

Joint Learning of Reward Machines and Policies in Environments with Partially Known Semantics

Christos K. Verginis

Uppsala University, Uppsala, Sweden

Cevahir Koprulu, Sandeep Chinchali, Ufuk Topcu

University of Texas at Austin, Austin, Texas, USA

Abstract

We study the problem of reinforcement learning for a task encoded by a reward machine. The task is defined over a set of properties in the environment, called atomic propositions, and represented by Boolean variables. One unrealistic assumption commonly used in the literature is that the truth values of these propositions are accurately known. In real situations, however, these truth values are uncertain since they come from sensors that suffer from imperfections. At the same time, reward machines can be difficult to model explicitly, especially when they encode complicated tasks. We develop a reinforcement-learning algorithm that infers a reward machine that encodes the underlying task while learning how to execute it, despite the uncertainties of the propositions' truth values. In order to address such uncertainties, the algorithm maintains a probabilistic estimate about the truth value of the atomic propositions; it updates this estimate according to new sensory measurements that arrive from [exploration](#) of the environment. Additionally, the algorithm maintains a hypothesis reward machine, which acts as an estimate of the reward machine that encodes the task to be learned. As the agent explores the environment, the algorithm updates the hypothesis reward machine according to the obtained rewards and

Email addresses: christos.verginis@angstrom.uu.se (Christos K. Verginis), cevahir.koprulu@utexas.edu (Cevahir Koprulu), sandeepc@utexas.edu (Sandeep Chinchali), utopcu@utexas.edu (Ufuk Topcu)

the estimate of the atomic propositions’ truth value. Finally, the algorithm uses a Q-learning procedure for the states of the hypothesis reward machine to determine [an optimal policy that accomplishes the task](#). We prove that the algorithm successfully infers the reward machine and asymptotically learns a policy that accomplishes the respective task.

Keywords: Reinforcement learning; Reward machines; Perception limitations;

1. Introduction

Reinforcement learning (RL) studies the problem of learning an optimal behavior for agents with unknown dynamics in potentially unknown environments. [When the task assigned to the agent admits a certain form of compositionality,](#)
5 [such as decomposition to temporally dependent sub-tasks, a variety of methods incorporate high-level knowledge that can help the agent explore the environment more efficiently \[1\].](#) This high-level knowledge is usually expressed through abstractions of the environment and sub-task sequences [2, 3, 4] or in linear temporal logic [5, 6]. [Intuitively, such methodologies inform the agent](#)
10 [about the compositionality of the task, including its decomposition to sub-tasks as well as the sub-task execution sequences that lead to its completion.](#) Consequently, the agent learns sub-task policies to execute such sequences and [accomplish the overall task](#). Recently, the authors in [7] proposed the concept of *reward machines* in order to provide high-level information to the agent in
15 the form of rewards. [Reward machines are finite-state structures that encode a possibly non-Markovian reward function.](#) They decompose a task into several temporally related subtasks, such as “get coffee and bring it to the office without encountering obstacles”. Reward machines allow the composition of such subtasks in flexible ways, including concatenations, loops, and conditional
20 rules, providing the agent access to high-level temporal relationships among the subtasks. These subtasks are defined over a set of properties in the environment, called atomic propositions, and represented by Boolean variables, such as “office” and “obstacle” in the aforementioned example. Intuitively, as the agent

navigates in the environment, it transits between the states within a reward
25 machine; such transitions are triggered by the truth values of the atomic propo-
sitions. After every such transition, the reward machine outputs the reward the
agent should obtain according to the encoded reward function. Furthermore,
[7] develops a q-learning method that learns the optimal policy associated with
the reward function that is encoded by the reward machine.

30 In many practical situations, the reward machines encode complex tempo-
ral relationships among the subtasks and are hence too difficult to construct.
Therefore, the need emerges for the learning agent to learn how to accomplish
the underlying task without having a priori access to the respective reward ma-
chine. The work [8] develops an algorithm that infers both the reward machine
35 and an optimal RL policy from the trajectories of the system.

The procedure of inferring and using a reward machine is tightly connected
to the values of the atomic propositions; such values trigger the reward ma-
chine’s transitions, which output the reward the agent should obtain. In related
literature, there are two main methods to encode the learning agent’s access to
40 the truth values of the atomic propositions. First, the agent has access to a
so-called *labelling function*, a map that provides which atomic propositions are
true in which areas of the environment, e.g., which room of an indoor environ-
ment has an office. Second, the agent accurately detects the truth values of the
atomic propositions from sensors as it navigates in the environment.

45 The aforementioned methods rely on assumptions that are unrealistic in
practical scenarios. Firstly, many situations involve agents operating in a priori
unknown environments, and hence accurate knowledge of a labelling function
cannot be obtained a priori. Secondly, autonomous agents are endowed with
sensors that naturally suffer from imperfections, such as misclassifications or
50 missed object detections. Consequently, the truth values of the atomic propo-
sitions are uncertain. When the reward machine that encodes the task is a
priori unknown, uncertainties in such values complicate significantly the RL
problem since they interfere with the estimation of the reward machine and,
consequently, the learning of a suitable policy. Therefore, it is necessary to

55 develop RL algorithms that are robust to such uncertainties.

In this paper, we investigate an RL problem for a task encoded by a reward machine over a set of atomic propositions. The reward machine and the labelling function associated with the atomic propositions are a priori unknown. The agent obtains the truth values of the atomic propositions via measurements from
60 sensors that suffer from imperfections, leading to uncertainties in the atomic propositions' truth values. Our main contribution lies in the development of an RL algorithm that learns a policy that achieves the given task despite the a priori unknown reward machine and the atomic proposition uncertainties. In order to account for such uncertainties, the algorithm holds a probabilistic
65 belief over the truth values of the atomic propositions and it updates this belief according to new sensory measurements that arrive from the exploration of the environment. Furthermore, the algorithm maintains a so-called hypothesis reward machine to estimate the unknown reward machine that encodes the task. It uses the rewards the agent obtains as it explores the environment and
70 the aforementioned belief on the atomic propositions to update this hypothesis reward machine. Finally, the algorithm deploys a q-learning procedure that is based on the hypothesis reward machine and the probabilistic belief to learn an optimal policy that accomplishes the underlying task. An illustrative diagram of the proposed algorithm is depicted in Fig. 1.

75 We establish theoretical guarantees on the algorithm's convergence to a policy that accomplishes the underlying task. Our guarantees rely on the following sufficient conditions. First, the belief updates lead to better estimation of the truth values of the atomic propositions. That is, there exists a finite time instant after which the probability that the estimated propositions' values match
80 the actual ones is more than 0.5. Second, the learning agent explores the environment sufficiently well in each episode of the proposed RL algorithm. More specifically, the length of each episode is larger than a pre-defined constant associated with the size of the operating environment and of the reward machine that encodes the task to be accomplished. Based on the aforementioned
85 assumptions, we prove that the proposed algorithm 1) asymptotically infers a

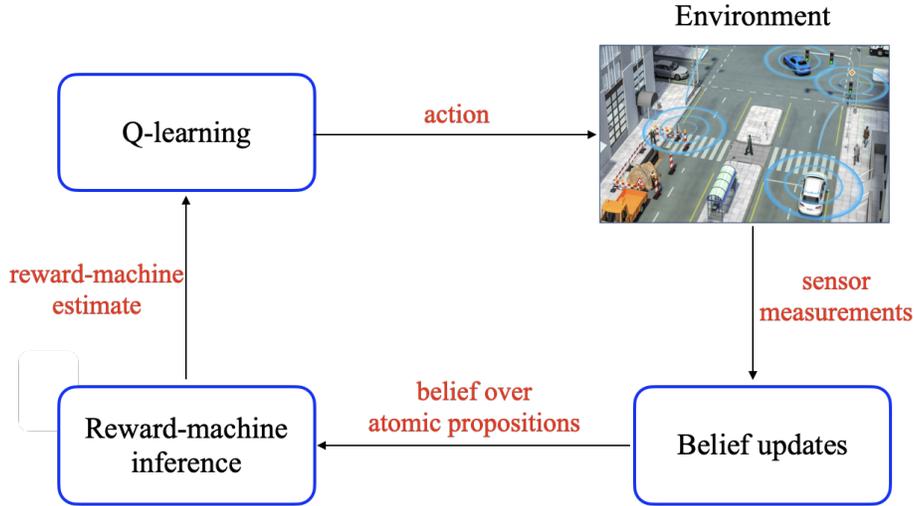


Figure 1: Overview of the proposed reinforcement-learning algorithm, consisting of perception updates, inference of the reward machine, and q-learning.

reward machine that is equivalent to the one encoding the underlying task, and 2) asymptotically learns a policy that accomplishes this task.

Finally, we carry out experiments comparing the proposed algorithm with a baseline version that does not update the beliefs of the propositions' truth values as well as with nominal q-learning and double deep q-network (DDQN) algorithms. The experiments show that the proposed algorithm outperforms the baseline version and the nominal algorithm in terms of convergence. The experiments further show the robustness of the proposed algorithm to inaccurate sensor measurements.

2. Related Work

Previous RL works on providing higher-level structure and knowledge about the reward function focus mostly on abstractions and hierarchical RL; in such cases, a meta-controller decides which subtasks to perform, and a controller decides which actions to take within a subtask (e.g., [9, 3, 10]). Other works use temporal logic languages to express a specification, and then generate the

corresponding reward functions, typically using deterministic automata (DA) [11, 5, 12, 6, 13, 14, 15, 16, 17, 18]. Inspired by [7], we use in this work the concept of reward machines, which encode more compactly and expressively high-level information on the reward function of the agent. In particular and
105 unlike DA, reward machines offer the ability of assigning separate state-action reward functions to each of their states; the learning agent obtains rewards depending on the transitions of the reward machine and hence each transition can be viewed as a separate sub-task. Therefore, the reward machine gives a certain structure to the agent in the sense that it informs it about the decomposition
110 of the task into sub-tasks and the sequence these sub-tasks must be executed. Consequently, the agent learns how and in what sequence to execute the different sub-tasks towards the completion of the overall task. Additionally, unlike the works in the related literature, we consider that the reward machine that encodes the task at hand is a priori *unknown*.

115 There exists a large variety of works dealing with perception uncertainty [19, 20, 21, 22, 23, 24, 25, 26, 27, 28]. However, most of these works consider the *planning* problem, which consists of computing an optimal policy by explicitly employing the underlying agent model. Furthermore, the considered uncertainty usually arises from probabilistic dynamics, which is resolved either
120 by active-perception procedures, belief-state propagation, or analysis with partially observable Markov decision processes (POMDPs). Besides, the works that consider high-level reward structure (like temporal logic [19, 20, 22, 24]), assume full availability of the respective reward models, unlike the problem setting considered in this paper.

125 Several recent works consider tasks encoded by Non-Markovian reward functions that are *a priori unknown*, using both automata [29, 30, 31, 32] as well as reward machines [33, 34, 8]. In such works, the agents infers the reward function from the executed trajectories in the environment. Similarly, [35] develops an algorithm that infers probabilistic reward machines. The aforementioned works,
130 however, do not consider uncertainty in the truth values of the atomic propositions; [29] and [33] assume partial environment state observability by employing

POMDP models, while considering the truth values of atomic propositions in the environment known. In contrast, in this work we consider uncertainty in the semantic representation of the environment, i.e., the truth value of the environment properties that define the agent’s task is a priori unknown.

3. Problem Formulation

This section formulates the considered problem. We first describe the model of the learning agent and the environment, the reward machine that encodes the task to be accomplished, and the observation model of the sensor the agent uses to detect the values of atomic propositions.

3.1. Agent Model

We model the interaction between the agent and the environment by a Markov decision process (MDP) [36, 37], formally defined below.

Definition 1. A *labelled MDP (l-MDP)* is a tuple $\mathcal{M} = (S, s_I, A, \mathcal{T}, R, \mathcal{AP}, L_G, \hat{L}, \gamma)$ consisting of the following: S is a finite state space, resembling the operating environment; $s_I \in S$ is an initial state¹; A is a finite set of actions; $\mathcal{T} : S \times A \times S \rightarrow [0, 1]$ is a probabilistic transition function, modelling the agent dynamics in the environment; $R : S^* \times A \times S \rightarrow \mathbb{R}$ is a reward function, specifying payoffs to the agent; $\gamma \in [0, 1)$ is a discount factor; \mathcal{AP} is a finite set of Boolean atomic propositions, each one taking values in $\{\text{True}, \text{False}\}$; $L_G : S \rightarrow 2^{\mathcal{AP}}$ is the ground-truth labelling function, assigning to each state $s \in S$ the atomic propositions that are true in that state; $\hat{L} : S \rightarrow 2^{\mathcal{AP}}$ is an estimate of the ground-truth labelling function L_G . We define the size of \mathcal{M} , denoted as $|\mathcal{M}|$, to be $|S|$ (i.e., the cardinality of the set S).

The atomic propositions \mathcal{AP} are Boolean-valued properties of the environment; we denote by $s \models p$ if an atomic proposition $p \in \mathcal{AP}$ is true at state

¹ s_I can also be a distribution of initial states.

$s \in S$. The ground-truth labelling function L_G provides which atomic propositions are true in the states of the environment, i.e., $L_G(s) = P \subseteq \mathcal{AP}$ is equivalent to $s \models p$, for all $p \in P$. We consider the setting where L_G is *unknown to the agent*, which detects the truth values of \mathcal{AP} through sensor units such as cameras or range sensors. *Since such sensors suffer from imperfections (e.g., noise), the agent maintains* a time-varying probabilistic belief $\hat{\mathcal{L}}_i : S \times 2^{\mathcal{AP}} \rightarrow [0, 1]$ about the truth values of \mathcal{AP} , where i denotes a time index. More specifically, for a state $s \in S$ and a subset of atomic propositions $P \subseteq \mathcal{AP}$, $\hat{\mathcal{L}}_i(s, P)$ represents the probability of the event that P is true at s , i.e., $\hat{\mathcal{L}}_i(s, P) = \Pr(\bigcap_{p \in P} s \models p)$. Notice that at each time index i and for every state $s \in S$, it holds that $\sum_{P \subseteq \mathcal{AP}} \hat{\mathcal{L}}_i(s, P) = 1$. The prior belief of the agent might be an uninformative prior distribution. We assume that the truth values of the propositions are mutually independent in each state, i.e.,

$$\Pr\left(\bigcap_{p \in P} s \models p\right) = \prod_{p \in P} \Pr(s \models p), \quad \forall s \in S, P \subseteq \mathcal{AP}$$

and also between every pair,

$$\Pr(s \models P \wedge s' \models P') = \Pr(s \models P) \Pr(s' \models P'), \quad \forall s, s' \in S, P, P' \subseteq \mathcal{AP}$$

The agent's belief $\hat{\mathcal{L}}_i$, for some time index $i \geq 0$, determines the agent's inference of the environment configuration, and hence the estimate \hat{L} of the labelling function. More specifically, we define $\hat{L} : (S \times 2^{\mathcal{AP}} \rightarrow [0, 1]) \times S \rightarrow 2^{\mathcal{AP}}$, with

$$\hat{L}(\hat{\mathcal{L}}_i, s) = \{p \in \mathcal{AP} : \hat{\mathcal{L}}_i(s, p) \geq 0.5\}, \quad (1)$$

155 i.e., the most probable outcomes. The following example illustrates the relation between L_G and \hat{L} .

Example. *Let the 10-state office workspace be shown in Fig. 2. The set of atomic propositions is $\mathcal{AP} = \{c, o, X\}$, corresponding to “coffee”, “office”, and “obstacle”, respectively. The ground-truth labelling function, shown at the top of*
 160 *Fig. 2, is defined as $L_G(s_{12}) = \{c\}$, $L_G(s_{22}) = L_G(s_{23}) = \{X\}$, $L_G(s_{24}) = \{o\}$,*

	1	2	3	4	5
1	▲	c			
2		X	X	o	

	1	2	3	4	5
1	▲	c, X		o	
2		c	o, X	o	

Figure 2: Top: A workspace with 10 states s_{ij} illustrating the ground-truth labelling function $L_G : S \rightarrow 2^{\{c,o,X\}}$, where $\{c,o,X\}$ stand for {"coffee", "office", "obstacle"}, respectively; Bottom: Perception uncertainty leads to the estimate $\hat{L}(\hat{\mathcal{L}}, s) : S \rightarrow 2^{\{c,o,X\}}$, according to which the truth values of several properties are incorrect (e.g. X in s_{12} or the absence of X in s_{22}).

and \emptyset in the rest of the states. Imperfect sensory measurements might lead to a distribution $\hat{\mathcal{L}}$ and an estimate $\hat{L}(\hat{\mathcal{L}}, \cdot)$ as shown at the bottom of Fig. 2, i.e., $\hat{L}(s_{12}) = \{c, X\}$, $\hat{L}(s_{22}) = \{c\}$, $\hat{L}(s_{14}) = \hat{L}(s_{24}) = \{o\}$, $\hat{L}(s_{23}) = \{o, X\}$, and \emptyset in the rest of the states.

