# **Detecting Hidden Triggers: Mapping Non-Markov Reward Functions to Markov**

Gregory Hyde<sup>a,\*</sup> and Eugene Santos, Jr.<sup>a</sup>

<sup>a</sup>Thayer School of Engineering at Dartmouth College, Hanover, NH, USA ORCID (Gregory Hyde): https://orcid.org/0000-0001-5608-5224, ORCID (Eugene Santos, Jr.): https://orcid.org/0000-0003-2923-5113

Abstract. Many Reinforcement Learning algorithms assume a Markov reward function to guarantee optimality. However, not all reward functions are Markov. This paper proposes a framework for mapping non-Markov reward functions into equivalent Markov ones by learning specialized reward automata, Reward Machines. Unlike the general practice of learning Reward Machines, we do not require a set of high-level propositional symbols from which to learn. Rather, we learn hidden triggers, directly from data, that construct them. We demonstrate the importance of learning Reward Machines over their Deterministic Finite-State Automata counterparts given their ability to model reward dependencies. We formalize this distinction in our learning objective. Our mapping process is constructed as an Integer Linear Programming problem. We prove that our mappings form a suitable proxy for maximizing reward expectations. We empirically validate our approach by learning black-box, non-Markov reward functions in the Officeworld domain. Additionally, we demonstrate the effectiveness of learning reward dependencies in a new domain, Breakfastworld.

### 1 Introduction

Reinforcement Learning (RL) is traditionally formulated as a sequential decision-making task modeled under a Markov Decision Process (MDP) [4]. This implies that both the reward and transition functions are Markov, depending solely on the current state and action. However, in complex environments, these functions are not always known and cannot always be assumed to be Markov. In fact, it has been shown that there are behaviors that cannot be encoded by any Markov reward function under certain state representations [1]. To address non-Markov dynamics, researchers often extend the state representation to enforce the Markov property, though this is usually at the cost of overcomplicating the state representation or making it excessively sparse. For example, Mnih et al. [19] stacked subsequent frames from the Atari games [3] to encode a suitable representation for training Deep Q-Networks (DQNs). Similarly, Recurrent DQNs (RDQNs) have been used to encode state histories for similar purposes [14]. However, these sparse, deep models obscure the interpretation of rewards. To better understand them, it is crucial to examine their dynamics under less sparse and more interpretable forms.

To this end, automata have been used to compactly represent the hidden or historical features of non-Markov rewards. While Deterministic Finite-State Automata (DFAs) have traditionally served this role [5, 12, 11], more recently, specialized reward automata, Reward Machines (RMs), have been used [32, 22, 23, 10]. We highlight two distinctions between DFAs and RMs: 1.) DFAs are inherently less expressive than RMs as they only *accept* or *reject* traces. Consequently, when rewards are complex, multiple DFAs must be learned [12]. In contrast, RMs can encode complex reward behavior in a single automaton. 2.) DFAs are only capable of representing patterns over  $H = (S \times A)^*$ , the history of states and actions. Conversely, RMs can represent patterns over  $H = (S \times A \times \mathbb{R})^*$ , which includes rewards. This affords us a unique learning objective where rewards serve as memory allowing us to model their dependencies.

However, prior work on learning RMs assumes access to P, a set of high-level propositions, and  $L: S \times A \times S \rightarrow 2^P$ , a labeling function mapping transitions to instantiations in P. This makes much of this prior work ungeneralizable. We, instead, infer L by identifying minimal patterns in  $H = (S \times A \times \mathbb{R})^*$ , called *hidden triggers*, that differentiate reward outcomes. Rather than learn RMs directly, we map observed non-Markov rewards onto the cross-product of U, the state space of a candidate RM, and S, the observed state space, to form a Markov representation called the Abstract Reward Markov Decision Process (ARMDP). We infer an ARMDP, and thus its decomposable RM, by solving an Integer Linear Program (ILP), lending itself to powerful off-the-shelf discrete optimization solvers.

**Contributions:** We introduce a novel algorithm for mapping non-Markov reward functions to equivalent Markov ones. We show how by leveraging  $H = (S \times A \times \mathbb{R})^*$  in our learning objective we can dramatically expedite learning in cases with interdependent reward signals. We validate our approach by demonstrating our ability to learn black-box RMs from the Officeworld domain, originally introduced in the first RM paper [16]. We then evaluate the representative power of the ARMDP by extending DQN agents with ARMDP features to compare their learning profiles against various sparse RDQN variants. Finally, we validate our approach to represent complex reward function behavior in a new domain, Breakfastworld, featuring black-box RMs with interdependent reward signals. We show how by leveraging  $H = (S \times A \times \mathbb{R})^*$  in these experiments we can expedite learning. Our main contributions are three-fold:

- 1. We introduce a novel algorithm for learning RMs without assumed access to *P* and *L*,
- 2. We prove that the ARMDP is a suitable proxy for maximizing reward expectations under non-Markov rewards,
- We demonstrate the effectiveness of learning RMs under multiple interdependent reward signals.

<sup>\*</sup> Corresponding Author. Email: gregory.m.hyde.th@dartmouth.edu.

### 2 Related Work

Decision-making with non-Markov reward functions is often represented as a Non-Markov Reward Decision Process (NMRDP) [27]. An NMRDP is similar to an MDP except that the reward function,  $R: H \to \mathbb{R}$ , maps state and action histories,  $H = (S \times A)^*$ , to reward values. However, as H can become unwieldy due to the curse of history, more succinct representations have been encoded using automata. The intuition of this approach lies in the fact that automata are finitely represented and offer a level of interpretation due to their structured form and symbolic representation.

Traditionally, the automata of choice has been the DFA, yielding promising results in identifying Markov representations for NM-RDPs [5, 12, 11]. More recently, specialized reward automata, RMs, have also been used [16]. While DFAs operate over an input alphabet and emit *true* or *false* over traces, RMs extend this expressiveness by emitting reward values along each transition. RMs are particularly appealing as they offer succinct encoding of non-Markov reward behavior and task decomposition. They have also been extended by a broad range of RL algorithms that exploit RM structure to learn effective or optimal (in the discrete case) policies [7]. These include but are not limited to, specialized Q-Learning (QRM), Counterfactual experiences (CRM), Hierarchical Reinforcement Learning (HRM) and Reward Shaping (RSRM) for RMs [7, 17].

Whether using a DFA or an RM, various strategies for learning them exist. One prominent strategy for learning automata is to use Angluin's L\* algorithm [2]. L\* has been applied for both DFA [12] and RM [22, 23, 10] learning. Other strategies include Inductive Logic Answer Set Programming (ILASP) [11] for DFAs and SAT solvers [32] or Tabu search [29] for RMs. Non-automata-based learning of non-Markov rewards has been accomplished via Linear Temporal Logic (LTL), [18, 5, 6]. LTL can also be later translated into automata. Camacho et al. provide a helpful pipeline for transforming LTL formulae into RMs [7].

