Reynouard, R.; and Marra, G. 2021. Online Learning of non-Markovian Reward Models. In *Proceedings of the 13th International Conference on Agents and Artificial Intelligence, ICAART, Volume 2, Online Streaming*, 74–86. SCITEPRESS.
- Sutton, R. S.; and Barto, A. G. 1998. *Reinforcement learning - an introduction*. Adaptive computation and machine learning. MIT Press. ISBN 978-0-262-19398-6.
- Tappler, M.; Aichernig, B. K.; Bacci, G.; Eichlseder, M.; and Larsen, K. G. 2021. L^{*}-based learning of Markov decision processes (extended version). In *Formal Aspects Comput.*, volume 33, 575–615.
- Velasquez, A.; Bissey, B.; Barak, L.; Beckus, A.; Alkhouri, I.; Melcer, D.; and Atia, G. K. 2021. Dynamic Automaton-Guided Reward Shaping for Monte Carlo Tree Search. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI*, 12015–12023. AAAI Press.
- Velasquez, A.; and Melcer, D. 2020. Verification-Guided Tree Search. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS*, 2026–2028. International Foundation for Autonomous Agents and Multiagent Systems.
- Weiss, G.; Goldberg, Y.; and Yahav, E. 2018. Extracting Automata from Recurrent Neural Networks Using Queries and Counterexamples. In *Proceedings of the 35th International Conference on Machine Learning, ICML*, volume 80 of *Proceedings of Machine Learning Research*, 5244–5253. PMLR.
- Weiss, G.; Goldberg, Y.; and Yahav, E. 2019. Learning Deterministic Weighted Automata with Queries and Counterexamples. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems, NeurIPS*, 8558–8569.
- Xu, Z.; Gavran, I.; Ahmad, Y.; Majumdar, R.; Neider, D.; Topcu, U.; and Wu, B. 2020. Joint Inference of Reward Machines and Policies for Reinforcement Learning. In *Proceedings of the Thirtieth International Conference on Automated Planning and Scheduling*, 590–598. AAAI Press.
- Xu, Z.; Wu, B.; Ojha, A.; Neider, D.; and Topcu, U. 2021. Active Finite Reward Automaton Inference and Reinforcement Learning Using Queries and Counterexamples. In *Machine Learning and Knowledge Extraction - 5th IFIP TC 5, TC 12, WG 8.4, WG 8.9, WG 12.9 International Cross-Domain Conference, CD-MAKE*, volume 12844 of *Lecture Notes in Computer Science*, 115–135. Springer.