165 A policy is a function that maps states in S to a probability distribution over
actions in A . At state $s \in S$, an agent, using policy π , picks an action a with
probability $\pi(s, a)$, and the new state s' is chosen with probability $\mathcal{T}(s, a, s')$,
defined in Def. 1. A policy π and the initial state s_I together determine a
stochastic process; we write $S_0 A_0 S_1 \dots$ for the random trajectory of states and
170 actions.

A trajectory is a realization of the agent's stochastic process $S_0 A_0 S_1 \dots$: a sequence of states and actions $s_0 a_0 s_1 \dots s_k a_k s_{k+1}$ with $s_0 = s_I$. Its corresponding label sequence is $\ell_0 \ell_1 \dots \ell_k$ with $\ell_j = L_G(s_j)$, for all $j \leq k$. Similarly, the reward sequence is $r_0 r_1 \dots r_k$, where $r_j = R(s_0 \dots s_j a_j s_{j+1})$, for all $j \leq k$. A trajectory

175 $s_0 a_0 s_1 \dots s_k a_k s_{k+1}$ achieves a reward $\sum_{j=0}^k \gamma^j R(s_0 \dots s_j a_j s_{j+1})$, where γ is the discount factor defined in Def. 1. Note that the definition of the reward function assumes that the reward is a function of the whole trajectory; this allows the reward function to be non-Markovian [7]. Given a trajectory $s_0 a_0 s_1 \dots s_{k+1}$ and a belief $\hat{\mathcal{L}}_i$, we naturally define then the *observed* label sequence by $\hat{\ell}_0 \hat{\ell}_1 \dots \hat{\ell}_k$,
 180 with $\hat{\ell}_j = \hat{L}(\hat{\mathcal{L}}_i, s_j)$ for all $j \leq k$ and some $i \geq 0$.

Example (Continued). *Let the 10-state office workspace be shown in Fig. 2. Assume that the agent receives award 1 if it brings coffee to the office without encountering states with obstacles, and zero otherwise. Let us now consider a state sequence $s_{11} s_{12} s_{13} s_{14} s_{24}$ of a trajectory. According to the ground-truth*
 185 *labelling function (see top of Fig. 2), such a trajectory would produce the label and reward sequence $(\emptyset, 0)(c, 0)(\emptyset, 0)(\emptyset, 0)(o, 1)$. However, the observed label and reward sequence, based on the estimate \hat{L} (see bottom of Fig. 2), is $(\emptyset, 0)((c, X), 0)(\emptyset, 0)(o, 0)(o, 1)$.*

3.2. Reward Machines

190 Encoding a (non-Markovian) reward in a type of finite state-machine is achieved by reward machines, introduced in [7] and formally defined below.

Definition 2. *A reward machine is a tuple $\mathcal{A} = (V, v_I, 2^{\mathcal{A}^{\mathcal{P}}}, \mathcal{R}, \delta, \sigma)$ that consists of the following: V is a finite set of states; $v_I \in V$ is an initial state; \mathcal{R} is an input alphabet; $\delta : V \times 2^{\mathcal{A}^{\mathcal{P}}} \rightarrow V$ is a deterministic transition function;*
 195 *$\sigma : V \times 2^{\mathcal{A}^{\mathcal{P}}} \rightarrow \mathcal{R}$ is an output function. We define the size of \mathcal{A} , denoted by $|\mathcal{A}|$, to be $|V|$ (i.e., the cardinality of the set V).*

The run of a reward machine \mathcal{A} on a sequence of labels $\ell_0 \dots \ell_k \in (2^{\mathcal{A}^{\mathcal{P}}})^*$ is a sequence $v_0(\ell_0, r_0)v_1(\ell_1, v_1) \dots v_k(\ell_k, v_k)$ of states and label-reward pairs such that $v_0 = v_I$ and for all $j \in \{0, \dots, k\}$, we have $\delta(v_j, \ell_j) = v_{j+1}$ and $\sigma(v_j, \ell_j) =$
 200 r_j . We write $\mathcal{A}(\ell_0 \dots \ell_k) = r_0 \dots r_k$ to connect the input label sequence to the sequence of rewards produced by the reward machine \mathcal{A} . We say that a reward machine \mathcal{A} encodes the reward function R of an **l-MDP** on the ground truth if, for every trajectory $s_0 a_0 \dots s_k a_k s_{k+1}$ and the corresponding label sequence

$\ell_0 \dots \ell_k$, the sequence of rewards equals $\mathcal{A}(\ell_0 \dots \ell_k)$. Moreover, given a fixed
 205 $i \geq 0$, we say that a reward machine \mathcal{A} *encodes the reward function R on $\hat{\mathcal{L}}_i$*
 if for every trajectory $s_0 a_0 \dots s_k a_k s_{k+1}$ and the corresponding *observed* label
 sequence $\hat{\ell}_0 \dots \hat{\ell}_k$ the sequence of rewards equals $\mathcal{A}(\hat{\ell}_0 \dots \hat{\ell}_k)$. Note, however,
 that there might not exist a reward machine that encodes the reward function
 on $\hat{\mathcal{L}}_i$, as the next example shows.

Example (Continued). *Let the 10-state office workspace be shown in Fig. 2. Assume that the agent receives award 1 if it brings coffee to the office without encountering states with obstacles, and zero otherwise. Such a task is encoded by the reward machine shown in Fig. 3. Let also the state sequence of a given trajectory be $s_{11}s_{12}s_{13}s_{14}s_{24}$. As stated before, the trajectory produces, via the ground-truth labelling function, the label and reward sequence $(\emptyset, 0)(c, 0)(\emptyset, 0)(\emptyset, 0) (o, 1)$. Such a label sequence produces the run of the reward machine*

$$v_0(\emptyset, 0)v_0(c, 0)v_1(\emptyset, 0)v_1(\emptyset, 0)v_1(o, 1)v_3.$$

The observed label and reward sequence, based on the estimate \hat{L} , is $(\emptyset, 0) ((c, X), 0) (\emptyset, 0)(o, 0)(o, 1)$, which produces the reward-machine run

$$v_0(\emptyset, 0)v_0((c, X), 0)v_2(\emptyset, 0)v_2(o, 0)v_2(o, 0)v_2,$$

210 *which clearly does not comply with the observed reward sequence. Hence, we conclude that the reward machine does not encode the reward function on $\hat{\mathcal{L}}$. In fact, note that we cannot construct a reward machine that encodes the reward function on $\hat{\mathcal{L}}$. The actual reward function of bringing coffee from s_{12} to the office in s_{24} cannot be expressed via the set \mathcal{AP} with this specific $\hat{\mathcal{L}}$, because of*
 215 *the ambiguity in the states satisfying “o” and “c”; a single accepting trajectory does not correspond to a unique observed label sequence.*

3.3. Observation Model

The probabilistic belief is updated based on an observation model \mathcal{O} . More specifically, at each time step, the agent’s perception module processes a set of

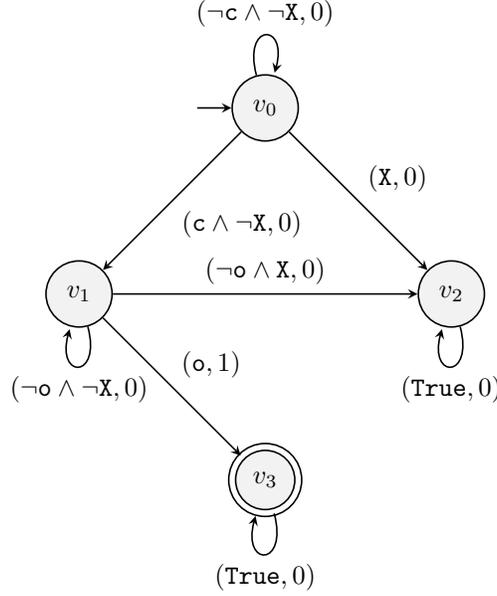


Figure 3: Reward machine for the task of Fig. 2.

220 sensory measurements regarding the atomic propositions \mathcal{AP} . In the Bayesian framework, the observation model is used for the update of the agent's belief, as specified in the following definition.

Definition 3. Let $\mathcal{Z}(s_1, s_2, p) \in \{\text{True}, \text{False}\}$ denote the perception output of the agent, when the agent is at state s_1 , for the atomic proposition $p \in \mathcal{AP}$ at
 225 s_2 . The joint observation model of the agent is $\mathcal{O} : S \times S \times \mathcal{AP} \times \{\text{True}, \text{False}\} \rightarrow [0, 1]$, where $\mathcal{O}(s_1, s_2, p, b)$ represents the probability that $\mathcal{Z}(s_1, s_2, p) = \text{True}$ if the truth value of p is identical to b .

More specifically, the joint observation describes the likelihoods $\mathcal{O}(s_1, s_2, p, \text{True}) = Pr(\mathcal{Z}(s_1, s_2, p) = \text{True} \mid s_2 \models p)$ and $\mathcal{O}(s_1, s_2, p, \text{False}) = Pr(\mathcal{Z}(s_1, s_2, p) =$
 230 $\text{True} \mid s_2 \not\models p)$. An accurate observation model is the one for which $\mathcal{O}(s_1, s_2, p, b) = 1$ for $b = \text{True}$ and $\mathcal{O}(s_1, s_2, p, b) = 0$ for $b = \text{False}$. In the Bayesian framework, the observation model is used for the update of the agent's belief. Nevertheless, in the absence of such observation model, one can perform the update in a

frequentist way.

235 *3.4. Problem Statement*

We consider an **l-MDP**-modeled autonomous agent, whose task is encoded by a reward machine that is unknown to the agent. The ground-truth labelling function L_G , providing the truth values of the atomic propositions, is also unknown to the agent. The formal definition of the problem statement is as follows.

240 **Problem 1.** *Let an **l-MDP** agent model be $\mathcal{M} = (S, s_I, A, \mathcal{T}, R, \mathcal{AP}, L_G, \hat{L}, \gamma)$ with unknown ground-truth labelling function L_G and transition function \mathcal{T} . Let also an unknown reward machine that encodes the reward function R on the ground truth be \mathcal{A} , representing a task to be accomplished. Develop an algorithm that learns a policy π that maximizes $Z(s_I) = \mathbb{E}_\pi[\sum_{i=0}^{\infty} \gamma^i R(S_0, A_0 \dots S_{i+1})]$.*

245

Remark 1. *We note that the problem of unknown ground-truth labelling function L_G and uncertain sensor observations does not apply only to cases with reward machines, but also to more general RL frameworks that employ labelling-function mappings to identify properties of the environment (e.g., [7, 42, 31, 250 35]). However, reward machines [7] constitute state-of-the-art tools when it comes to symbolic RL with task decomposition. Additionally, reward machines naturally depend on the atomic propositions of the environment and hence on the respective labelling function. Therefore, inference of an unknown reward machine poses a unique challenge when such a labelling function is uncertain. 255 For the aforementioned reasons, we focus on the problem of RL with unknown reward machines. Nevertheless, we illustrate how the proposed algorithm, given in the next section, applies to more general RL frameworks.*

4. Joint Inference of Reward machines and Policies for RL

Algorithm 1 JIRP Algorithm

```
1:  $\mathcal{H} \leftarrow \text{Initialize}(V)$ 
2:  $Q = \{q^v | v \in V\} \leftarrow \text{Initialize}()$ 
3:  $X \leftarrow \emptyset$ 
4: for episode  $n = 1, 2, \dots$  do
5:    $(\lambda, \rho, Q) \leftarrow \text{QRM\_episode}(\mathcal{H}, Q, \mathcal{L})$ 
6:   if  $\mathcal{H}(\lambda) \neq \rho$  then           // The trace  $(\lambda, \rho)$  is inconsistent with  $\mathcal{H}$ 
7:      $\text{Add}(X, (\lambda, \rho))$            // Append the set  $X$  with the trace  $(\lambda, \rho)$ 
8:      $\mathcal{H} \leftarrow \text{Infer}(X)$        // Infer a new reward machine from  $X$ 
9:      $Q \leftarrow \text{Initialize}()$ 
10:  end if
11: end for
```

260 The main results of the paper consist of the development of a methodology
that accommodates the uncertainty in the ground-truth labelling function L_G
and the environmental observations. One can incorporate such a methodology to
standard RL algorithms to account for the aforementioned uncertainty. For the
sake of presentation, we illustrate the incorporation of our methodology to the
265 JIRP algorithm [8], which considers the joint inference of reward machines and
policies with perfect knowledge of the environment semantics. Consequently,
we provide in this section an overview of JIRP, before proceeding to the main
results of the paper.

The JIRP algorithm [8] (see Algorithm 1) aims at learning the optimal policy for maximizing $Z(s_I)$ by maintaining a hypothesis reward machine (RM) \mathcal{H} (initialized with a set of states V in line 1). Its main component is the `QRM_episode()` (shown in Algorithm 2), originally proposed in [7] for accurately known reward machines. The QRM episode outputs a set (λ, ρ) , called *trace*, and a set Q of q-functions q^v (initialized in line 2), one for each state $v \in V$ of the hypothesis reward machine; recall that a reward machine decomposes the overall task into sub-tasks, one for each one of its states. The current state v of the hypothesis reward machine guides the exploration by determining

Algorithm 2 QRM_episode($\mathcal{H}, Q, \mathcal{L}$)

Input: A reward machine $\mathcal{H} = (V, v_I, 2^{\mathcal{P}}, \mathbb{R}, \delta, \sigma)$, a set of q -functions $Q = \{q^v | v \in V\}$, a labelling function \mathcal{L} .

```
1:  $s \leftarrow \text{InitialState}()$ ;  $v \leftarrow v_I$ ,  $\lambda \leftarrow []$ ;  $\rho \leftarrow []$ 
2: for  $0 \leq t < \text{eplength}$  do
3:    $a \leftarrow \text{GetEpsilonGreedyAction}(q^v, s)$ 
4:    $s' \leftarrow \text{ExecuteAction}(s, a)$  // Transition in l-MDP  $\mathcal{M}$  based on
    $\mathcal{T}(s', a, s)$ 
5:    $v' \leftarrow \delta(v, \mathcal{L}(s, a, s'))$  // Transition in  $\mathcal{H}$ 
6:    $r \leftarrow \text{ObserveReward}()$ 
7:    $q^v(s, a) \leftarrow \text{Update}(r, s')$  // Update of  $q$ -function for state  $v$  of  $\mathcal{H}$ 
8:   for  $\hat{v} \in V \setminus v$  do
9:      $\hat{v}' \leftarrow \delta(\hat{v}, \mathcal{L}(s, a, s'))$  // Transition in  $\mathcal{H}$ 
10:     $\hat{r} \leftarrow \sigma(\hat{v}, \mathcal{L}(s, a, s'))$  // Reward calculation from  $\mathcal{H}$ 
11:     $q^{\hat{v}}(s, a) \leftarrow \text{Update}(\hat{r}, s')$  // Update of  $q$ -function for state  $\hat{v}$  of  $\mathcal{H}$ 
12:   end for
13:    $\text{Append}(\lambda, \mathcal{L}(s, a, s'))$ ;  $\text{Append}(\rho, r)$ 
14:    $s \leftarrow s'$ ;  $v \leftarrow v'$ 
15: end for
16: return  $(\lambda, \rho, Q)$ 
```

which q -function is used to choose the next action (line 3 of Alg. 2). However, in each single exploration step, the algorithm updates the q -functions that correspond to all reward-machine states (lines 7 and 11 of Alg. 2). Furthermore, the algorithm uses the sets λ and ρ to collect the labelling and reward sequences $\mathcal{L}(s, a, s')$ and r , respectively, from the agent trajectory (line 13 of Alg. 2)). Note that the returned rewards in ρ are *observed* (line 6 of Alg. 2), since the reward machine encoding the actual reward function is not known. Instead, JIRP operates with a hypothesis reward machine, which is updated using the traces of QRM_episode(). In particular, the episodes of QRM are used to collect traces

and update q -functions. As long as the traces are consistent with the current hypothesis reward machine, QRM explores more of the environment using the reward machine to guide the search. However, if a trace (λ, ρ) is detected that is inconsistent with the hypothesis reward machine (i.e., $\mathcal{H}(\lambda) \neq \rho$, line 6 of Alg. 1), JIRP stores it in a set X (line 7 of Alg. 1) - the trace (λ, ρ) is called a counterexample and the set X a sample. Once the sample is updated, the algorithm re-learns a new hypothesis reward machine (line 8 of Alg. 1) and proceeds. More information can be found in [8, 7]. We further detail the various functions used in Alg. 2: The function `GetEpsilonGreedyAction` (q^v, s) chooses an action based on an ε_a -greedy method, i.e., it chooses the action $\operatorname{argmax}_{a \in A}(q^v(s, a))$ with probability $1 - \varepsilon_a$ and a random action $a \in A$ with probability ε_a for a small positive constant ε_a . The function `Update` (r, s') updates $q^v(s, a)$ based on

$$q^v(s, a) \leftarrow q^v(s, a) + \alpha_{lr} \left(r + \gamma \max_{a' \in A} q^v(s', a') - q^v(s, a) \right)$$

for a constant α_{lr} representing the learning rate.