A notable limitation of these prior works is that they assume access to a set of high-level symbols (the input language for the automata) and the relationship of the underlying process to those symbols. In RM literature, these are denoted P and  $L: S \times A \times S \rightarrow 2^P$  where P is a set of symbols, separate from the NMRDP, and L is a mapping of transitions from the NMRDP to instantiations in P, respectively. A similar assumption is often made for DFA learning, with exceptions being the work of Gaon and Brafman [12] as well as Christoffersen et al. [9] who learn symbolic representations from  $H = (S \times A)^*$ to build their automata. Gaon and Brafman learned their DFAs by applying L\*, however, their approach was shown to be sample inefficient. Christoffersen et al. improved upon this by solving a discrete optimization problem for DFA learning, outlined by Shvo et al. [24], which regularized automata based on their size to enhance generalization. We improve upon these approaches by, instead, learning an RM without assumed access to P and L. We show how, by learning RMs, we can model reward dependencies in a single automaton by representing patterns in  $H = (S \times A \times \mathbb{R})^*$ .

### **3** Preliminaries

### 3.1 Reinforcement Learning and Non-Markov Reward

RL is a type of self-supervised learning in which an agent learns to interact with its environment to maximize its long-term reward expectation [25]. It is an iterative process in which the agent receives feedback in the form of rewards or punishments to adjust its behavior. RL is formalized using the MDP. An MDP is a 6-tuple,  $(S, A, T, R, \gamma, \rho)$ , where S is the state space, A is the action space,  $T: S \times A \times S \rightarrow [0, 1]$  is a Markov probabilistic transition function,  $R: S \times A \times S \rightarrow \mathbb{R}$  is a Markov reward function mapping states, actions and states to rewards,  $\gamma \in [0, 1)$  is the discount factor and  $\rho: S \rightarrow [0, 1]$  is the initial state distribution. Briefly, NMRDPs are similar to MDPs in that they share  $(S, A, T, R, \gamma, \rho)$ ; however,  $R: H \rightarrow \mathbb{R}$  maps histories of states and actions to rewards, instead.

The solution for an MDP is a policy,  $\pi : S \to A$ , that maps states to actions. The value of a policy for any state,  $s \in S$ , at time point t, is the expected discounted return, defined as follows:

$$V_{\pi}(s) = \mathbb{E}_{\pi} \left[ \sum_{k=0}^{\infty} \gamma^k r_{t+k+1} \mid s_t = s \right]$$
(1)

where  $\mathbb{E}_{\pi}$  is the expected reward from starting in *s* and following  $\pi$ . The optimal policy, denoted  $\pi_*$ , is determined by solving for the optimal state value function:

$$V_*(s) = \max V_\pi(s) \tag{2}$$

for all  $s \in S$ . We can write the optimal state-action value function in terms of the state value function as follows:

$$Q_*(s,a) = \mathbb{E}_{\pi} \left[ r_{t+1} + \gamma V_*(s_{t+1}) \mid s_t = s, a_t = a \right]$$
(3)

for all  $s \in S$  and  $a \in A$ . Hence:

$$V_*(s) = \max_{a} Q_*(s, a)$$
 (4)

Equation (4) is often solved through a process called value iteration [25]. The optimal policy is realized by selecting actions according to  $\max_a Q_*(s, a)$ , otherwise known as a greedy policy. A popular model-free approach for learning Equation (4) is Q-Learning [30]:

$$q(s_t, a_t) \xleftarrow{\alpha} [r_{t+1} + \gamma V(s_{t+1})] \tag{5}$$

where  $q(s_t, a_t)$  is updated slowly according to  $\alpha \in [0, 1)$ . When updating Equation (5) in an online fashion, typically an  $\epsilon$ -greedy strategy is employed where the agent follows a greedy policy but chooses a random action  $\epsilon$ -percent of the time. Both strategies are guaranteed to converge to the optimal policy in the limit as each state-action pair is visited infinitely often, but only under the assumption that T and R are Markov. Otherwise, expectations have no stable solution.

### 3.2 Reward Machines

RMs are a specialized type of automata imposed over an underlying state and action space, S and A [16]. RMs operate over a set of propositional symbols, P, representing high-level events, separate from S. RMs are defined as a 5-tuple,  $(U, u_1, F, \delta_u, \delta_r)$ , where U is a finite state space,  $u_1 \in U$  is the initial state, F is a set of terminal states where  $F \cap U = \emptyset$ ,  $\delta_u : U \times 2^P \to U \cup F$  is a transition function and  $\delta_r : U \times 2^P \times (U \cup F) \to \mathbb{R}$  is a reward function. A labeling function,  $L: S \times A \times S \rightarrow 2^P$ , connects S and A to the RM by mapping transitions to instantiations in P. RMs are capable of modeling reward dependencies based on histories that can be represented as a regular language (e.g., loops, conditions, and interleaving). RMs were initially demonstrated on the Officeworld domain [16] shown in Figure 1 (a). Tasks (b-e) represent RMs sensitive to  $P = \{c, m, o, *, A, B, C, D\}$  where c is a coffee tile, m, the mail tile, o, the office tile, \*, a plant tile, and A, B, C, D are room IDs.  $\neg$ ,  $\vee$  and  $\wedge$  are logical operators for *negation*, *or*, and *and*, respectively. The cross-product of an NMRDP and an RM can be used to form a Markov representation for predicting both R and T.



**Figure 1.** The Officeworld domain (a) with four Reward Machines, (b) deliver coffee to office, (c) deliver mail to office, (d) deliver coffee and mail to office and (e) patrol task sequencing A, B, C, and D.

### 4 Approach

We wish to map non-Markov reward signals into succinct Markov representations. We do so by inferring *minimal* RMs that represent hidden or historical features that render the representation Markov. To make our approach generalizable to practical decision-making scenarios, we drop the assumed access to P and L. For simplicity of notation, we drop notation on terminal states, F.

**Problem Statement** Given a finite set of finite-length trajectories,  $T_o$ , sampled from an NMRDP,  $(S, A, T, R, \gamma, \rho)$ , we learn a *minimal* RM,  $(U, u_1, \delta_u, \delta_r)$ , that renders rewards Markov to S, A and U.

While R in the NMRDP is unknown, we assume it can be sampled from and is representable as an RM. Hereafter, we will refer to R as a black-box RM. Our approach for solving this problem is to learn an ARMDP, which is simply the MDP cross-product between an RM and an NMRDP that enables the extraction of the associated RM. Learning of the ARMDP is formulated as an ILP. A *minimal* RM is an RM with the smallest hidden state space |U|, and the least number of transition edges between different RM states.