270

5. Joint Learning of Reward Machines and Policies with Partially Known Semantics

This section gives the main results of this paper, which is joint perception and inference of policies and reward machines. In particular, we extend the
 275 JIRP algorithm, presented in Section 4, to take into account the uncertainty in the labelling function. More specifically, our algorithm tries to infer the reward machine \mathcal{A} that encodes the reward function on the most probable outcomes, $\hat{\mathcal{L}}(\hat{\mathcal{L}}_j, \cdot)$, i.e., $\mathcal{A}(\hat{\ell}_0, \dots, \hat{\ell}_k) = r_0 \dots r_k$ for some $j \geq 0$. At the same time, the agent's belief is updated at every time step based on the observation model of
 280 Definition 3. Similarly to [19], if the agent's knowledge about the environment has changed significantly, the agent then updates its belief $\hat{\mathcal{L}}$ and aims to infer a new reward machine.

Algorithms 3 and 4 illustrate the aforementioned procedure. The algorithms are similar to 1 and 2, respectively, with three main differences: first, Alg. 3

285 aims to approximate a reward machine \mathcal{A} that encodes the reward functions
on the “current” belief $\hat{L}(\hat{\mathcal{L}}_h, \cdot)$; second, Alg. 4 updates this belief using the
observation obtained by the agent in its state (line 14 of Alg. 4); third, Alg.
3 updates the belief $\hat{L}(\hat{\mathcal{L}}_h, \cdot)$, and consequently the reward machine \mathcal{H} , if it is
different enough from what is dictated by the obtained observations (lines 13-18
290 of Alg. 3). We present next the reasoning of Alg. 3 and 4 in more detail.

Algorithm 3 uses a hypothesis reward machine \mathcal{H} (line 1 of Alg. 3) to guide
the learning process. Similarly to [8], it runs multiple QRM episodes to update
the Q functions, and uses the collected traces to update the counterexamples
in the set X and \mathcal{H} (lines 7-12 of Alg. 3). Nevertheless, in our case, \mathcal{H} aims
295 to approximate the reward machine \mathcal{A} that encodes the reward function on
the most probable outcomes, $\hat{L}(\hat{\mathcal{L}}_h, \cdot)$, since that is the available information to
the agent; $\hat{\mathcal{L}}_h$ is the “current” belief that the agent uses in its policy learning
and reward-machine inference. This is illustrated via a modified version of the
QRM episode algorithm (QRM_episode_mod invoked in line 7 of Alg. 3), shown
300 in Algorithm 4. In particular, QRM_episode_mod uses the probabilistic belief
 $\hat{\mathcal{L}}_h$ in order to navigate in the hypothesis reward machine \mathcal{H} (lines 5, 9 of Alg. 4).
Moreover, $\hat{\mathcal{L}}_h$ is used to update the counterexample set X with the inconsistent
traces (line 15 of Alg. 4 and lines 8-12 of Alg. 3). Hence, Alg. 3 aims to infer the
reward machine that encodes the reward function on the belief $\hat{\mathcal{L}}_h$. Of course, as
305 illustrated in Example 3.1, there might not exist a reward machine that encodes
the reward function on $\hat{\mathcal{L}}_h$. In that case, the algorithm cannot infer the “correct”
reward machine. For this reason, the modified QRM-episode algorithm updates
a recurring probabilistic belief $\hat{\mathcal{L}}_j$ (line 14) in a Bayesian manner and based
on the observation model of Def. 3; we provide the function BayesUpdate() in
310 Section 5.1. The environment probabilistic belief $\hat{\mathcal{L}}_h$ used in the policy- and
reward-machine-learning is updated with the current belief $\hat{\mathcal{L}}_j$ only if the two
are significantly different (lines 13-18 of Algorithm 3), which is evaluated by the
function SignifChange(). More specifically, if $\hat{\mathcal{L}}_j$ is significantly changed with
respect to the estimate $\hat{\mathcal{L}}_h$, the algorithm updates $\hat{\mathcal{L}}_h$ accordingly, re-initializes
315 the Q functions, the hypothesis reward machine \mathcal{H}_h , and the counterexample

Algorithm 3 Joint Perception and Learning Algorithm

```
1:  $\mathcal{H} \leftarrow \text{Initialize}(V)$ 
2:  $Q = \{q^v | v \in V\} \leftarrow \text{Initialize}(V)$ 
3:  $X \leftarrow \emptyset$ 
4:  $j \leftarrow 0$ 
5:  $\hat{\mathcal{L}}_h \leftarrow \hat{\mathcal{L}}_0$ ;
6: for episode  $n = 1, 2, \dots$  do
7:    $(\lambda, \rho, Q, \hat{\mathcal{L}}_j) \leftarrow \text{QRM.episode.mod}(\mathcal{H}, Q, \hat{\mathcal{L}}_h, \hat{\mathcal{L}}_j)$ 
8:   if  $\mathcal{H}(\lambda) \neq \rho$  then           // The trace  $(\lambda, \rho)$  is inconsistent with  $\mathcal{H}$ 
9:     add  $(\lambda, \rho)$  to  $X$            // Append the set  $X$  with the trace  $(\lambda, \rho)$ 
10:     $\mathcal{H} \leftarrow \text{Infer}(X)$        // Infer a new reward machine from  $X$ 
11:     $Q \leftarrow \text{Initialize}()$ 
12:   end if
13:   if  $\text{SignifChange}(\hat{\mathcal{L}}_h, \hat{\mathcal{L}}_j)$  then   //  $\hat{\mathcal{L}}_j$  has changed significantly
14:     with respect to  $\hat{\mathcal{L}}_h$ 
15:      $\mathcal{H} \leftarrow \text{Initialize}(V)$ 
16:      $Q \leftarrow \text{Initialize}(V)$ 
17:      $X \leftarrow \emptyset$ 
18:      $\hat{\mathcal{L}}_h \leftarrow \hat{\mathcal{L}}_j$ 
19:   end if
20: end for
```

set X . We provide the function $\text{SignifChange}()$, which evaluates the difference among $\hat{\mathcal{L}}_h$ and $\hat{\mathcal{L}}_j$ using a divergence test, in Section 5.2.

5.1. Information Processing

This section provides the $\text{BayesUpdate}(\hat{\mathcal{L}}_{j-1}, s')$ function (line 14 of Alg. 4) that updates the probability distribution $\hat{\mathcal{L}}_j$ based on the observation made by the agent at state s' .

When at state s' , the agent will receive new perception outputs according to the observation model $\mathcal{O}(s', \dots)$ for all states and atomic propositions. The agent then employs the observations to update its learned model of the envi-

Algorithm 4 QRM_episode_mod($\mathcal{H}, Q, \hat{\mathcal{L}}_h, \hat{\mathcal{L}}_j$)

Input: A reward machine $\mathcal{H} = (V, v_I, 2^{\mathcal{P}}, \mathbb{R}, \delta, \sigma)$, a set of q -functions $Q = \{q^v | v \in V\}$, two probabilistic beliefs $\hat{\mathcal{L}}_h, \hat{\mathcal{L}}_j$.

- 1: $s \leftarrow \text{InitialState}(); v \leftarrow v_I, \lambda \leftarrow []; \rho \leftarrow []$
- 2: **for** $1 \leq t < \text{eplength}$ **do**
- 3: $a \leftarrow \text{GetEpsilonGreedyAction}(q^v, s)$
- 4: $s' \leftarrow \text{ExecuteAction}(s, a)$ // Transition in l-MDP \mathcal{M} based on $\mathcal{T}(s', a, s)$
- 5: $v' \leftarrow \delta(v, \hat{L}(\hat{\mathcal{L}}_h, s'))$ // Transition in \mathcal{H} based on $\hat{\mathcal{L}}_h$
- 6: $r \leftarrow \text{ObserveReward}()$
- 7: $q^v(s, a) \leftarrow \text{Update}(r, s')$ // Update of q -function for state v of \mathcal{H}
- 8: **for** $\hat{v} \in V \setminus v$ **do**
- 9: $\hat{v}' \leftarrow \delta(\hat{v}, \hat{L}(\hat{\mathcal{L}}_h, s'))$ // Transition in \mathcal{H} based on $\hat{\mathcal{L}}_h$
- 10: $\hat{r} \leftarrow \sigma(\hat{v}, \hat{L}(\hat{\mathcal{L}}_h, s'))$ // Reward calculation from \mathcal{H} based on $\hat{\mathcal{L}}_h$
- 11: $q^{\hat{v}}(s, a) \leftarrow \text{Update}(\hat{r}, s')$ // Update of q -function for state \hat{v} of \mathcal{H}
- 12: **end for**
- 13: $j \leftarrow (n - 1) \cdot \text{eplength} + t$
- 14: $\hat{\mathcal{L}}_j \leftarrow \text{BayesUpdate}(\hat{\mathcal{L}}_{j-1}, s')$ // Update probabilistic belief \mathcal{L}_j
- 15: Append($\lambda, \hat{L}(\hat{\mathcal{L}}_h, s')$); Append(ρ, r)
- 16: $s \leftarrow s'; v \leftarrow v'$
- 17: **end for**
- 18: **return** ($\lambda, \rho, Q, \hat{\mathcal{L}}_j$)

ronment in a Bayesian approach. For ease of notation, let $\hat{\mathcal{L}}_j = \hat{\mathcal{L}}_j(s, p)$ and $\mathcal{O}(b) = \mathcal{O}(s', s, p, b)$. Given the prior belief of the agent $\hat{\mathcal{L}}_{j-1}$ and the received observations \mathcal{Z} , the posterior belief follows

$$Pr(s \models p | \mathcal{Z}(s', s, p) = \text{True}) = \frac{\hat{\mathcal{L}}_{j-1} \mathcal{O}(\text{True})}{\hat{\mathcal{L}}_{j-1} \mathcal{O}(\text{True}) + (1 - \hat{\mathcal{L}}_{j-1}) \mathcal{O}(\text{False})} \quad (2a)$$

$$Pr(s \models p | \mathcal{Z}(s', s, p) = \text{False}) = \frac{\hat{\mathcal{L}}_{j-1} (1 - \mathcal{O}(\text{True}))}{\hat{\mathcal{L}}_{j-1} (1 - \mathcal{O}(\text{True})) + (1 - \hat{\mathcal{L}}_{j-1}) (1 - \mathcal{O}(\text{False}))} \quad (2b)$$

for all $s \in S$ and $p \in \mathcal{AP}$. Depending on the truth value observed for p , $\text{BayesUpdate}(\hat{\mathcal{L}}_{j-1}, s')$ updates $\hat{\mathcal{L}}_j$ according to one of the aforementioned expressions. Besides, for any $P \subseteq \mathcal{AP}$, $\hat{\mathcal{L}}_j(s, P) = \prod_{p \in P} \hat{\mathcal{L}}_j(s, p)$.

325 5.2. Divergence Test on the Belief

This section provides the $\text{SignifChange}(\hat{\mathcal{L}}_h, \hat{\mathcal{L}}_j)$ function (line 13 of Alg. 3) that compares the distributions $\hat{\mathcal{L}}_h$ and $\hat{\mathcal{L}}_j$.

As mentioned before, if the agent's current estimate about the atomic propositions, encoded by the current belief $\hat{\mathcal{L}}_j$, has not changed significantly from its previous estimate (encoded by $\hat{\mathcal{L}}_h$), then the agent will continue executing the algorithm using the current $\hat{\mathcal{L}}_h$. However, if its estimate has changed significantly, the agent will update its belief and aim to infer a new reward machine (lines 13-18 of Alg. 3). The algorithm compares the two distributions, $\hat{\mathcal{L}}_h$ and $\hat{\mathcal{L}}_j$ via the function $\text{SignifChange}()$. The function uses the Jensen-Shannon divergence to quantify the change in the belief distribution between two consecutive time steps. We note here that alternative divergence methods would also work without altering the correctness of the algorithm (proven in Sec. 5.4). We choose the Jensen-Shannon divergence due to its symmetry, finite value, and its efficiency for generic distributions over atomic propositions. The cumulative Jensen-Shannon divergence over the states and the propositions can be expressed as

$$\begin{aligned} D_{\mathcal{JSD}}(\hat{\mathcal{L}}_h || \hat{\mathcal{L}}_j) &= \frac{1}{2} D_{\mathcal{KL}}(\hat{\mathcal{L}}_h || \hat{\mathcal{L}}_m) + \frac{1}{2} D_{\mathcal{KL}}(\hat{\mathcal{L}}_j || \hat{\mathcal{L}}_m) = \\ & \frac{1}{2} \sum_{s \in S} \sum_{p \in \mathcal{AP}} \hat{\mathcal{L}}_h(s, p) \log \frac{\hat{\mathcal{L}}_h(s, p)}{\hat{\mathcal{L}}_m(s, p)} + (1 - \hat{\mathcal{L}}_h(s, p)) \log \frac{1 - \hat{\mathcal{L}}_h(s, p)}{1 - \hat{\mathcal{L}}_m(s, p)} \\ & + \frac{1}{2} \sum_{s \in S} \sum_{p \in \mathcal{AP}} \hat{\mathcal{L}}_j(s, p) \log \frac{\hat{\mathcal{L}}_j(s, p)}{\hat{\mathcal{L}}_m(s, p)} + (1 - \hat{\mathcal{L}}_j(s, p)) \log \frac{1 - \hat{\mathcal{L}}_j(s, p)}{1 - \hat{\mathcal{L}}_m(s, p)} \end{aligned}$$

where $\hat{\mathcal{L}}_m = \frac{1}{2}(\hat{\mathcal{L}}_h + \hat{\mathcal{L}}_j)$ is the average distribution. One of the input parameters to the algorithm is a threshold γ_d on the above divergence; $\text{SignifChange}(\hat{\mathcal{L}}_h, \hat{\mathcal{L}}_j)$ returns **False** if $D_{\mathcal{JSD}}(\hat{\mathcal{L}}_h || \hat{\mathcal{L}}_j) < \gamma_d$ and **True** if $D_{\mathcal{JSD}}(\hat{\mathcal{L}}_h || \hat{\mathcal{L}}_j) \geq \gamma_d$, determining whether the agent will use its previous estimate $\hat{L}(\hat{\mathcal{L}}_h, \cdot)$ or update $\hat{\mathcal{L}}_h$ to guide the learning and reward machine inference.

330

5.3. Inference of Reward Machines

The goal of reward-machine inference is to find a reward machine \mathcal{H} that
335 is consistent with all the counterexamples in the sample set X , i.e., such that
 $\mathcal{H}(\lambda) = \rho$ for all $(\lambda, \rho) \in X$, for the current belief $\hat{\mathcal{L}}_h$. Unfortunately, such a task
is computationally hard in the sense that the corresponding decision problem
“given a sample X and a natural number $k > 0$, does a consistent Mealy machine
with at most k states exist?” is NP-complete [8, 38]. In this paper we follow
340 the approach of [8], which is to learn minimal consistent reward machines with
the help of highly-optimized SAT solvers [39, 40, 41]. The underlying idea is to
generate a sequence of formulas ϕ_k^X in propositional logic for increasing values
of $k \in \mathbb{N}$ (starting with $k = 1$) that satisfy the following two properties:

- ϕ_k^X is satisfiable if and only if there exists a reward machine with k states
345 that is consistent with X ;
- a satisfying assignment of the variables in ϕ_k^X contains sufficient informa-
tion to derive such a reward machine.

By increasing k by one and stopping once ϕ_k^X becomes satisfiable (or by
using a binary search), we construct an algorithm that learns a minimal reward
350 machine that is consistent with the given sample.

As stressed in Example 3.1, a reward machine that encodes R on $\hat{\mathcal{L}}_h$ might
not exist and the inference procedure might end up creating an arbitrarily large
Mealy machine and yielding very high computation times (trying to compute a
Mealy machine that does not exist). To prevent such cases, one can limit the
355 maximum allowable number of states of the Mealy machine and the maximum
number of episodes used to generate the counterexamples, and use the last valid
inferred Mealy machine (or the initial estimate if there is no valid one). We note
that such a technique requires knowledge of an upper bound on the number of
states of the reward machine. Practically, this bound could be a conservatively
360 large number that could be derived from the total number of atomic propositions
 \mathcal{AP} . Nevertheless, large bounds could lead to high computational times in cases

the algorithm is trying to infer a reward machine that does not exist due to inaccuracies in $\hat{\mathcal{L}}_h$. Such complexity stems from the inherent scalability issues of the employed SAT-based optimizers and can be alleviated by modifying the underlying algorithms to obtain polynomial-time heuristics [8, Sec. 4.3]. In the experimental results, the proposed algorithm, such cases did not prevent the successful convergence of the proposed algorithm.

A sufficient condition for a reward machine that encodes R on $\hat{\mathcal{L}}_h$ to exist is sufficient accuracy of $\hat{\mathcal{L}}_h$, i.e., $\hat{L}(\hat{\mathcal{L}}_h, s) = \{p \in \mathcal{AP} : \hat{\mathcal{L}}_h(s, p) \geq 0.5\} = L_G(s)$, for all $s \in S$ (see (1)). Intuitively, such a property implies that the most probably atomic-proposition estimates are the correct ones. We show later that such a property can be achieved with an unambiguous observation model and sufficient exploration of the environment. In case $\hat{L}(\hat{\mathcal{L}}_h, s) \neq L_G$, it is not straightforward to derive conditions for the existence of a reward machine that encodes R on $\hat{\mathcal{L}}_h$. Based on intuition, one could argue that small inaccuracies in $\hat{\mathcal{L}}_h$, such as wrong labels in just one of the environment states, could lead to such a property. Consider, however, the grid example of Fig. 2 and the case where $\hat{L}(\hat{\mathcal{L}}_h, s_{ij}) = L_G(s_{ij})$ for all $i \in \{1, 2\}$, $j \in \{1, \dots, 5\}$, except for $\hat{L}(\hat{\mathcal{L}}_h, s_{14}) = \{“X”\} \neq L_G(s_{14}) = \{\}$, i.e., the sensor falsely detects an obstacle at state s_{14} . In that case, there is no reward machine that encodes R on $\hat{\mathcal{L}}_h$: if such a reward machine \mathcal{A} existed, its reward sequence $\mathcal{A}(\hat{\ell}_0, \dots, \hat{\ell}_k)$, given an input label sequence from any trajectory $s_0 a_0 \dots a_k s_{k+1}$ and based on $\hat{\mathcal{L}}_h$, would need to match the observed reward sequence. Consider, however, the trajectory $s_{11} s_{12} s_{13} s_{14} s_{24}$, which gives the *observed* label sequence $\{“c”\}, \{“X”\}, \{“o”\}$. In order for $\mathcal{A}(\{“c”\}, \{“X”\}, \{“o”\})$ to match the observed reward sequence, which is 0, 0, 0, 1, \mathcal{A} would have to neglect the obstacle label “X”. In that case, however, the reward machine would output a reward sequence of 0, 0, 0, 1 also for the trajectory $s_{11} s_{12} s_{22} s_{23} s_{24}$, which gives the observed label sequence $\{“c”\}, \{“X”\}, \{“X”\}, \{“o”\}$, i.e., $\mathcal{A}(\{“c”\}, \{“X”\}, \{“X”\}, \{“o”\}) = 0, 0, 0, 1$. The *observed* reward sequence, however, is a sequence of zeros due to encounter of obstacles. Hence, there is no reward machine that encodes R on $\hat{\mathcal{L}}_h$.

Similarly, one might argue that if $\hat{L}(\hat{\mathcal{L}}_h, s)$ is not equal to L_G , then there is no reward machine that encodes R on $\hat{\mathcal{L}}_h$. Consider, however, the scenario where the task is to bring coffee *or* water to the office without encountering any obstacles and the extra atomic proposition “water” is found in s_{14} , i.e., $L_G(s_{14}) = \{“w”\}$. Further assume that $\hat{L}(\hat{\mathcal{L}}_h, s_{ij}) = L_G$ for all $i \in \{1, 2\}$, $j \in \{1, \dots, 5\}$, except for $\hat{L}(\hat{\mathcal{L}}_h, s_{14}) = \{“c”\} \neq L_G(s_{14}) = \{“w”\}$, i.e., the sensor falsely detects “coffee” instead of “water” in s_{14} . In that case, it can be verified that the observed reward sequence of any trajectory matches the sequence $\mathcal{A}(\hat{\ell}_0, \dots, \hat{\ell}_k)$, where \mathcal{A} is a reward machine similar to the one Fig. 3. Therefore, the existence of a reward machine that encodes R on $\hat{\mathcal{L}}_h$ depends on the underlying task and its relation to the atomic propositions.

5.4. Convergence in the limit

In this section, we establish the correctness of Algorithm 3. More specifically, we provide theoretical guarantees on (i) the convergence of the hypothesis reward machine \mathcal{H} to the actual reward machine \mathcal{A} , and (ii) the asymptotic learning of the optimal policy that achieves accomplishment of the task encoded by \mathcal{A} .

The aforementioned guarantees are based on two sufficient conditions. First, the estimate $\hat{L}(\hat{\mathcal{L}}_h, \cdot)$ becomes identical with the ground-truth labelling function L_G after a finite number of episodes. Second, the length of each episode is larger than $2^{|\mathcal{M}|+1}(|\mathcal{A}| + 1) - 1$, where \mathcal{A} is the unknown reward machine that encodes the task to be accomplished. The first condition is tightly connected to the observation model \mathcal{O} of the agent (see Def. 3) and its effect on the Bayesian updates (2). Intuitively, when there is unambiguity in the observation model, i.e., it is sufficiently accurate or inaccurate, the Bayesian updates (2), along with consistent observations, lead to accurate estimation of L_G . In the next section, we show that such unambiguity implies $\mathcal{O}(s_1, s_2, p, \text{True}) = Pr(\mathcal{Z}(s_1, s_2, p) = \text{True} | s_2 \models p) > \mathcal{O}(s_1, s_2, p, \text{False}) = Pr(\mathcal{Z}(s_1, s_2, p) = \text{False} | s_2 \not\models p)$ or $\mathcal{O}(s_1, s_2, p, \text{True}) = Pr(\mathcal{Z}(s_1, s_2, p) = \text{True} | s_2 \models p) < \mathcal{O}(s_1, s_2, p, \text{False}) = Pr(\mathcal{Z}(s_1, s_2, p) = \text{False} | s_2 \not\models p)$ and we illustrate that, in such cases, observations that are consistent with the

truth values of \mathcal{AP} and the observation model lead to convergence of $\hat{L}(\hat{\mathcal{L}}_i, \cdot)$ to $L_G(\cdot)$. Nevertheless, such conditions are only sufficient for accurate estimation
 425 of L_G . The experiments of Section 6 illustrate that the proposed algorithm learns how to accomplish the underlying task even with random observation models.

The convergence results are given in the following Theorem, whose proof can be found in the Appendix:

430 **Theorem 1.** *Let $\mathcal{M} = (S, s_I, A, \mathcal{T}, R, \mathcal{AP}, \hat{L}, \gamma)$ be an l -MDP and \mathcal{A} be a reward machine that encodes the reward function R on the ground truth. Further, assume that there exists a constant $n_f > 0$ such that $\hat{L}(\hat{\mathcal{L}}_h, \cdot) = L_G(\cdot)$, for all episodes $n > n_f$ of Algorithm 3². Then, Algorithm 3, with eplength $\geq 2^{|\mathcal{M}|+1}(|\mathcal{A}| + 1) - 1$, converges almost surely to an optimal policy in the limit.*

435

Remark 2. *It can be concluded that the main differences of the proposed algorithm with respect to JIRP [8] consist of the updates of the labelling-function estimate (Sec. 5.1) and the divergence test (Sec. 5.2). As mentioned before, however, these techniques can be applied to more general RL frameworks that
 440 rely on labelling functions [42, 31]. Such an application follows a procedure similar to the extension of the JIRP algorithm in the proposed methodology. More specifically, all the operations of the respective algorithms can be implemented based on a labelling-function estimate $\hat{\mathcal{L}}_h$. Since the algorithms in these works rely on exploration of the environment by the agents, as commonly done in RL,
 445 sensor observations can be used to update a running estimate $\hat{\mathcal{L}}_i$ based on an observation model, as done in Sec. 5.1. Then, a divergence test can determine the update of $\hat{\mathcal{L}}_h$ with $\hat{\mathcal{L}}_i$, as done in Sec. 5.2.*

²Note that this implies that $\{p \in \mathcal{AP} : \hat{L}_h(s, p) \geq 0.5\} = L_G(s)$ for all $s \in S$.

5.5. Accurate estimation of L_G

450 This section provides further insights on the relation of the Bayesian updates (2) and the observation model \mathcal{O} towards the accurate estimation of L_G by \hat{L} given in (1). For ease of exposition, we assume that the agent receives observations about atomic proposition $p \in \mathcal{AP}$ in a state $s \in S$ only when it is located at that state. In the following, we use $\mathcal{O}(s, s, p, \text{True}) = \mathcal{O}_T$ and
 455 $\mathcal{O}(s, s, p, \text{False}) = \mathcal{O}_F$.

We derive a closed-form expression for the probabilistic belief $\hat{\mathcal{L}}_N(s, p)$ regarding an atomic proposition $p \in \mathcal{AP}$ in a state $s \in S$ after N observations. From these N observations, we consider that N_T observations correspond to the detection $\mathcal{Z}(s, s, p) = \text{True}$ and N_F observations correspond to $\mathcal{Z}(s, s, p) = \text{False}$, with $N = N_T + N_F$. We first show that $\hat{\mathcal{L}}_N(s, p)$ will be the same regardless of the order of the observations. Indeed, let $\hat{\mathcal{L}}_i(s, p)$ be the belief at the i th iteration, with $i \geq 0$. Assume that the agent then obtains first an observation $\mathcal{Z}(s, s, p) = \text{True}$ followed by $\mathcal{Z}(s, s, p) = \text{False}$. According to (2), it holds that

$$\begin{aligned} \hat{\mathcal{L}}_{i+1} &= \frac{\hat{\mathcal{L}}_i \mathcal{O}_T}{\hat{\mathcal{L}}_i \mathcal{O}_T + (1 - \hat{\mathcal{L}}_i) \mathcal{O}_F} \\ \hat{\mathcal{L}}_{i+2} &= \frac{\hat{\mathcal{L}}_{i+1} (1 - \mathcal{O}_T)}{\hat{\mathcal{L}}_{i+1} (1 - \mathcal{O}_T) + (1 - \hat{\mathcal{L}}_{i+1}) (1 - \mathcal{O}_F)} \\ &= \frac{\hat{\mathcal{L}}_i \mathcal{O}_T (1 - \mathcal{O}_T)}{\hat{\mathcal{L}}_i \mathcal{O}_T (1 - \mathcal{O}_T) + \mathcal{O}_F (1 - \hat{\mathcal{L}}_i) (1 - \mathcal{O}_F)} \end{aligned} \quad (3)$$

where we omit the argument (s, p) for brevity. On the contrary, assume that the agent obtains first an observation $\mathcal{Z}(s, s, p) = \text{False}$ followed by $\mathcal{Z}(s, s, p) = \text{True}$. According to (2), it holds that

$$\begin{aligned} \hat{\mathcal{L}}_{i+1} &= \frac{\hat{\mathcal{L}}_i (1 - \mathcal{O}_T)}{\hat{\mathcal{L}}_i (1 - \mathcal{O}_T) + (1 - \hat{\mathcal{L}}_i) (1 - \mathcal{O}_F)} \\ \hat{\mathcal{L}}_{i+2} &= \frac{\hat{\mathcal{L}}_{i+1} \mathcal{O}_T}{\hat{\mathcal{L}}_{i+1} \mathcal{O}_T + (1 - \hat{\mathcal{L}}_{i+1}) \mathcal{O}_F} \\ &= \frac{\hat{\mathcal{L}}_i \mathcal{O}_T (1 - \mathcal{O}_T)}{\hat{\mathcal{L}}_i \mathcal{O}_T (1 - \mathcal{O}_T) + \mathcal{O}_F (1 - \hat{\mathcal{L}}_i) (1 - \mathcal{O}_F)}, \end{aligned} \quad (4)$$

from which we conclude that $\hat{\mathcal{L}}_{i+2}$ is the same in both cases. By induction, after

$N = N_T + N_F$ observations, $\hat{\mathcal{L}}_N$ will be the same regardless of the sequence of N_T and N_F observations.

We now proceed to derive a general closed-form expression for $\hat{\mathcal{L}}_N$. Since the sequence of observations does not alter such an expression, we assume, without loss of generality, that the agent obtains first N_T observations followed by N_F observations. After N_T observations $\mathcal{Z}(s, s, p) = \text{True}$, iterative application of (2a) yields

$$\hat{\mathcal{L}}_{N_T} = \frac{\alpha^{N_T}}{\hat{\mathcal{L}}_0^{-1} + (\alpha - 1) \sum_{i=0}^{N_T-1} \alpha^i}$$

where we denote $\alpha = \frac{\mathcal{O}_T}{\mathcal{O}_F}$. After another N_F observations $\mathcal{Z}(s, s, p) = \text{False}$, iterative application of (2b) yields

$$\hat{\mathcal{L}}_N = \frac{(1 - \alpha \mathcal{O}_F)^{N_F}}{D_F}$$

where

$$\begin{aligned} D_F = & (1 - \mathcal{O}_F)^{N_F} \hat{\mathcal{L}}_{N_T}^{-1} + \mathcal{O}_F (1 - \alpha) (1 - \alpha \mathcal{O}_F)^{N_F-1} \\ & + \mathcal{O}_F (1 - \mathcal{O}_F) (1 - \alpha) \sum_{i=0}^{N_F-2} (1 - \alpha \mathcal{O}_F)^i (1 - \mathcal{O}_F)^{N_F-2-i} \end{aligned}$$

Substituting $\hat{\mathcal{L}}_{N_T}$ in $\hat{\mathcal{L}}_N$ yields

$$\hat{\mathcal{L}}_N = \frac{\alpha^{N_T} (1 - \alpha \mathcal{O}_F)^{N_F}}{D_N} \quad (5)$$

where

$$\begin{aligned} D_N = & \alpha^{N_T} \mathcal{O}_F (1 - \alpha) (1 - \alpha \mathcal{O}_F)^{N_F-1} + (1 - \mathcal{O}_F)^{N_F} \left(\hat{\mathcal{L}}_0^{-1} + (\alpha - 1) \sum_{i=0}^{N_T-1} \alpha^i \right) \\ & + \alpha^{N_T} \mathcal{O}_F (1 - \mathcal{O}_F) (1 - \alpha) \sum_{i=0}^{N_F-2} (1 - \alpha \mathcal{O}_F)^i (1 - \mathcal{O}_F)^{N_F-2-i} \end{aligned}$$

Note that, the case $\mathcal{O}_T = \mathcal{O}_F$, i.e., $\alpha = 1$, leads to $\hat{\mathcal{L}}_N = \hat{\mathcal{L}}_0$ for all $N > 0$. Such cases entail an *ambiguous* observation model and are equivalent to cases without the updates (2); the probability of observing $\mathcal{Z}(s, s, p) = \text{True}$ is the same regardless of whether $s \models p$ or $s \not\models p$ and the updates (2) fail to estimate L_G .