The following sections are organized as follows: We begin by defining the ARMDP as the cross-product space between an RM and an NMRDP in Section 4.1. Section 4.2 describes a passive learning approach for learning ARMDPs from  $T_o$ , by solving an ILP. We later make this an active learning approach in Section 4.5 by incorporating a Q-Learning agent. Section 4.3 briefly outlines how to extract an RM from an ARMDP. Finally, Section 4.4 describes the theoretical insights of ARMDPs as a proxy for the NMRDP.

### 4.1 Abstract Reward Markov Decision Process

An ARMDP is an MDP with additional properties sufficient for extracting/learning a valid RM. **Definition 1.** Given an NMRDP,  $(S, A, T, R, \gamma, \rho)$ , RM,  $(U, u_1, \delta_u, \delta_r)$ , and labeler,  $L : S \times A \times S \rightarrow 2^P$ , an AR-MDP is an MDP,  $(\tilde{S}, A, \tilde{T}, \tilde{R}, \gamma, \tilde{\rho})$ , satisfying:

1. 
$$\tilde{S} = S \times U$$
  
2.  $\{u \mid (s, u) \in \tilde{S}, \tilde{\rho}((s, u)) > 0\} = \{u_1\}$   
3.

$$\tilde{T}((s,u),a,(s',u')) = \begin{cases} T(s,a,s') & \text{if } \delta_u(u,L(s,a,s')) = u' \\ 0 & \text{otherwise} \end{cases}$$

4. 
$$\tilde{R}((s,u), a, (s', u')) = \delta_r(u, L(s, a, s'), u')$$

Property 1 defines the state space of the ARMDP as the crossproduct of S (what we observe) and U (what we do not observe). Property 2 imposes an initial state distribution compliant with  $u_1 \in U$  being the initial RM state. Property 3 defines the transition function of the ARMDP as a composition function of T and  $\delta_u$ . Finally, Property 4 defines the reward function of the ARMDP as  $\delta_r$ .

The intuition behind the ARMDP is that if we assume rewards in  $T_o$  are Markov to S and A, conflicts are bound to arise. Consider  $\delta_r : U \times 2^P \times U \to \mathbb{R}$ , for some black-box RM, being rewritten as  $\delta_r : U \times S \times A \times S \times U \to \mathbb{R}$  by invoking L. Because we do not observe U, we lack the specificity to predict  $\delta_r$ . Our goal is to find some candidate ARMDP, that renders rewards Markov to S, A, and U, by identifying *hidden triggers* in  $H = (S \times A \times \mathbb{R})^*$  that resolve all reward conflicts in  $T_o$ .

### 4.2 Formulation

We infer candidate ARMDPs by solving an ILP that resolves reward conflicts over  $T_o$ . Let  $\tau_m \in T_o$  be the m-th trace in  $T_o$  of length  $l_m$ . Further, let  $\tau_{m,n} = (s_n^m, a_n^m, r_{n+1}^m, s_{n+1}^m)$  be the n-th tuple of  $\tau_m$ . We define a reward conflict to be any two  $\tau_{m,n}$  and  $\tau_{\tilde{m},\tilde{n}}$  where  $(s_n^m, a_n^m, s_{n+1}^m) = (\tilde{s}_{\tilde{n}}^n, a_{\tilde{n}}^{\tilde{n}}, \tilde{s}_{\tilde{n}+1}^n)$ , but  $r_{n+1}^m \neq r_{\tilde{n}+1}^{\tilde{n}}$ . The motivation of our ILP is to map  $T_o$  onto a valid ARMDP such that we resolve all reward conflicts observed in  $T_o$ . Intuitively,  $\tilde{S}$  provides a degree of freedom over U to do just that.

We define the matrix,  $O = M \times N \times K \times K$ , where  $M = |T_o|$ , N reflects the largest  $l_m$ ,  $\forall \tau_m \in T_o$ , and K = |U| reflects the size of the finite state space for some candidate RM. Each  $O_{m,n}$ , then, is a  $K \times K$  matrix that lays the foundation for *Property l* in the ARMDP:

$$O_{m,n} = \begin{pmatrix} O_{m,n,1,1} & O_{m,n,1,2} & \cdots & O_{m,n,1,K} \\ O_{m,n,2,1} & O_{m,n,2,2} & \cdots & O_{m,n,2,K} \\ \vdots & \vdots & \ddots & \vdots \\ O_{m,n,K,1} & O_{m,n,K,2} & \cdots & O_{m,n,K,K} \end{pmatrix}$$

Each  $O_{m,n,i,j} \in O_{m,n}$  is a binary variable reflecting the decision to map  $s_n^m$  to RM state  $u_i \in U$  and  $s_{n+1}^m$  to RM state  $u_j \in U$ . That is, if  $O_{m,n,i,j} = 1$ , then  $\tau_{m,n} = (s_n^m, a_n^m, r_{n+1}^m, s_{n+1}^m)$ is mapped onto the corresponding ARMDP tuple of the form,  $((s_n^m, u_i), a_n^m, r_{n+1}^m(s_{n+1}^m, u_j))$ . As such, each  $\tau_{m,n}$  introduces  $K^2$  variables through its respective  $O_{m,n}$  matrix. We constrain each  $O_{m,n}$  to choose a single mapping by enforcing that  $\forall m =$  $1, 2, ..., M, \forall n = 1, 2, ..., l_m$ :

$$\sum_{i=1}^{K} \sum_{j=1}^{K} O_{m,n,i,j} = 1$$
(6)

We further constrain  $O_{m,n}$  where n = 1 to satisfy the initial state requirement of *Property 2* in the ARMDP. That is,  $\forall m = 1, 2, ..., M$ : 1360

$$\sum_{j=1}^{K} O_{m,1,1,j} = 1 \tag{7}$$

Next, we ensure the continuity of each trajectory in the mapping process. That is, we encode agreement between subsequent mappings of all  $O_{m,n}$  and  $O_{m,n+1}$  pairs by enforcing the following constraint  $\forall m = 1, 2, ..., M, \forall n = 1, 2, ..., l_m - 1, \forall j = 1, 2, ..., K$ :

$$\sum_{i=1}^{K} O_{m,n,i,j} = \sum_{j'=1}^{K} O_{m,n+1,j,j'}$$
(8)

Equation (8) simply forces  $s_{n+1}^m$  to be mapped to the same RM state in both  $\tau_{m,n}$  and  $\tau_{m,n+1}$ . Equations (6-8) form the scaffolding of our mapping.