In cases of an unambiguous observation model, however, satisfying $\mathcal{O}_T > \mathcal{O}_F$ or $\mathcal{O}_T < \mathcal{O}_F$, observations that are consistent with \mathcal{O} and the truth value of p in s lead to accurate estimation of L_G . By consistent observations, we mean
465 observations satisfying $N_T = \mathcal{O}_T N$ and $N_F = (1 - \mathcal{O}_T)N$ in case $s \models p$ and $N_T = \mathcal{O}_F N$ and $N_F = (1 - \mathcal{O}_F)N$ in case $s \not\models p$.

In order to formally prove the aforementioned conjecture, one would ideally need to show that $\hat{\mathcal{L}}_{N+1} > \hat{\mathcal{L}}_N$ in case $s \models p$ and $\hat{\mathcal{L}}_{N+1} < \hat{\mathcal{L}}_N$ in case $s \not\models p$. The
470 complexity of the expression (5), however, renders such a procedure significantly difficult. Instead, we plot $\hat{\mathcal{L}}_N$ in (5) as a function of N for different values of the observation model satisfying $\mathcal{O}_T > \mathcal{O}_F$ and $\mathcal{O}_T < \mathcal{O}_F$. In particular, we consider four different cases formed by $\mathcal{O}_T > \mathcal{O}_F$, $\mathcal{O}_T < \mathcal{O}_F$ and $s \models p$, $s \not\models p$. In each of these cases, we generate 500 different combinations of \mathcal{O}_T and \mathcal{O}_F
475 randomly from a uniform distribution. In the cases where $\mathcal{O}_T > \mathcal{O}_F$, we choose \mathcal{O}_T randomly in the interval $(0, 1)$ and \mathcal{O}_F in the interval $(0, \mathcal{O}_T)$. Reversely, in the cases where $\mathcal{O}_T < \mathcal{O}_F$, we choose \mathcal{O}_F randomly in the interval $(0, 1)$ and \mathcal{O}_T in the interval $(0, \mathcal{O}_F)$. In all cases we choose $\hat{\mathcal{L}}_0$ randomly from a uniform distribution in $(0, 1)$. To obtain observations consistent with the observation
480 model, we further choose $N_T = \mathcal{O}_T N$, $N_F = (1 - \mathcal{O}_T)N$ when $s \models p$ and $N_T = \mathcal{O}_F N$, $N_F = (1 - \mathcal{O}_F)N$ when $s \not\models p$. The results for the four different cases are depicted in Fig. 4. One can conclude from the figure that the $\hat{\mathcal{L}}_N(s, p)$ converges to 1 and 0 for when $s \models p$ and $s \not\models p$, respectively.

485 5.6. Extension to continuous spaces

In cases where the state S or action space A is continuous, the tabular q-updates of the QRM algorithm (Alg. 4) cannot be implemented. In such cases, one can employ a parametrization $\theta^a \in \mathbb{R}^\ell$, for some $\ell \in \mathbb{N}$, in order to approximate the q function for each state of the hypothesis reward machine.
490 In particular, we define $q^v(s, a | \theta^a) : S \times A \times \bar{V} \times \mathbb{R}^\ell \rightarrow \mathbb{R}$ as a parametric approximation of $q^v(s, a)$ for each state-action pair $(s, a) \in S \times A$ and reward-machine state $v \in \bar{V}$; \bar{V} is defined as $\bar{V} = \{1, \dots, \bar{V}\}$, where \bar{V} is an upper bound on the number of reward-machine states. Such an approximation is typically

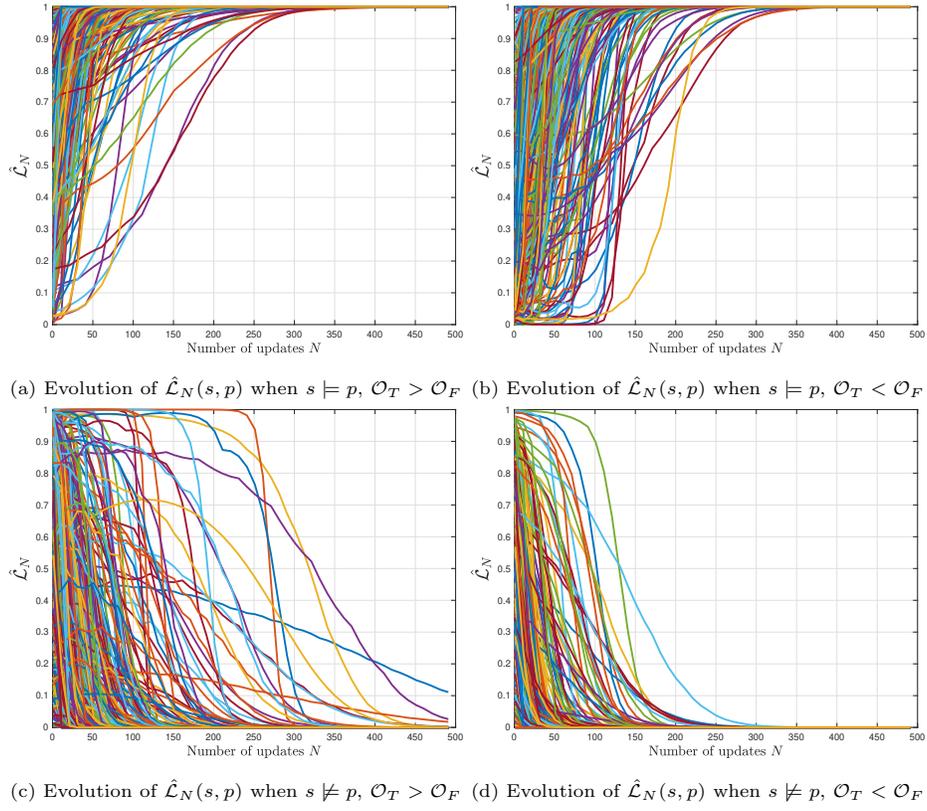


Figure 4: Evolution of $\hat{\mathcal{L}}_N(s, p)$ according to the Bayesian updates (2).

modelled using deep neural networks. The q -learning update of Alg. 4 (lines 7
 495 and 11) is then performed by minimizing the loss function $\mathbb{E}_{s \sim pr^\beta} [(q^v(s, a | \theta^q) - y)^2]$ [31], where pr^β is the probability distribution of state-visit over S under an arbitrary stochastic policy β and $y = R(s_0 \dots s') + \gamma \max_{a' \in A} q^v(s', a' | \theta^q)$. One could also append an intrinsic reward r^{in} to the aforementioned update; such an intrinsic reward takes positive values when the agent observes a new label
 500 during exploration in an episode and zero otherwise and has been shown to aid with learning [31, 42].

We note that, although the underlying spaces are continuous, the atomic propositions naturally induce a discrete abstraction $S_A = \{S_{A_1}, \dots, S_{A_M}\}$ of the continuous state space S . Such an abstraction is formed by grouping states
 505 where the same atomic propositions hold true. That is, given $s, s' \in S$ and $S_{A_i} \in S_A$ a discrete state, it holds $s, s' \in S_{A_i}$ if and only if there exists a subset $P \subset \mathcal{AP}$ of atomic propositions satisfying $s \models p$ and $s' \models p$ for all $p \in P$. Therefore, the ground-truth labelling function $L_G : S \rightarrow 2^{\mathcal{AP}}$ is naturally extended from S to S_A , i.e., $L_G : S_A \rightarrow 2^{\mathcal{AP}}$. Finally, we note that one could use
 510 appropriate tools, such as the options framework [42], to abstract the continuous state and action spaces to discrete ones and apply the proposed algorithm.

6. Experimental Results

We test the proposed algorithm in three different environments, described in the following. In all environments, we employ the RC2 SAT solver [44] from
 515 the PySAT library [45] towards reward-machine inference. All experiments were conducted on a Lenovo laptop with 2.70-GHz Intel i7 CPU and 8-GB RAM.

6.1. HouseExpo floor plans

We first test the algorithm in floor plan environments based on the indoor layouts collected in the HouseExpo dataset [43]. HouseExpo consists of indoor
 520 layouts, built from 35,126 houses with a total of 252,550 rooms. There are 25 room categories, such as bedroom, garage, office, boiler room, etc. The dataset provides bounding boxes for the rooms with their corresponding categories,

along with 2D images of layouts. We pre-process these layouts to generate grid-world-like floor plans, where every grid is labelled with the type of the room it is in (see Fig. 5). We use these labels as atomic propositions to define tasks encoded as reward machines. In particular, we use the set of atomic propositions $\mathcal{AP} = \{\text{kitchen, bathroom, bedroom, office, indoor}\}$.

In the experiments, an episode terminates when the agents receives a non-zero reward or the maximum number of steps allowed in one episode is exceeded. This maximum number of steps per episode is 1000, whereas the maximum number of total training steps is 500,000. Furthermore, we set the Jensen-Shannon divergence threshold γ_d of Section 5.2 to 10^{-5} , and Algorithm 4 selects a random action with probability $\varepsilon_a = 0.3$.

We compare the following experimental settings regarding the observation model and the belief updates:

1. Time-varying Bayesian Observation Model (TvBNN): We use a number of training HouseExpo layouts to train a Bayesian neural network, which the algorithm uses as an observation model in test layouts for the experiments. The algorithm samples observation probabilities of the atomic propositions' truth value from the neural network at the beginning of every training episode.
2. Fixed Bayesian Observation Model (FiBNN): We use a number of training HouseExpo layouts to train a Bayesian neural network, which the algorithm uses as an observation model in test layouts for the experiments. The algorithm samples observation probabilities of the atomic propositions' truth value from the neural network only once, at the beginning of training.
3. Random Observation Model: The algorithm uses an observation model whose observation probabilities are uniformly sampled from the interval $[0.1, 0.9]$.
4. Random Observation Model-2: The algorithm uses an observation model whose observation probabilities are uniformly sampled from the interval

[0.4,0.5].

5. No Belief Updates: The algorithm uses a belief function that is randomly
555 initialized, without updating it.

The Bayesian observation model in the first two settings is a neural network of three densely connected layers with flipout estimators [46]. The loss function is the expected lower-bound loss, and a combination of Kullback-Leibler divergence and categorical cross entropy. Given the information about a grid
560 in a floor plan, the Bayesian neural network (BNN) predicts category probabilities of the room where the grid is located. The input vector consists of the 2D coordinates of the grid, the size and the number of neighbours of the room. The output is the predicted probabilities of the room type. The training set consists of 3,000 pre-processed layouts from the HouseExpo dataset, which we
565 select based on the floor plan size and the number of different room categories. After pre-processing, to eliminate variance with respect to the initial position of grids, we rotate every training layout 3 times by 90 degrees, and then add each rotated version to the training set. In all aforementioned cases, we simulate the observations to be consistent with the respective observation model,
570 i.e., observe that $\mathcal{Z}(s_1, s_2, p) = \text{True}$ and $\mathcal{Z}(s_1, s_2, p) = \text{False}$ with probabilities $\mathcal{O}(s_1, s_2, p, \text{True})$ and $1 - \mathcal{O}(s_1, s_2, p, \text{True})$, respectively, when $s_2 \models p$, and with probabilities $\mathcal{O}(s_1, s_2, p, \text{False})$ and $1 - \mathcal{O}(s_1, s_2, p, \text{False})$, respectively, when $s_2 \not\models p$.

For the reward-machine-inference part of the algorithm, we limit the maximum
575 allowable number of states for the hypothesis reward machine to 4. Similarly, we limit the maximum allowable size of the trace sets, which is the number of episodes whose traces are recorded, to 20. In case the inference exceeds the allowed number of states, the process is stopped and the training continues with the last valid hypothesis reward machine until the next inference step.

580 We first specify the following task to the agent: $\phi_1 = \text{“Go to bedroom or office, and then to the kitchen”}$, which is encoded by the reward machine shown in Fig. 6.

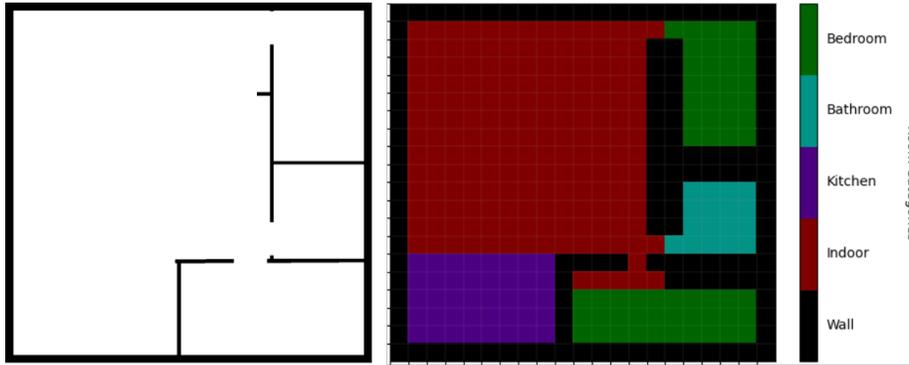


Figure 5: Left: Sample indoor layout from HouseExpo dataset; Right: Pre-processed sample layout

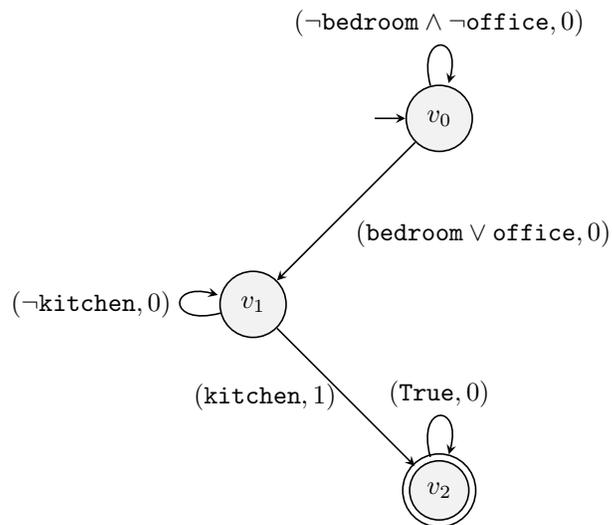


Figure 6: Ground truth reward machine of the task ϕ_1 .

We execute the proposed algorithm in 75 test floor plans. For each one of the first four experimental settings, where updates according to an observation model are performed, the learning agent is trained 5 times in each test floor plan, each one corresponding to a different random seed, leading to a total of 375 runs. For the last setting that does not include belief updates, we removed 12 test floor plans where the inference procedure took unreasonably long time to finish, resulting in 63 test floor plans. This is attributed to the fact that the belief functions are randomly initialized and not updated throughout the training. For this last setting, the learning agent is also trained 5 times in each test floor plan, each one corresponding to a different random seed, leading to a total of 290 runs. In each one of the aforementioned runs, we perform an evaluation episode every 100 training steps and register the obtained rewards. Figure 7 and Table 1 show the results. In particular, Fig. 7a depicts the progression of the rewards during training for the time-varying BNN (TvBNN), separated in the ones belonging to the 25th and 75th percentiles, and the ones belonging to the 10th and 90th percentiles; The figure also depicts the median values with solid line. We omit the plots for the other observation models since they converge after a similar amount of steps, as can be verified by Table 1. In particular, Table 1 shows the training steps required for convergence for the 25th (Q1), 50th (Q2), and 75th (Q3) percentiles of the attained rewards for all cases of observation models. More specifically, we compute the z^{th} percentile as follows: for each training step, we write the obtained rewards for all the 375 runs in increasing order and calculate the z^{th} position; subsequently, the training steps for the z^{th} percentile are the training steps required for this position to become 1. The table also depicts the ratio of successful inferences of reward machines (RS), and the average number of belief updates (BU) of $\hat{\mathcal{L}}_h$ after the JSD divergence test.

Figures 7b-7d demonstrate the progression of belief updates throughout the training; Fig. 7b shows the number of training runs (y-axis) in which the belief function $\hat{\mathcal{L}}_h$ is updated for the k^{th} time (x-axis). For instance, in the time-varying BNN (“TvBNN”) setting, the learning agent updates its belief for the

third time in around 20 runs out of 375 training runs (5 runs per every floor
 615 plan where a reward machine is successfully inferred). Fig. 7c shows the number
 of the step (y-axis) in which the belief function is updated for the k^{th} time on
 average (x-axis), and Fig. 7d displays the average Jensen-Shannon divergence
 value that led to the k^{th} belief update. The time-varying BNN (“TvBNN”)
 setting allows the learning agent to obtain a belief function that represents the
 620 environment well in a couple of steps. When sampling observation probabilities
 once, as in the fixed BNN (“FiBNN”) setting, it slows down the process since
 the sampled observation might not be informative. Regarding the random ob-
 servation model, due to the wide uniform distribution range, the agent samples
 probabilities that are close to zero or one, resulting to a more unambiguous
 625 observation model; that is, the agent is more “certain” that an atomic proposi-
 tion won’t or will be observed, respectively. This is the reason why the second
 random observation model performs more belief updates until the divergence
 score drops below the threshold. Overall, it is expected to see that, in the BNN
 settings, the algorithm performs fewer belief updates compared to the random
 630 settings.

In view of the aforementioned results, one concludes that the labelling-
 function belief converges after much fewer steps than the ones needed for the
 learning of the optimal policy. Therefore, the different observation models used
 do not affect significantly the inference of the reward machine or the speed of
 635 the q-learning’s convergence to the optimal policy, as long as such observation
 models lead to accurate estimation of the ground-truth labelling function. This
 fact can be verified by the experiments of Sec. 5.4, where the labelling-function
 estimate converges after approximately 500 updates (training steps), which is
 much fewer than the thousands of steps needed for learning the optimal policy
 640 (see Fig. 7 and Table 1).

To test the capabilities of our method in a more complicated scenario, we
 set a second task as $\phi_2 =$ “Go to `kitchen` while avoiding the `bathroom`, and
 then go to the `bathroom` while avoiding the `bedroom`”, which is encoded by the

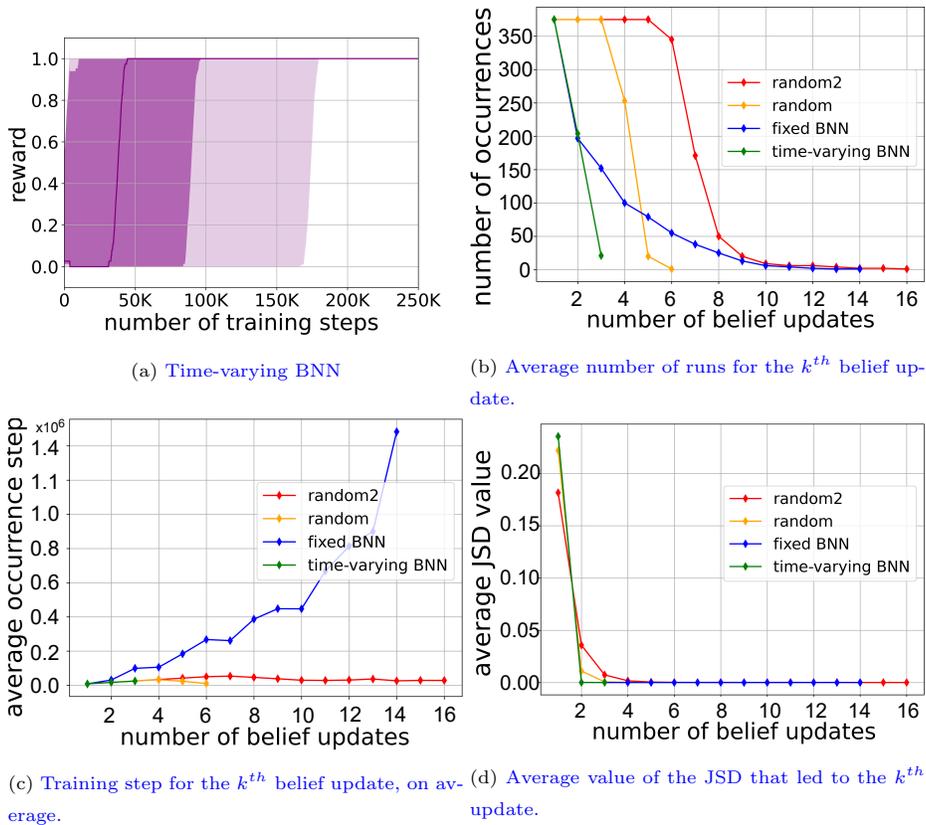


Figure 7: (a): Attained rewards of 5 independent runs for every 100 training steps in 75 test indoor layouts for task ϕ_1 using TvBNN: The darker region is bounded by 25th and 75th percentiles, while the lighter region is bounded by 10th and 90th. The solid line represents the median. (b)-(d): Progression of belief updates during training for task ϕ_1 .

Table 1: Convergence results of the proposed algorithm under different observation models in 5 independent runs for 75 test indoor layouts in the first HouseExpo floorplan experiment.

Settings	TvBNN	FiBNN	Random	Random2	No Update
Q1	95,800	109,600	103,000	107,800	500,000
Q2	44,400	48,900	40,100	45,700	189,900
Q3	9,900	22,700	11,200	17,200	35,000
RS	75/75	75/75	75/75	75/75	63/75
BU	1.6	2.79	3.73	6.64	-

645 reward machine shown in Fig. 8.

For this second task, we execute the proposed algorithm in 7 test floor plans with 5 different random seeds, resulting to a total of 35 runs. Figure 9 and Table 2 show the convergence results and belief updates, similar to the ones for ϕ_1 , for the first four settings; we omit the last one (no belief updates) since training 650 fails to yield an inferred reward machine. We conclude that all observation models lead to accurate estimation of the ground-truth labelling function and hence to successful inference of the reward machine and learning of the optimal policy. According to Table 2, the belief update patterns are similar to the ones observed for ϕ_1 . The average number of belief updates is smaller for the BNN 655 observation settings compared to the random observation settings. The TvBNN setting requires the fewest number of belief updates on average, whereas the Random-2 setting requires the highest number of updates. Finally, one can still conclude that the proposed algorithm outperforms the nominal setting, where no belief updates are performed.

660

6.2. Office world

We next test the proposed algorithm in the office-world environment shown in Fig. 10a. The environment is a 12×9 grid world with four atomic propositions $\mathcal{AP} = \{a, b, c, d\}$ corresponding to coffee a , mail b , obstacle c , and office

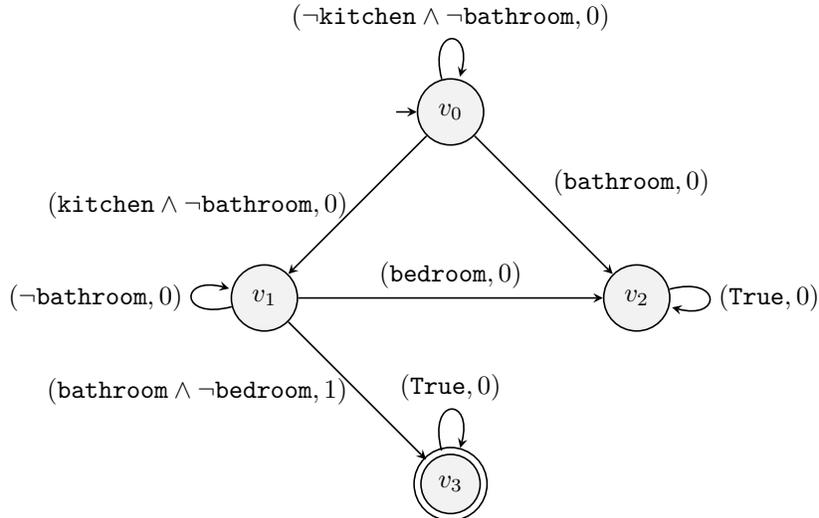


Figure 8: Ground truth reward machine of the task ϕ_2 .

Table 2: Convergence results of the proposed algorithm under different observation models in 5 independent runs for 7 test indoor layouts in the second HouseExpo floorplan experiment.

Settings	TvBNN	FiBNN	Random	Random2
Q1	255,500	263,100	341,100	231,100
Q2	146,700	141,800	137,100	122,200
Q3	81,500	80,000	74,200	91,100
BU	1.57	2.85	3.45	6.22

665 *d.* The task assigned to the agent is to get coffee, then mail, and eventually
arrive at the office without encountering the obstacle. The reward machine that
encodes the aforementioned task is given in Figure 10b. In the experiments, an
episode terminates when the agent reaches the accepting state or it exceeds the
maximum number of 2000 steps per episode. The maximum number of train-
670 ing steps is 1,500,000. We set the Jensen-Shannon divergence threshold γ_d to
 $\gamma_d = 10^{-5}$ and Algorithm 4 selects a random action with probability $\varepsilon_a = 0.3$.

We consider the following experimental settings, where we compare different
observation models as well as with standard learning-algorithms of the related

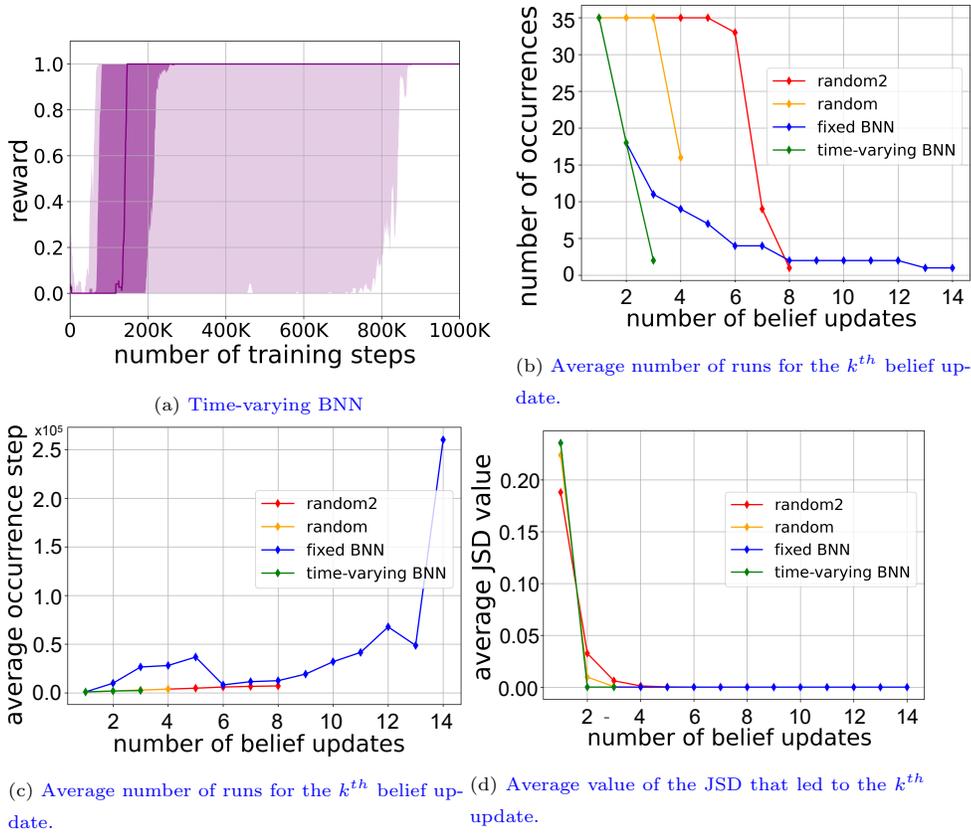
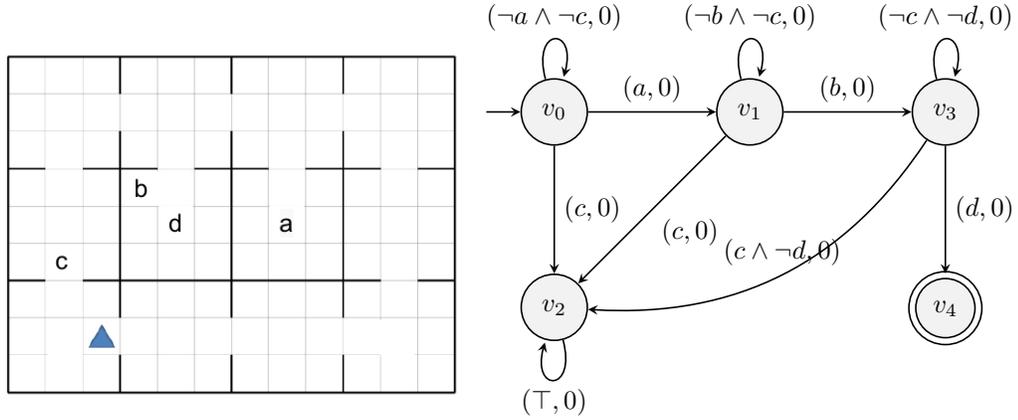


Figure 9: (a): Attained rewards of 5 independent runs for every 100 training steps in 7 test indoor layouts for task ϕ_2 using TvBNN: The darker region is bounded by 25th and 75th percentiles, while the lighter region is bounded by 10th and 90th. The solid line represents the median. (b)-(d): Progression of belief update during training for task ϕ_2 .



(a) Office world with four atomic propositions: coffee a , mail b , obstacle c , and office d . The blue triangle indicates the initial state of an episode. (b) Reward machine for the office-world environment. The task is to fetch coffee a , then mail b , and eventually arrive at the office d without encountering the obstacle c .

Figure 10: Office-world environment

literature. In what follows, we use the acronym JPL (Joint Perception and Learning) for the proposed algorithm.

1. True Observation model (JPL-true): The proposed algorithm uses an accurate and unambiguous observation model that gives the ground-truth probability of observing an atomic proposition at a state, i.e., $\mathcal{O}(s_1, s_2, p, \text{True}) = 1$ and $\mathcal{O}(s_1, s_2, p, \text{False}) = 0$.
2. False Observation model (JPL-false): The proposed algorithm uses an inaccurate and unambiguous observation model that gives false probabilities of observing an atomic proposition at a state, i.e., $\mathcal{O}(s_1, s_2, p, \text{True}) = 0$ and $\mathcal{O}(s_1, s_2, p, \text{False}) = 1$.
3. Random Observation model (JPL-random): The algorithm uses an observation model whose observation probabilities are uniformly sampled from the interval $[0.1, 0.9]$.
4. Random Observation model-2 (JPL-random2): The algorithm uses an observation model whose observation probabilities are uniformly sampled

690

from the interval $[0.4, 0.6]$.

695

5. Q-learning: We test the traditional Q-learning algorithm [47] with a randomly initialized labelling-function belief, without updating it or inferring a reward machine. Our implementation extends the original algorithm by providing a vector of binary values indicating whether each atomic proposition is observed by the agent (1) or not (0) according to its belief function.

700

6. DDQN: We test the double deep Q-learning (DDQN) algorithm [?], with a randomly initialized labelling-function belief, without updating it or inferring a reward machine. Our implementation augments the state space with the sequence of the last 200 labels observed according to the agent’s belief function. We use a neural network of 6 hidden layers, each consisting of 64 neurons.

705

In all aforementioned cases, we simulate the observations to be consistent with the respective observation model, i.e., observe that $\mathcal{Z}(s_1, s_2, p) = \text{True}$ and $\mathcal{Z}(s_1, s_2, p) = \text{False}$ with probabilities $\mathcal{O}(s_1, s_2, p, \text{True})$ and $1 - \mathcal{O}(s_1, s_2, p, \text{True})$, respectively, when $s_2 \models p$, and with probabilities $\mathcal{O}(s_1, s_2, p, \text{False})$ and $1 - \mathcal{O}(s_1, s_2, p, \text{False})$, respectively, when $s_2 \not\models p$.