Thus far, all constraints have been applied over individual or subsequent tuples of  $T_o$ . However, it is necessary to form a consensus on the mapping for all observed traces onto the ARMDP to ensure *Property 3* and *Property 4*. We begin with *Property 3*.

Property 3 encodes a transition function,  $\tilde{T}$ , constructed over components T, from the NMRDP, and  $\delta_u$ , from the RM. Even though  $\delta_u$  is unknown, we know that it must be deterministic. That is, for any transition,  $(s, a, s') \in S \times A \times S$ , and for any  $u_i \in U$ ,  $\delta_u(u_i, L(s, a, s'))$  must map to a single  $u_j \in U$ . Therefore, we must constrain O to agree on the assignments of  $u_i$  and  $u_j$  over shared transition tuples. We construct the following indicator variable:

$$I_{s,a,s',i,j} = \left[\sum_{m=1}^{M} \sum_{\substack{n=1, \\ (s,a,s')=(s_{n}^{m},a_{n}^{m},s_{n+1}^{m})}}^{l_{m}} O_{m,n,i,j} \ge 1\right]$$

 $I_{s,a,s',i,j}$  is an indicator variable representing the sum of all  $O_{m,n,i,j}$ in O sharing the transition tuple (s, a, s') with assignments  $i, j \in$ 1, 2, ..., K. When  $I_{s,a,s',i,j} = 1$ , it acts as a *hidden trigger* that highlights the sensitivity of L to (s, a, s') and enforces the transition  $\delta_u(u_j, L(s, a, s')) = u_j$ , in the RM, thus governing  $\tilde{T}$ . We later consider the sensitivity of *hidden triggers* to rewards in *Property 4*.

Let  $(S, A, S)_o = \{(s_n^m, a_n^m, s_{n+1}^m) \mid m \in \{1, 2, ..., M\}, n \in \{1, 2, ..., l_m\}\}$  be the set of all unique transitions in  $T_o$ . Then to enforce the determinism of  $\delta_u$ , we impose that  $\forall (s, a, s') \in (S, A, S)_o, \forall i = 1, 2, ..., K$ :

$$\sum_{j=1}^{K} I_{s,a,s',i,j} \le 1 \tag{9}$$

Consider that for any  $(s, a, s') \in (S, A, S)_o$  and for any  $i \in \{1, 2, ..., K\}$  that when we constrain the sum  $\forall j = 1, 2, ..., K$ , in Equation (9), to be less than or equal to 1, we limit  $\delta_u(u_i, L(s, a, s'))$  to have at most one outcome. With this, Equation (9) satisfies the determinism of  $\delta_u$  and *Property 3* of the ARMDP.

The encoding of *Property 4* enforces the resolution of reward conflicts in  $T_o$ . While conflicts exist under the perspective of  $S \times A \times S$ , by mapping conflicting reward tuples onto  $S \times U \times A \times S \times U$  we can separate them. Similar to our encoding of *Property 3*, we do this by constraining O to agree on the assignments of  $u_i$  and  $u_j$  over shared reward tuples by way of indicator variables:

$$I_{s,a,r,s',i,j} = \left[ \sum_{m=1}^{M} \sum_{\substack{n=1, \\ (s,a,r,s') = \tau_{m,n}}}^{l_m} O_{m,n,i,j} \ge 1 \right]$$

 $I_{s,a,r,s',i,j}$  is an indicator variable representing the sum over all  $O_{m,n,i,j}$  in O sharing the reward tuple (s, a, r, s') with assignments  $i, j \in \{1, 2, ..., K\}$ . When  $I_{s,a,r,s',i,j} = 1$ , it imposes the following reward emission,  $\delta_r(u_i, L(s, a, s'), u_j) = r$ , in the RM.

Let  $r(s, a, s') = \{r_{n+1}^m \mid m \in \{1, 2, ..., M\}, n \in \{1, 2, ..., l_m\}, (s_n^m, a_n^m, s_{n+1}^m) = (s, a, s')\}$  be the set of all observed rewards for some (s, a, s'). Then to enforce the determinism of  $\delta_r$ , we impose that  $\forall (s, a, s') \in (S, A, S)_o, \forall i = 1, 2, ..., K, \forall j = 1, 2, ..., K$ :

$$\sum_{\in r(s,a,s')} I_{s,a,r,s',i,j} \le 1 \tag{10}$$

Consider that for any  $(s, a, s') \in (S, A, S)_o$  and for any  $i, j \in \{1, 2, ..., K\}$  that when we constrain the sum  $\forall r \in r(s, a, s')$ , in Equation (10), to be less than or equal to 1, we limit  $\delta_r(u_i, L(s, a, s'), u_j)$  to have at most one reward outcome, thus resolving our reward conflicts in  $T_o$ . Importantly,  $I_{s,a,r,s',i,j} = 1$ contextualizes its corresponding *hidden trigger*,  $I_{s,a,s',i,j}$ , with reward. When rewards are interdependent (e.g., some upstream reward conflict determines a downstream reward conflict), rewards serve as memory, hence, we capture relationships in  $H = (S \times A \times \mathbb{R})^*$ . By Equation (10), we satisfy the determinism of  $\delta_r$  and *Property 4* of the ARMDP.

Equations (6-10) form the entirety of the constraints for ARMDP learning. However, there are potentially many ARMDP solutions for any given  $T_o$ . As such, we are motivated to derive simple and interpretable models for representing reward dynamics. Therefore, we orient our objective and learning procedure toward finding the *minimal* ARMDP, and thus the *minimal* RM, that satisfies  $T_o$ . This is done in two ways. First, for any inference of an ARMDP, we start by assuming K = |U| = 2, thus dramatically limiting our search space. If K = 2 is determined to be insufficient to resolve reward conflicts (ILP infeasibility) we increment K by 1 and try again. This is known as a *deepening* strategy employed in various related works [9, 6]. Second, we limit the number of *hidden triggers* that dictate transitions between different RM states (i.e.,  $\delta_u(u_i, L(s, a, s')) = u_j$ where  $u_i \neq u_j$ ). We do so by minimizing the following objective:

$$z = \sum_{(s,a,s)\in(S,A,S)_o} \sum_{i=1}^{K} \sum_{\substack{j=1,\\i\neq j}}^{K} I_{s,a,s',i,j}$$
(11)

We solve Equations (6-11) using out-of-the-box ILP solvers in Gurobi [13] to infer viable ARMDPs. We note that all notation in this section was constructed assuming  $L: S \times A \times S \rightarrow 2^P$  though it is straightforward to rewrite it as being sensitive to  $L: S \rightarrow 2^P$ . We provide the full ILP in Appendix A<sup>1</sup>.