710

We run the aforementioned cases in the office environment for 10 independent runs, each corresponding to a different random seed. In each one of the aforementioned runs, we perform an evaluation episode every 100 training steps and register the obtained rewards. The results are shown in Fig. 11 and Table 3. In particular, Fig. 11a shows the progression of rewards collected in test episodes for the random observation model (JPL-random) as well as for the Q-learning and DDQN algorithms. One concludes that the proposed algorithm successfully learns the optimal policy, whereas Q-learning and DDQN fail to learn a policy that accomplishes the task; such a result demonstrates that simply augmenting the state space with observed labels is not sufficient neither with a Q-table nor with a Q-function approximated by a neural network. We omit the reward-progression plots for the other cases of observation models, since

715

720 they yield similar convergence results with the first random observation model. Table 3 shows the training steps required for convergence for the 25th (Q1), 50th (Q2), and 75th (Q3) percentiles of the attained rewards for the aforementioned cases, as in Sec. 6.1. The table also depicts the average number of belief updates (BU) of $\hat{\mathcal{L}}_h$.

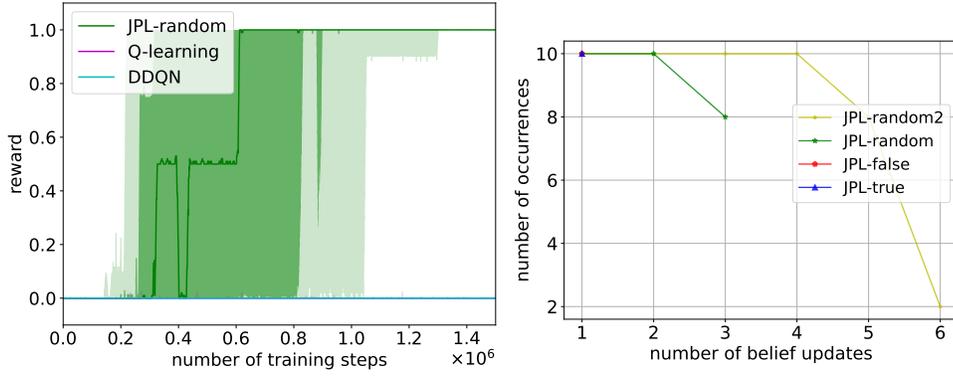
725 Figures 11b-11d further show how the information-processing mechanism of the proposed algorithm handles the belief updates. On average, JPL-true and JPL-false obtain a belief function that satisfies the JSD constraint, $\gamma_d = 10^{-5}$, in one episode, hence Figs. 11b, 11c, and 11d have only one data point for these algorithms; JPL-random2 satisfies the JSD threshold in 6 episodes, 3
730 more than what JPL-random, since the observation model is more ambiguous; i.e., it carries less information due to a narrow range of values used to draw probabilities ([0.4, 0.6] vs [0.1, 0.9]).

Table 3: Convergence results of the proposed algorithm under different observation models in 10 independent runs in the office-world environment.

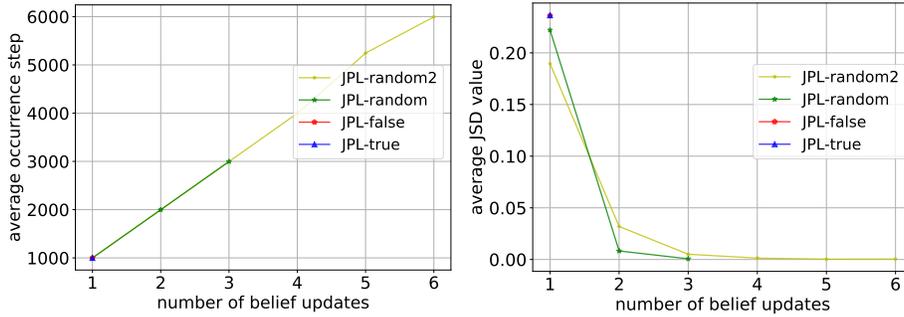
Setting	JPL-true	JPL-false	JPL-random	JPL-random2
Q1	1,029,000	1,019,800	958,500	1,003,700
Q2	776,000	776,000	818,300	737,600
Q3	412,100	423,300	315,600	412,100
BU	1	1	2.8	5

735 6.3. Craft-world Environment

Finally, we test the proposed algorithm in the craft-world environment shown in Fig. 12a. The environment is a 12×9 grid world with five atomic propositions $\mathcal{AP} = \{w, i, t, h, f\}$ corresponding to wood w , iron i , toolshed t , workbench h , and factory f . The task assigned to the agent is to build stairs by following the
740 instructions below:



(a) Progression of reward using (i) the proposed algorithm with a random observation model, (ii) a Q-learning algorithm, and a (iii) DDQN algorithm. (b) Average number of runs for the k^{th} belief update.



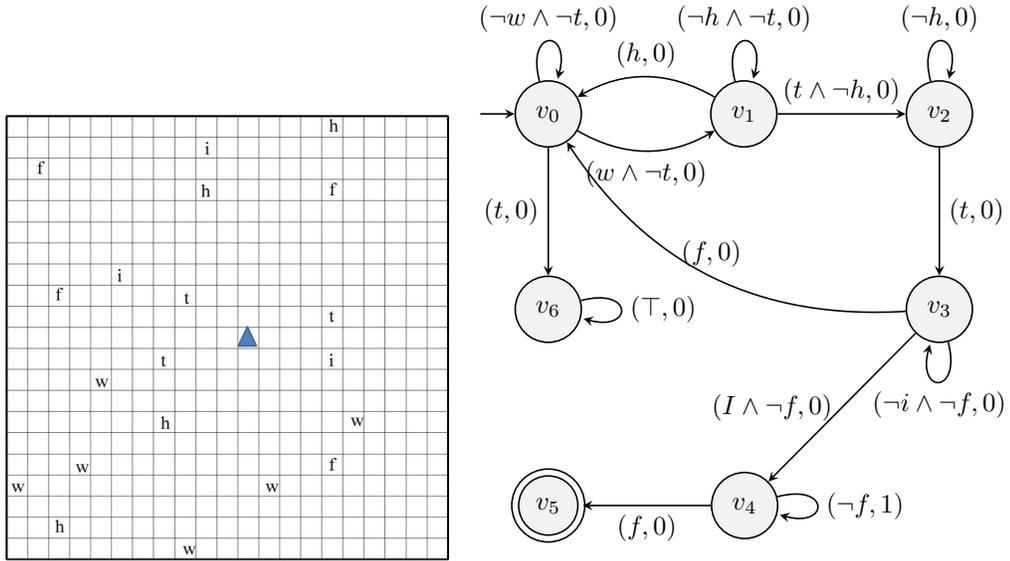
(c) Average number of runs for the k^{th} belief update. (d) Average value of the JSD that led to the k^{th} update.

Figure 11: (a): Attained rewards of 10 independent runs for every 100 training steps for the office-world task using the first random observation model: The darker region is bounded by 25th and 75th percentiles, while the lighter region is bounded by 10th and 90th. The solid line represents the median. (b)-(d): Progression of belief update during training for the office-world task.

1. Get wood w (toolshed t cannot be used before getting wood),
2. Use toolshed t (using workbench h before toolshed resets again),
3. Use workbench h ,
4. Get iron i (using factory f before collecting iron resets task), and
- 745 5. Use factory f .

The reward machine for such a task is shown in Fig. 12b. In the experiments, an episode terminates when the agent reaches the accepting state or it exceeds the maximum number of 400 steps per episode. The maximum number of training steps is 2,000,000. We set the Jensen-Shannon divergence threshold γ_d to $\gamma_d = 10^{-5}$ and Algorithm 4 selects a random action with probability $\varepsilon_a = 0.3$. We evaluate the same algorithmic observation-model cases as in the office-world case of Sec. 6.2.

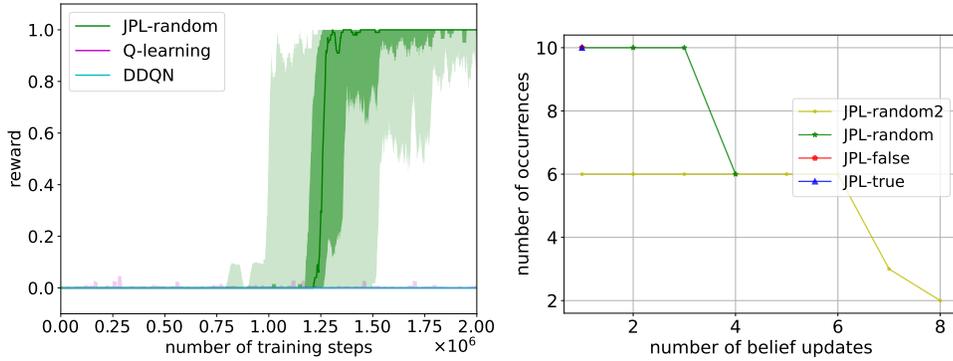
The results for 10 independent runs, each corresponding to a different random seed, are shown in Fig. 13. Figure 13a show the progression of rewards obtained in test episodes for the first random observation model of the proposed algorithm as well as the Q-learning and DDQN algorithms. We omit the reward-progression plots for other observation models since they achieve similar speed of convergence. We observe that, since the craft-world scenario involves a more complicated task and larger state space than the office-world scenario, the proposed algorithm requires more than 1.2 million training steps, in median values, to converge. Additionally, as in the office-world case, the Q-learning and DDQN algorithms fail to learn the policy that achieves the task. Furthermore, Table 4 shows the training steps required for convergence for the 25th (Q1), 50th (Q2), and 75th (Q3) percentiles of the attained rewards for the aforementioned cases, as in the previous cases. The table also depicts the average number of belief updates (BU) of $\hat{\mathcal{L}}_h$. Finally, Fig. 13b-13d show how the information-processing mechanism of the proposed algorithm handles the belief updates, where we observe a behaviour similar to the previous cases in Sec. 6.1 and 6.2.



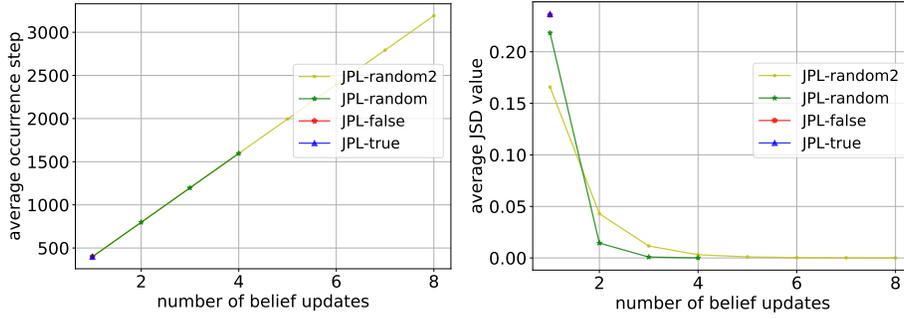
(a) Craft world with five atomic propositions: wood w , iron i , toolshed t , workbench h , and factory f . The blue triangle indicates the initial state of an episode.

(b) Reward machine for the craft-world environment. The task is to build stairs by getting wood w (toolshed t cannot be used before getting wood), then using toolshed t (using workbench h before toolshed resets) and workbench h consecutively. Finally, the agent should obtain iron i (using factory f before collecting iron resets task) and process all items in the factory f .

Figure 12: Craft-world environment.



(a) Progression of reward using (i) the proposed algorithm with a random observation model, (ii) a Q-learning algorithm, and a (iii) DDQN algorithm. (b) Average number of runs for the k^{th} belief update.



(c) Average number of runs for the k^{th} belief update. (d) Average value of the JSD that led to the k^{th} update.

Figure 13: (a): Attained rewards of 10 independent runs for every 100 training steps for the craft-world task using the first random observation model: The darker region is bounded by 25th and 75th percentiles, while the lighter region is bounded by 10th and 90th. The solid line represents the median. (b)-(d): Progression of belief update during training for the office-world task.

Table 4: Convergence results of the proposed algorithm under different observation models in 10 independent runs in the craft-world environment.

Setting	JPL-true	JPL-false	JPL-random	JPL-random2
Q1	1,937,600	1,907,600	1,950,800	1,985,800
Q2	1,435,600	1,242,800	1,531,200	1,825,600
Q3	1,196,400	1,129,600	1,256,000	1,484,400
BU	1	1	3.6	6.83

7. Conclusion and Discussion

770 We develop a reinforcement algorithm subject to reward-machine-encoded tasks and perceptual limitations. The algorithm holds probabilistic beliefs over the truth value of the atomic propositions and uses a hypothesis reward machine to estimate the reward machine that encodes the task to be accomplished. Both the beliefs and the hypothesis reward machine are updated by using the agent’s sensor measurements and the obtained rewards from the agent’s trajectories. 775 The algorithm uses the aforementioned beliefs and the hypothesis reward machine in a q-learning procedure to learn an optimal policy that accomplishes this task.

Currently, the theoretical guarantees, provided by Theorem 1, presuppose 780 that the algorithm obtains a good enough estimate of the atomic propositions’ value, as encoded in the ground truth labelling function L_G , i.e., $\hat{L}(\hat{\mathcal{L}}_h, \cdot) = L_G(\cdot)$. The definition of $\hat{L}(\hat{\mathcal{L}}_h, \cdot)$ in (1), based on the most probable atomic propositions p for state s , i.e., $\hat{\mathcal{L}}_h(s, p) \geq 0.5$, mildens such an assumption, which is further verified to hold by the experimental results. [Additionally, Sec. 5.4](#) illustrates that unambiguous observation models, i.e., when the probability of observing an atomic proposition when it is true is different than the probability of observing it when it is false, lead to accurate estimation of the ground-truth labelling function when the obtained observations are consistent with these observation models. Such an estimation is achieved even by the 785

790 time-varying observation models used in Sec. 6, illustrating the robustness of
the proposed approach to the choice of observation model. Along the same lines,
the experimental results show that the proposed algorithm, used with a large
variety of unambiguous observation models, achieves estimation of the ground-
truth labelling function in much fewer steps than the ones needed for inferring
795 the reward machine and learning an optimal policy. Therefore, we conclude that
the selection of observation model does not practically affect the convergence of
the algorithm to the optimal policy. Larger sets of atomic propositions would
naturally lead to more complex reward machines and a proportional increase in
convergence is expected both in the labelling-function estimation as well as the
800 reward-machine inference and learning of the optimal policy.

We further note that larger reward machines, possibly describing more com-
plex tasks yield higher computational times. In our implementation, inference of
reward machines with more than 7 states resulted in unreasonably high compu-
tation times. As mentioned in Sec. 5.3, modification of the inference algorithm
805 would be able to alleviate such issues. We argue, however, that reward machines
up to 7 states can describe significantly complicated tasks, such as Minecraft
[31] or Atari tasks [42]. In the future, we will consider uncertain observation
models that are not consistent with the obtained observations and cases where
the ground-truth values of the atomic propositions change with time.

810 **8. Acknowledgements**

This work was supported in part by ONR N00014-22-1-2254 and NSF 1652113.

References

- [1] M. E. Taylor, P. Stone, Cross-domain transfer for reinforcement learn-
ing, Proceedings of the 24th international conference on Machine learning
815 (2007) 879–886.
- [2] S. P. Singh, Reinforcement learning with a hierarchy of abstract models, in:

Proceedings of the National Conference on Artificial Intelligence, no. 10, Citeseer, 1992, p. 202.

- [3] T. D. Kulkarni, K. Narasimhan, A. Saeedi, J. Tenenbaum, Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation, *Advances in neural information processing systems* 29 (2016) 3675–3683.
- [4] D. Abel, D. Arumugam, L. Lehnert, M. Littman, State abstractions for life-long reinforcement learning, *International Conference on Machine Learning* (2018) 10–19.
- [5] X. Li, C.-I. Vasile, C. Belta, Reinforcement learning with temporal logic rewards, *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2017) 3834–3839.
- [6] D. Aksaray, A. Jones, Z. Kong, M. Schwager, C. Belta, Q-learning for robust satisfaction of signal temporal logic specifications, *IEEE 55th Conference on Decision and Control (CDC)* (2016) 6565–6570.
- [7] R. T. Icarte, T. Klassen, R. Valenzano, S. McIlraith, Using reward machines for high-level task specification and decomposition in reinforcement learning, *International Conference on Machine Learning* (2018) 2107–2116.
- [8] Z. Xu, I. Gavran, Y. Ahmad, R. Majumdar, D. Neider, U. Topcu, B. Wu, Joint inference of reward machines and policies for reinforcement learning, *Proceedings of the International Conference on Automated Planning and Scheduling* 30 (2020) 590–598.
- [9] G. Konidaris, On the necessity of abstraction, *Current opinion in behavioral sciences* 29 (2019) 1–7.
- [10] J. Andreas, D. Klein, S. Levine, Modular multitask reinforcement learning with policy sketches, *International Conference on Machine Learning* (2017) 166–175.