We also note a structural assumption that we make in later experimentation. That is, we only use  $O_{m,n,i,j}$  where  $j \ge i$  (i.e., the upper triangle matrix of each  $O_{m,n}$ ). The implication of this is that we assume the RM to never return to a previous state, once left. While this might seem to limit the expressiveness of any resulting ARMDP, we note that our ILP simply unrolls these cycles expanding U as needed. Because  $T_o$  represents a finite set of finite length trajectories, learned ARMDPs will be equivalently expressive, but larger in some cases. For all RMs considered in this work, this structural assumption holds resulting in the *minimal* ARMDP being learned while dramatically reducing learning time. We provide more details on this assumption in Appendix B.

<sup>&</sup>lt;sup>1</sup> All Appendix items and code can be found in the full paper [15].

### Extracting RMs from ARMDPs 4.3

We provide a brief outline for extracting an RM from an ARMDP. Given that K = |U| is determined by our *deepening* strategy, we assume U is given. Additionally, by Equation (7) we designate  $u_1 \in$ U to be the initial state. To extract the rules of  $\delta_u$  and  $\delta_r$  we simply iterate  $\forall m = 1, 2, ..., M, \forall n = 1, 2, ..., l_m, \forall i = 1, 2, ..., K, \forall j = 1, 2, ..., K$ 1, 2, ..., K. If  $O_{m,n,i,j} = 1$  then  $\delta_u(u_i, L(s_n^m, a_n^m, s_{n+1}^m)) = u_j$ and  $\delta_r(u_i, L(s_n^m, a_n^m, s_{n+1}^m), u_j) = r_{n+1}^m$ .

#### 4.4 Theoretical Considerations

Our ILP is constructed such that it infers an ARMDP that resolves reward conflicts in  $T_{o}$ . Because the ARMDP is simply just an MDP, we can apply out-of-the-box RL algorithms to solve it. However, it remains to be shown how the ARMDP might be a suitable proxy for the NMRDP for the purpose of maximizing reward expectations. We start with the weighted Reward Sum (RS) for some  $\tau_m \in T_o$ :

$$RS(\tau_m) = \sum_{n=1}^{l_m} r_{n+1}^m \cdot T(s_n^m, a_n^m, s_{n+1}^m)$$
(12)

Given that we assume the form of R to be an RM and that  $\delta_u$  and  $\delta_r$ , in an RM, are deterministic functions,  $RS(\tau_m)$  reliably captures the reward observations essential for constructing a reward expectation, assuming the trajectory sequence is preserved. Similarly, we construct a weighted Abstracted Reward Sum (ARS) for some  $\tau_m \in T_o$ under the representation of the ARMDP to be:

$$ARS(\tau_m) = \sum_{n=1}^{l_m} \sum_{i=1}^{K} \sum_{j=1}^{K} r_{n+1}^m \cdot T(s_n^m, a_n^m, s_{n+1}^m) \cdot O_{m,n,i,j}$$
(13)

We remind readers that every  $O_{m,n,i,j}$  reflects a binary variable that, if 1, indicates  $\delta_u(u_i, L(s_n^m, a_n^m, s_{n+1}^m)) = u_j$ . By referencing *Prop*erty 3 of the ARMDP, multiplying  $O_{m,n,i,j}$  into  $T(s_n^m, a_n^m, s_{n+1}^m)$ to determine  $\tilde{T}$  should be evident. We introduce Theorem 1:

**Theorem 1.**  $RS(\tau_m) = ARS(\tau_m) \ \forall \tau_m \in T_o.$ 

For space reasons, we provide the proof of Theorem 1 in Appendix C, but briefly state that by ensuring  $RS(\tau_m) = ARS(\tau_m)$  $\forall \tau_m \in T_o$ , the ARMDP is a suitable proxy for maximizing reward expectations for the NMRDP, assuming  $T_o$  is representative. Furthermore, even if  $T_o$  is not representative, we posit that any ARMDP derived from  $T_o$  serves as a useful hypothesis for RL agents to interrogate in an active learning paradigm. We explore that intuition in the next section.

#### 4.5 RL for Non-Markov Rewards

Our ILP formulation for learning ARMDPs is a form of passive learning. That is to say, while the ARMDP will sufficiently resolve reward conflicts found in  $T_o$ ,  $T_o$  may not be representative of the true reward dynamics. Therefore, we are interested in extending our approach into an active learning framework by applying Q-Learning over inferred ARMDPs. Our intuition is that while any instance of  $T_o$ , or any ARMDP inferred from it, may not be truly representative, Q-Learning agents might gradually expand  $T_o$  in their pursuit to maximize reward expectations. In this way, each inferred ARMDP acts as a hypothesis that can be interrogated for consistency.

We present ARMDPQ-Learning, with pseudocode available in Algorithm 1. The agent applies a *deepening* strategy by initially assuming S and A form an MDP (i.e., |U| = 1). We denote this ARMDP

### Algorithm 1 ARMDPQ-Learning

19:

end if

20: end while

Input: NMRDP 1:  $ARMDP = ARMDP_0$ 2:  $\pi_{ARMDP} = \pi_0$ 3:  $T_o = []$ 4: K = |U| = 25: while !done do  $\tau_m$ , conflict = NMRDP.sim(ARMDP,  $\pi_{ARMDP}$ ) 6: 7: if conflict then 8:  $T_{o}.append(\tau_{m})$ 9: solved = False10: while !solved do 11:  $ARMDP, solved = solve\_ILP(T_o, K)$ 12: if !solved then 13: K = K + 1end if 14: 15: end while 16:  $\pi_{ARMDP} = \pi_0$ 17: else 18:  $\pi_{ARMDP} = update_{qsa}(\tau_m)$ 

as ARMDP<sub>0</sub> (line 1). K = |U| = 2, is used for subsequent ARMDP inference. The agent policy,  $\pi_{ABMDP}$ , is initially set to  $\pi_0$ , a policy with all Q-values set to 0 (line 2). We use the NMRDP to simulate some  $\tau_m$  under the perspective of the current ARMDP, selecting actions according to  $\pi_{ARMDP}$  (line 6). If a conflict is found we add  $\tau_m$ into  $T_o$  (line 8) and solve for a new ARMDP (line 11). By adding only conflicting  $\tau_m$ , we maintain a smaller  $T_o$  for more efficient ILP solving. If the problem is infeasible we increment K by 1 (line 13) until there is a solution. As each ARMDP presents a new perspective, we reset  $\pi_{ARMDP}$  back to  $\pi_0$  (line 16). Assuming no conflict is found for  $\tau_m$ , we update  $\pi_{ARMDP}$  with  $\tau_m$  according to Q-Learning updating (line 18). Note that line 18 is shorthand. In reality, tuples are updated in an online fashion while simulating the trajectory.

#### 5 **Experiments**

We briefly outline our experiments. Experiment 1 tests the efficacy of ARMDPQ-Learning in the Officeworld domain. Experiment 2 assesses the representative power of ARMDPs by comparing learning profiles of DQNs using S versus various RDQN architectures only using S. Finally, Experiment 3 demonstrates the efficacy of learning RMs over DFAs given their ability to learn over  $H = (S \times A \times \mathbb{R})^*$ .

#### 5.1 Experiment 1 - Officeworld

We validated our approach by applying ARMDPQ-Learning on the four black-box RMs (Tasks (b-e)) from the Officeworld domain visualized in Figure 1. We hid the RM as well as P and L from the agent. We refer readers to Appendix D for details on hyper-parameter selection, but note that we used an  $\epsilon$ -greedy policy and set  $\epsilon$  to 0% after 90,000 episodes to exploit learned behavior. Results for Experiment 1 can be seen in Figure 2.

The agent learned to perform optimally in all tasks, suggesting that a meaningful representation was learned. Task (b), (c), and (e) all arrived at |U| reflecting the ground truth model (2, 2, and 4, respectively). Conversely, Task (d) arrived at  $|U| \approx 3.2$ . Because the coffee and the mail tiles can be achieved in any order, occasionally



**Figure 2.** Average rewards with standard deviations, evaluated over 10 trials using ARMDPQ-Learning on Tasks (b-e) in the Officeworld domain.

the agent would lock on to one solution, and not the other, requiring less RM states.  $T_o$  remained relatively small for all tasks, with  $|T_o| \approx 4, 8, 39$ , and 44 for Tasks (b-e), respectively. We also tracked the cumulative solve times for all ILP solves under each task to be  $\approx 0.1, 0.6, 4,200$ , and 12,000 seconds, respectively. We include an extension to Figure 2 in Appendix D, containing figures for |U|,  $|T_o|$  and cumulative solve times.

In all experiments, we noticed high variability in early performance reflected in high standard deviations between episodes 0-20,000. We attribute this to identifying the wrong pattern in H during early conflict resolution. Interestingly, because these early solutions were so poor, they were easily testable by the Q-Learning agent. This lends some credence to our initial intuition of the Q-Learning agent serving as an important active learner capable of interrogating its own perspectives and updating in the face of new evidence.

## 5.2 Experiment 2 - ARMDP Representations

Previously, Christoffersen et al. compared the learning rates of Q-Learning agents over NMRDPs, using DFAs, to sparse, recurrent, deep RL models [9]. However, in smaller discrete state spaces, Q-Learning has the advantage of updating state-action values directly given their simpler, tabular representation. In contrast, deep RL models must learn approximate solutions using complex computations over large batches of data. It remained to be shown how beneficial it could be to represent the state space of a DQN with  $\tilde{S}$  instead of S. This would allow for a more even comparison of training profiles against RDQN architectures that have to learn a Markov representation over histories of S.

For each Task (b-e), in Experiment 1, we extracted S, from the finalized ARMDP. We used  $\tilde{S}$  as the feature space for an Abstracted DQN (ADQN). As the ADQN explored each task we updated  $\tilde{S}$  according to its respective ARMDP. In parallel, we trained recurrency-based models using S, namely, an RDQN, a Long-Short Term Memory DQN (LSTMDQN), and a Gated Recurrent DQN (GRUDQN). For a baseline, we also trained a DQN using S.

Rewards were very sparse for the Officeworld tasks and exploration via an  $\epsilon$ -greedy policy was often insufficient. As we didn't seek to demonstrate the efficacy of RL strategies under sparse rewards, we seeded each model with an optimal trajectory every 10th trajectory. Results for Experiment 2 can be seen in Figure 3.

In all tasks, the ADQN model learned at the rate of, or out-learned, all other models. This became more apparent as the complexity of the task increased (Task (e) being the most complex). We note that for Task (b), the baseline DQN performed just as well as the other mod-



Figure 3. Average rewards with standard deviations, evaluated over 10 trials using DQN models on Tasks (b-e) in the Officeworld domain.

els only because the shortest path to the goal included the coffee tile. Interestingly, for Task (e), the RDQN performed just as poorly as the baseline DQN model while the LSTMDQN and GRUDQN models (while learning slowly) appeared to be uncovering the relevant hidden mechanisms of the underlying RM. We hypothesize that the *forgetting* component of the GRUDQN and LSTDQN is a key differentiating factor in their learning ability. Still, we remind readers that uncovering U, as discussed in Experiment 1, involves a non-trivial cost before ADQN training can commence.

### 5.3 Experiment 3 - Breakfastworld

To evaluate the effectiveness of learning our RMs over  $H = (S \times$  $A \times \mathbb{R}$ )<sup>\*</sup>, rather than  $H = (S \times A)^*$ , we introduce a new domain, Breakfastworld, seen in Figure 4 (a). We include two tasks, Tasks (b) and (c), containing multiple interdependent reward signals. The domain uses  $P = \{c, w, e, l\}$  representing propositions for *cooking*, washing, eating and leaving, respectively. Task (b) contains an RM wherein the agent is meant to cook, eat, and then leave.  $(u_1, c, u_2)$ returns a reward of -0.1, reflecting a time cost. This time cost is experienced in a non-linear fashion as the agent grows bored and spending another turn cooking returns a reward of -0.3,  $(u_2, c, u_3)$ . However, this patience pays off in  $(u_3, e, u_4)$ , where eating a well cooked meal returns a reward of 2. In contrast,  $(u_2, e, u_5)$  returns 1, reflecting an undercooked meal. Finally,  $(u_4, l, term)$  and  $(u_5, l, term)$  reflect the last reward dependency, returning the cumulative reward for each path in the RM, encouraging episode completion. Task (c) builds on Task (b) by adding a reward dependency for washing. The penalty for washing after cooking once,  $(u_6, w, u_7)$ , is less than after cooking twice,  $(u_4, w, u_5)$ .

Both tasks reflect RMs more complex in terms of |U| than Task (e) in Experiment 1. However, we demonstrate that RM learning benefits in cases where rewards are interdependent as upstream reward con-



**Figure 4.** The Breakfastworld domain (a) with two RMs, (b) to cook, eat, and then leave, (c) to cook, eat, wash, and then leave.

flict resolution results in natural splits for downstream reward conflict resolution. Conversely, DFA learning suffers under such conditions as each reward conflict must be represented by its own DFA which is learned independently from any other DFA [12], hence they only learn according to  $H = (S \times A)^*$ .