- [11] M. Wen, R. Ehlers, U. Topcu, Correct-by-synthesis reinforcement learning
845 with temporal logic constraints, IEEE/RSJ International Conference on
Intelligent Robots and Systems (IROS) (2015) 4983–4990.
- [12] D. Sadigh, E. S. Kim, S. Coogan, S. S. Sastry, S. A. Seshia, A learning
based approach to control synthesis of markov decision processes for linear
temporal logic specifications, IEEE Conference on Decision and Control
850 (2014) 1091–1096.
- [13] M. Hasanbeig, A. Abate, D. Kroening, Logically-constrained reinforcement
learning, arXiv preprint arXiv:1801.08099.
- [14] R. Toro Icarte, T. Q. Klassen, R. Valenzano, S. A. McIlraith, Teaching
multiple tasks to an rl agent using ltl, Proceedings of the 17th International
855 Conference on Autonomous Agents and MultiAgent Systems (2018) 452–
461.
- [15] Y. Wang, A. K. Bozkurt, M. Pajic, Reinforcement learning with temporal
logic constraints for partially-observable markov decision processes, arXiv
preprint arXiv:2104.01612.
- 860 [16] A. K. Bozkurt, Y. Wang, M. M. Zavlanos, M. Pajic, Control synthesis from
linear temporal logic specifications using model-free reinforcement learn-
ing, IEEE International Conference on Robotics and Automation (ICRA)
(2020) 10349–10355.
- [17] E. M. Hahn, M. Perez, S. Schewe, F. Somenzi, A. Trivedi, D. Wojtczak,
865 Omega-regular objectives in model-free reinforcement learning, Interna-
tional Conference on Tools and Algorithms for the Construction and Anal-
ysis of Systems (2019) 395–412.
- [18] S. Carr, N. Jansen, U. Topcu, Verifiable rnn-based policies for pomdps
under temporal logic constraints, arXiv preprint arXiv:2002.05615.

- 870 [19] M. Ghasemi, E. Bulgur, U. Topcu, Task-oriented active perception and
planning in environments with partially known semantics, *International
Conference on Machine Learning* (2020) 3484–3493.
- [20] R. R. da Silva, V. Kurtz, H. Lin, Active perception and control from tem-
poral logic specifications, *IEEE Control Systems Letters* 3 (4) (2019) 1068–
875 1073.
- [21] A.-A. Agha-Mohammadi, S. Chakravorty, N. M. Amato, Firm: Sampling-
based feedback motion-planning under motion uncertainty and imperfect
measurements, *The International Journal of Robotics Research* 33 (2)
(2014) 268–304.
- 880 [22] M. Guo, K. H. Johansson, D. V. Dimarogonas, Revising motion planning
under linear temporal logic specifications in partially known workspaces,
2013 *IEEE International Conference on Robotics and Automation* (2013)
5025–5032.
- [23] M. Lahijanian, M. R. Maly, D. Fried, L. E. Kavraki, H. Kress-Gazit, M. Y.
885 Vardi, Iterative temporal planning in uncertain environments with partial
satisfaction guarantees, *IEEE Transactions on Robotics* 32 (3) (2016) 583–
599.
- [24] S. C. Livingston, R. M. Murray, J. W. Burdick, Backtracking temporal
890 logic synthesis for uncertain environments, *IEEE International Conference
on Robotics and Automation* (2012) 5163–5170.
- [25] F. J. Montana, J. Liu, T. J. Dodd, Sampling-based reactive motion plan-
ning with temporal logic constraints and imperfect state information, *Critical
Systems: Formal Methods and Automated Verification* (2017) 134–149.
- [26] A. M. Ayala, S. B. Andersson, C. Belta, Temporal logic motion planning in
895 unknown environments, *IEEE/RSJ International Conference on Intelligent
Robots and Systems* (2013) 5279–5284.

- [27] A. Camacho, O. Chen, S. Sanner, S. A. McIlraith, Non-markovian rewards expressed in ltl: guiding search via reward shaping, Tenth Annual Symposium on Combinatorial Search.
- 900 [28] B. Ramasubramanian, A. Clark, L. Bushnell, R. Poovendran, Secure control under partial observability with temporal logic constraints, American Control Conference (ACC) (2019) 1181–1188.
- [29] X. Zhang, B. Wu, H. Lin, Learning based supervisor synthesis of pomdp for pctl specifications, 54th IEEE Conference on Decision and Control (CDC) (2015) 7470–7475.
- 905 [30] D. Furelos-Blanco, M. Law, A. Russo, K. Broda, A. Jonsson, Induction of subgoal automata for reinforcement learning, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 34, 2020, pp. 3890–3897.
- [31] M. Hasanbeig, N. Y. Jeppu, A. Abate, T. Melham, D. Kroening, Deep-synth: Automata synthesis for automatic task segmentation in deep reinforcement learning, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 35, 2021, pp. 7647–7656.
- 910 [32] A. Abate, Y. Almula, J. Fox, D. Hyland, M. Wooldridge, Learning task automata for reinforcement learning using hidden markov models, arXiv preprint arXiv:2208.11838.
- 915 [33] R. Toro Icarte, E. Waldie, T. Klassen, R. Valenzano, M. Castro, S. McIlraith, Learning reward machines for partially observable reinforcement learning, Advances in Neural Information Processing Systems 32 (2019) 15523–15534.
- 920 [34] G. Rens, J.-F. Raskin, R. Reynouad, G. Marra, Online learning of non-markovian reward models, arXiv preprint arXiv:2009.12600.
- [35] T. Dohmen, N. Topper, G. Atia, A. Beckus, A. Trivedi, A. Velasquez, Inferring probabilistic reward machines from non-markovian reward signals

- for reinforcement learning, in: Proceedings of the International Conference
925 on Automated Planning and Scheduling, Vol. 32, 2022, pp. 574–582.
- [36] R. Bellman, A markovian decision process, *Journal of mathematics and mechanics* (1957) 679–684.
- [37] M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*, John Wiley & Sons, 2014.
- 930 [38] E. M. Gold, Complexity of automaton identification from given data, *Information and control* 37 (3) (1978) 302–320.
- [39] D. Neider, N. Jansen, Regular model checking using solver technologies and automata learning, *NASA Formal Methods Symposium* (2013) 16–31.
- [40] M. J. Heule, S. Verwer, Exact dfa identification using sat solvers, *International Colloquium on Grammatical Inference* (2010) 66–79.
935
- [41] D. Neider, *Applications of automata learning in verification and synthesis*, Ph.D. thesis, Hochschulbibliothek der Rheinisch-Westfälischen Technischen Hochschule Aachen (2014).
- [42] M. Jin, Z. Ma, K. Jin, H. H. Zhuo, C. Chen, C. Yu, Creativity of ai:
940 Automatic symbolic option discovery for facilitating deep reinforcement learning, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 36, 2022, pp. 7042–7050.
- [43] L. Tingguang, H. Danny, L. Chenming, Z. Delong, W. Chaoqun, M. Q.-H. Meng, Houseexpo: A large-scale 2d indoor layout dataset for learning-based
945 algorithms on mobile robots, arXiv preprint arXiv:1903.09845.
- [44] A. Morgado, C. Dodaro, J. Marques-Silva, Core-guided maxsat with soft cardinality constraints (2014) 564–573.
- [45] A. Ignatiev, A. Morgado, J. Marques-Silva, PySAT: A Python toolkit for prototyping with SAT oracles (2018) 428–437doi:10.1007/

950

978-3-319-94144-8_26.

URL https://doi.org/10.1007/978-3-319-94144-8_26

- [46] Y. Wen, P. Vicol, J. Ba, D. Tran, R. Grosse, Flipout: Efficient pseudo-independent weight perturbations on mini-batches.

URL <https://openreview.net/forum?id=rJNpifWAb>

955

- [47] C. J. Watkins, P. Dayan, Q-learning, Machine learning 8 (3-4) (1992) 279–292.

Appendix A.

We provide here the proof of Theorem 1. We first need some necessary
960 concepts, starting with the *attainable trajectory*:

Definition 4. Let $\mathcal{M} = (S, s_I, A, \mathcal{T}, R, \mathcal{AP}, \hat{L}, \gamma)$ be an *l-MDP* and $m \in \mathbb{N}$ a natural number. We call a trajectory $\zeta = s_0 a_0 s_1 \dots s_k a_k s_{k+1} \in (S \times A)^* \times S$ *m-attainable* if (i) $k \leq m$ and (ii) $\mathcal{T}(s_i, a_i, s_{i+1}) > 0$ for each $i \in \{0, \dots, k\}$. Moreover, we say that a trajectory ζ is *attainable* if there exists an $m \in \mathbb{N}$ such
965 that ζ is *m-attainable*.

Since the function `GetEpsilonGreedyAction()` (line 3 of Algorithm 4) follows an ε_a -greedy policy, we can show that the proposed algorithm almost surely explores every attainable trajectory in the limit, i.e., with probability 1 as the number of episodes goes to infinity.

970

Lemma 1. [8, Lemma 1] Let $m \in \mathbb{N}$ be a natural number. Then, Algorithm 3, with *eplength* $\geq m$, almost surely explores every *m-attainable trajectory* at least once in the limit.

Similar to Def. 4, we call an observed label sequence $\hat{\lambda} = \hat{\ell}_0, \dots, \hat{\ell}_k$ (*m-attainable* on $\hat{\mathcal{L}}_h$ if there exists an (*m-attainable trajectory* $s_0 a_0 s_1 \dots s_k a_k s_{k+1}$ such that $\hat{\ell}_j = \hat{L}(\hat{\mathcal{L}}_h, s)$ for each $j \in \{0, \dots, k\}$.)
975

Consider now Algorithm 3 and assume that the condition of `SignifChange`($\hat{\mathcal{L}}_h, \hat{\mathcal{L}}_j$) (line 13) is not satisfied after a certain number of episodes, i.e., $\hat{\mathcal{L}}_h$ remains

fixed. Then, Lemma 1 implies that Algorithm 3 almost surely explores every m -attainable label sequences on $\hat{\mathcal{L}}_h$ in the limit, formalized in the following
 980 corollary.

Corollary 1. *Assume that there exists an $n_r > 0$ such that $\text{SignifChange}(\hat{\mathcal{L}}_h, \hat{\mathcal{L}}_j) = \text{False}$ (line 13 of Algorithm 3), for all episodes $n > n_r$. Then Algorithm 3, with $\text{eplength} \geq m$, explores almost surely every m -attainable label sequence on $\hat{\mathcal{L}}_h$ in the limit.*

985 Therefore, if Algorithm 3 explores sufficiently many m -attainable label sequences on some distribution $\hat{\mathcal{L}}_h$ and for a large enough value of m , it is guaranteed to infer a reward machine that is “good enough” in the sense that it is equivalent to the reward machine encoding the reward function R on $\hat{\mathcal{L}}_h$ and on all attainable label sequences on $\hat{\mathcal{L}}_h$, assuming that such a reward machine
 990 exists (see Example 3.1). This is formalized in the next lemma.

Lemma 2. *[8, Lemma 2] Let $\mathcal{M} = (S, s_I, A, \mathcal{T}, R, \mathcal{AP}, \hat{L}, \gamma)$ be an l -MDP and assume that there exists a reward machine \mathcal{A}_h that encodes the reward function R on $\hat{L}(\hat{\mathcal{L}}_h, \cdot)$ for some belief $\hat{\mathcal{L}}_h$. Assume that there exists an $n_r > 0$ such that $\text{SignifChange}(\hat{\mathcal{L}}_h, \hat{\mathcal{L}}_j) = \text{False}$ (line 13 of Alg. 3), for all episodes $n > n_r$. Then,
 995 Algorithm 3, with $\text{eplength} \geq 2^{|\mathcal{M}|+1}(|\mathcal{A}_h| + 1) - 1$, almost surely learns a reward machine in the limit that is equivalent to \mathcal{A}_h on all attainable label sequences on $\hat{\mathcal{L}}_h$.*

Following Lemma 2, Algorithm 3 will eventually learn the reward machine encoding the reward function on $\hat{\mathcal{L}}_h$, when $\hat{\mathcal{L}}_h$ is fixed. Intuitively, Lemma
 1000 2 suggests that, eventually, the algorithm is able to learn a reward machine that encodes the reward function on a *fixed* estimate $\hat{L}(\hat{\mathcal{L}}_h, \cdot)$, i.e., without the updates from the new observations. Nevertheless, the potential perception updates (line 14 of Algorithm 4) might prevent the algorithm **from learning a reward machine**, since $\hat{\mathcal{L}}_h$ might be changing after a finite number of episodes.
 1005 However, if the observation model is accurate (or inaccurate) enough and the agent explores sufficiently many label sequences, intuition suggests that $\hat{\mathcal{L}}_j$ will

be constantly improving, leading to less frequent updates with respect to the divergence test (line 13 of Algorithm 3). Therefore, there exists a finite number of episodes, after which $\hat{\mathcal{L}}_h$ will be fixed to some belief $\hat{\mathcal{L}}_f$ and $\hat{L}(\hat{\mathcal{L}}_f, \cdot)$ will be “close” to the ground-truth function L_G . By applying then Lemma 2, we conclude that Algorithm 3 will learn the reward machine that encodes the reward function on $\hat{\mathcal{L}}_f$, and consequently converge to the q -function that defines an optimal policy. This is used in the proof of Theorem 1 that follows.

Proof of Theorem 1. Note first that the ε_a -greedy action policy, imposed by the function `GetEpsilonGreedyAction()`, and $eplength \geq |\mathcal{M}|$ imply that every action-pair of \mathcal{M} will be visited infinitely often. Moreover, `SignifChange`($\hat{\mathcal{L}}_h, \hat{\mathcal{L}}_j$) will be always false for all episodes $n > n_f$, since $\hat{\mathcal{L}}_h$ will be an accurate enough estimate. By definition, the reward machine \mathcal{A} that encodes the reward function on the ground truth encodes the reward function on $\hat{\mathcal{L}}_h$ as well. Therefore, Lemma 2 guarantees that Algorithm 3 eventually learns a reward machine \mathcal{H} that is equivalent to \mathcal{A} on all attainable label sequences.

According to Observation 1 of [7], given the **l-MDP** $\mathcal{M} = (S, s_I, A, \mathcal{T}, R, \mathcal{AP}, \hat{L}, \gamma)$ and the reward machine $\mathcal{A} = (V, v_I, 2^{\mathcal{AP}}, \mathcal{R}, \delta, \sigma)$, one can construct an **l-MDP** $\mathcal{M}_{\mathcal{H}} = (S \times V, (s_I, v_I), A, \mathcal{T}', R', \mathcal{AP}, \hat{L}, \gamma)$ with

$$\mathcal{T}'((s, v), a, (s', v')) = \begin{cases} \mathcal{T}(s, a, s'), & \text{if } v' = \delta(v, L(s')) \\ 0 & \text{otherwise} \end{cases}$$

$$R'((s, v), a, (s', v')) = \sigma(v, L(s'))$$

and a Markovian reward function such that every attainable label sequence of $\mathcal{M}_{\mathcal{H}}$ receives the same reward as in \mathcal{M} . Thus, an optimal policy for $\mathcal{M}_{\mathcal{H}}$ will be also optimal for \mathcal{M} . Due to the ε_a -greedy action policy, the episode length being $eplength \geq |\mathcal{M}|$, and the fact that the updates are done in parallel for all states of the reward machine \mathcal{H} , every state-action pair of the **l-MDP** $\mathcal{M}_{\mathcal{H}}$ will be visited infinitely often. Therefore, convergence of q -learning for $\mathcal{M}_{\mathcal{H}}$ is guaranteed [47]. Finally, since an optimal policy for $\mathcal{M}_{\mathcal{H}}$ is optimal for \mathcal{M} , Algorithm 3 converges to an optimal policy too.

1030