We demonstrate how advantageous it can be that our ILP naturally learns according to  $H = (S \times A \times \mathbb{R})^*$  by applying ARMDPQ-Learning to Tasks (b) and (c). We then compare this to applying ARMDPQ-Learning on RMs equivalent in structure to Task (b) and (c), but that only emits cumulative rewards on terminal transitions. All other rewards are made 0. We refer to Tasks (b) and (c), in Figure 4, as Task (b) Full and Task (c) Full, respectively. Conversely, the RMs that only emit the cumulative rewards are referred to as Task (b) Cumulative and Task (c) Cumulative, respectively. The purpose of these cumulative tasks is to form a problem that produces reward conflicts only on the *l* tile that can not be contextualized by upstream reward conflicts. This would be similar to learning DFAs to parse the different possible rewards on the *l* tile only using  $H = (S \times A)^*$ . We show results for Experiment 3 in Table 1.

In all solvable cases of ARMDPQ-Learning, the agent achieved optimal performance and identified the correct |U|, for each task. We include an extension figure to Table 1 in Appendix F containing such results. Using Task (b) Full our cumulative solve time for all ILPs was dramatically smaller ( $\approx 4.4$  seconds) than when using Task (b) Cumulative ( $\approx 1,102.9$  seconds). This can be attributed to Task (b) Full requiring only a fourth of the number of trajectories to solve than Task (b) Cumulative, likely due to the leveraged context of upstream rewards. More impressive were the results for Task (c) Full which resulted in the cumulative solve time of  $\approx 153.8$  seconds. In no instance of attempting to learn in Task (c) Cumulative did we arrive at a stable solution in 100,000 seconds (>27 hours).

Task	$ T_o $	Solve Time (s)
Task (b) Full Task (b) Cumulative Task (c) Full Task (c) Cumulative	$6.2 \pm 1.1$ 22.4 ± 8.8 10.8 ± 3.0 n/a	$\begin{array}{c} 4.4 \pm 3.3 \\ 1102.9 \pm 1068.8 \\ 153.8 \pm 197.8 \\ \text{n/a} \end{array}$

 Table 1. Results over 10 trials using ARMDPQ-Learning on Tasks (b) and (c) from the Breakfastworld domain.

### 6 Conclusions

In this work, we introduced a formulation for learning RMs from non-Markov rewards by solving an ILP to learn an ARMDP. The ARMDP represents the cross-product between an RM and the underlying state space of an NMRDP. Importantly, we learn the AR-MDP, and thus the RM, without knowledge of P or L, by identifying hidden triggers that govern the rules of  $L: (S \times A \times S) \to 2^{P}$ . We also demonstrated how effective it can be to learn our automata according to  $H = (S \times A \times \mathbb{R})^*$ , rather than  $H = (S \times A)^*$ , a limitation of DFAs. We proved that the ARMDP is a suitable proxy for maximizing the reward expectations of the NMRDP for  $T_o$ . We validated our approach by applying ARMDPQ-Learning to train Q-Learning agents over black-box RMs. We also evaluated the representative power of the ARMDP by representing the state space of a DQN as the state space of an ARMDP. We compared the training profiles of these DQNs to the training profiles of various RDQN architectures that had to learn Markov representations over the sparse histories they encode.

An important aspect of this work is that the RMs we learn are inherently more interpretable than the sparse representation of histories in  $H = (S \times A \times \mathbb{R})^*$ . This is in part because we learn *minimal* RMs that encode H under the finite, structured, and symbolic representation of an automaton, but also because interpreting an RM doesn't require a comprehensive understanding of S and A outside of what the labeler encodes. This is a consequence of the RM being entirely decomposable from the ARMDP. For example, the RM in Figure 1 Task (b) conveys the need to bring coffee to the office without ever needing to describe the sequence of states and actions to achieve that goal. Since RL is predicated on the assumption that the reward function is the most concise, robust, and transferable definition of a task [20], we posit that RMs emerge as strong candidates for reward function representation and interpretation.

RM representations also carry important weight for the field of Inverse RL (IRL). In IRL, the goal is to learn a reward function from behavioral data that motivated that behavior [20, 26, 8]. While most approaches assume a Markov reward function, some of this work has been focused on learning rewards in Partially Observable MDPs (POMDPs) [8] and in Semi MDPs (SMDPs) [26]. We think that relaxing the Markov assumption in a principled way, analogous to this work, could yield insight into hidden abstractions made by a decision-maker. Moreover, given their interpretable form, it might shed light on what these hidden abstractions represent.

Lastly, we generally make note that this work is limited in that it assumes S and A are discrete. As such, we would like to consider automata extraction methods, in the future, using recurrent models [28, 21, 31] as this would extend our approach into the continuous case. Weiss et al. have demonstrated promising results in this regard [31]. Similarly, this work assumes R to be deterministic to some history. In scenarios where R is stochastic or noisy, we posit that we might be able to handle such cases by sampling histories to assess if the reward distributions differ in some statistical sense.

### Acknowledgements

This research was funded in part by the Air Force Office of Scientific Research Grant No. FA9550-20-1-0032 and the Office of Naval Research Grant No. N00014-19-1-2211.

### References

- D. Abel, W. Dabney, A. Harutyunyan, M. K. Ho, M. Littman, D. Precup, and S. Singh. On the Expressivity of Markov Reward. In *Advances in Neural Information Processing Systems*, volume 34, pages 7799–7812. Curran Associates, Inc., 2021. URL https://papers.neurips.cc/paper/ 2021/hash/4079016d940210b4ae9ae7d41c4a2065-Abstract.html.
- D. Angluin. Learning regular sets from queries and counterexamples. *Information and Computation*, 75(2):87–106, 1987. ISSN 08905401. doi: 10.1016/0890-5401(87)90052-6.
- [3] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling. The Arcade Learning Environment: An Evaluation Platform for General Agents. *Journal of Artificial Intelligence Research*, 47:253–279, 2013. ISSN 1076-9757. doi: 10.1613/jair.3912.
- [4] R. Bellman. A Markovian Decision Process. Journal of Mathematics and Mechanics, 6(5):679–684, 1957. ISSN 0095-9057. URL https: //www.jstor.org/stable/24900506.
- [5] R. Brafman, G. D. Giacomo, and F. Patrizi. LTLf/LDLf Non-Markovian Rewards. In *Proceedings of the AAAI Conference on Artificial Intelli*gence, volume 32, 2018. doi: 10.1609/aaai.v32i1.11572.
- [6] A. Camacho and S. A. McIlraith. Learning Interpretable Models Expressed in Linear Temporal Logic. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 29, pages 621–630, 2021. doi: 10.1609/icaps.v29i1.3529.
- [7] A. Camacho, R. Toro Icarte, T. Q. Klassen, R. Valenzano, and S. A. McIlraith. LTL and Beyond: Formal Languages for Reward Function Specification in Reinforcement Learning. In *Proceedings of the International Joint Conference on Artificial Intelligence*, volume 28, pages 6065–6073, 2019. ISBN 978-0-9992411-4-1. doi: 10.24963/ijcai.2019/840.
- [8] J. Choi and K.-E. Kim. Inverse Reinforcement Learning in Partially Observable Environments. *Journal of Machine Learning Research*, 12 (21):691–730, 2011. ISSN 1533-7928. URL http://jmlr.org/papers/v12/ choi11a.html.
- [9] P. J. K. Christoffersen, A. C. Li, R. T. Icarte, and S. A. McIlraith. Learning Symbolic Representations for Reinforcement Learning of Non-Markovian Behavior, 2023. URL http://arxiv.org/abs/2301.02952.
- [10] T. Dohmen, N. Topper, G. Atia, A. Beckus, A. Trivedi, and A. Velasquez. Inferring Probabilistic Reward Machines from Non-Markovian Reward Signals for Reinforcement Learning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 32, pages 574–582, 2022. doi: 10.1609/icaps.v32i1.19844.
- [11] D. Furelos-Blanco, M. Law, A. Jonsson, K. Broda, and A. Russo. Induction and Exploitation of Subgoal Automata for Reinforcement Learning. *Journal of Artificial Intelligence Research*, 70:1031–1116, 2021. ISSN 1076-9757. doi: 10.1613/jair.1.12372.
- [12] M. Gaon and R. I. Brafman. Reinforcement Learning with Non-Markovian Rewards. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 2019. doi: 10.1609/aaai.v34i04.5814.
- [13] Gurobi Optimization, LLC. Gurobi Optimizer Reference Manual, 2023. URL https://www.gurobi.com.
- [14] M. Hausknecht and P. Stone. Deep Recurrent Q-Learning for Partially Observable MDPs. In *Proceedings of the AAAI Fall Symposium*, 2015. URL http://arxiv.org/abs/1507.06527.
- [15] G. Hyde and E. Santos, Jr. Detecting Hidden Triggers: Mapping Non-Markov Reward Functions to Markov. arXiv preprint, 2024. URL http: //arxiv.org/abs/2401.11325. Full version with appendix and code.
- [16] R. T. Icarte, T. Klassen, R. Valenzano, and S. McIlraith. Using Reward Machines for High-Level Task Specification and Decomposition in Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pages 2107–2116, 2018. URL https://proceedings.mlr.press/v80/icarte18a.html.
- [17] R. T. Icarte, T. Q. Klassen, R. Valenzano, and S. A. McIlraith. Reward Machines: Exploiting Reward Function Structure in Reinforcement Learning. *Journal of Artificial Intelligence Research*, 73:173–208, 2022. ISSN 1076-9757. doi: 10.1613/jair.1.12440.
- [18] X. Li, C.-I. Vasile, and C. Belta. Reinforcement learning with temporal logic rewards. In 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 3834–3839. IEEE, 2017. ISBN 978-1-5386-2682-5. doi: 10.1109/IROS.2017.8206234.

- [19] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller. Playing Atari with Deep Reinforcement Learning, 2013. URL http://arxiv.org/abs/1312.5602.
- [20] A. Y. Ng and S. J. Russell. Algorithms for Inverse Reinforcement Learning. In Proceedings of the Seventeenth International Conference on Machine Learning, pages 663–670, 2000. ISBN 978-1-55860-707-1. doi: 10.5555/645529.657801.
- [21] C. W. Omlin and C. Giles. Extraction of rules from discrete-time recurrent neural networks. *Neural Networks*, 9(1):41–52, 1996. ISSN 08936080. doi: 10.1016/0893-6080(95)00086-0.
- [22] G. Rens and J.-F. Raskin. Learning Non-Markovian Reward Models in MDPs, 2020. URL http://arxiv.org/abs/2001.09293.
- [23] G. Rens, J.-F. Raskin, R. Reynouard, and G. Marra. Online Learning of non-Markovian Reward Models:. In *Proceedings of the 13th International Conference on Agents and Artificial Intelligence*, pages 74–86, 2021. ISBN 978-989-758-484-8. doi: 10.5220/0010212000740086.
- [24] M. Shvo, A. C. Li, R. T. Icarte, and S. A. McIlraith. Interpretable Sequence Classification via Discrete Optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021. doi: 10.1609/aaai. v35i11.17161.
- [25] R. S. Sutton and A. Barto. *Reinforcement Learning: An Introduction*. Adaptive computation and machine learning. The MIT Press, Cambridge, Massachusetts London, England, second edition edition, 2020. ISBN 978-0-262-03924-6.
- [26] C. Tan, Y. Li, and Y. Cheng. An inverse reinforcement learning algorithm for semi-Markov decision processes. In 2017 IEEE Symposium Series on Computational Intelligence (SSCI), pages 1–6, 2017. doi: 10.1109/SSCI.2017.8280816.
- [27] S. Thiébaux, C. Gretton, J. Slaney, D. Price, and F. Kabanza. Decision-Theoretic Planning with non-Markovian Rewards. *Journal of Artificial Intelligence Research*, 25(1):17–74, 2006. ISSN 1076-9757.
- [28] P. Tiňo and J. Šajda. Learning and Extracting Initial Mealy Automata with a Modular Neural Network Model. *Neural Computation*, 7(4): 822–844, 1995. ISSN 0899-7667, 1530-888X. doi: 10.1162/neco.1995. 7.4.822.
- [29] R. Toro Icarte, E. Waldie, T. Klassen, R. Valenzano, M. Castro, and S. McIlraith. Learning Reward Machines for Partially Observable Reinforcement Learning. In Advances in Neural Information Processing Systems, volume 32. Curran Associates, Inc., 2019. URL https://proceedings.neurips.cc/paper\_files/paper/2019/file/ 532435c44bec236b471a47a88d63513d-Paper.pdf.
- [30] C. J. C. H. Watkins and P. Dayan. Q-learning. *Machine Learning*, 8(3): 279–292, 1992. ISSN 1573-0565. doi: 10.1007/BF00992698.
- [31] G. Weiss, Y. Goldberg, and E. Yahav. Extracting Automata from Recurrent Neural Networks Using Queries and Counterexamples. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pages 5247–5256, 2018. URL https://proceedings.mlr.press/ v80/weiss18a.html.
- [32] Z. Xu, I. Gavran, Y. Ahmad, R. Majumdar, D. Neider, U. Topcu, and B. Wu. Joint Inference of Reward Machines and Policies for Reinforcement Learning. *Proceedings of the International Conference on Automated Planning and Scheduling*, 30:590–598, 2020. ISSN 2334-0843. doi: 10.1609/icaps.v30i1.6756.