

Learning Task Automata for Reinforcement Learning Using Hidden Markov Models

Alessandro Abate^a, Yousif Almulla^{a,b}, James Fox^a, David Hyland^a and Michael Wooldridge^a

^aUniversity of Oxford

^bMicrosoft Azure Quantum

{aabate, james.fox, david.hyland, mjlw}@cs.ox.ac.uk, yalmulla@microsoft.com

Abstract. Training reinforcement learning (RL) agents using scalar reward signals is often infeasible when an environment has sparse and non-Markovian rewards. Moreover, handcrafting these reward functions before training is prone to misspecification. We learn non-Markovian finite task specifications as finite-state ‘task automata’ from episodes of agent experience within environments with unknown dynamics. First, we learn a product MDP, a model composed of the specification’s automaton and the environment’s MDP (both initially unknown), by treating it as a partially observable MDP and employing a hidden Markov model learning algorithm. Second, we efficiently distil the task automaton (assumed to be a deterministic finite automaton) from the learnt product MDP. Our automaton enables a task to be decomposed into sub-tasks, so an RL agent can later synthesise an optimal policy more efficiently. It is also an interpretable encoding of high-level task features, so a human can verify that the agent’s learnt tasks have no misspecifications. Finally, we also take steps towards ensuring that the automaton is environment-agnostic, making it well-suited for use in transfer learning.

1 Introduction

Reinforcement Learning (RL) can be prohibitively slow (sample inefficient) at learning an optimal policy when the reward signal is sparse and non-Markovian because of the credit assignment problem. However, this setting is common: any task where no reward is given until a sequence of sequential sub-tasks is completed. Consider a medical procedure where the ordering of sub-task completion matters or where one can only enter a room if one already has the key.

Three existing approaches for improving learning in this setting are hierarchical RL [23], which allows agents to plan at various levels of abstraction; transfer learning, which utilises knowledge learnt from similar tasks [29]; and temporal logic planning approaches, which guide the agent’s exploration by focusing it on the MDP fragment that satisfies a linear temporal logic (LTL) property [15, 30]. The latter is similar to Icarte et al [19]’s reward machines (finite-state machines that represent non-Markovian reward functions) because an LTL property is often represented as an automaton.

Task specifications are often separable, so an automaton (or reward machine) allows an optimal policy to be found more efficiently by breaking the task into Markovian sub-tasks. Recent work (see Section 5) recognises that this automaton is often a priori unknown, so they learn it concurrently with finding an optimal policy. We focus on *how to best learn the task automaton (TA) representing a task*

specification in sparse, non-Markovian reward environments. To our knowledge, ours is the first model-based approach; we learn a model of the underlying MDP along with the TA. The aforementioned work on temporal logic planning, sample-efficient model-based RL [35], or other methods [6] can then synthesise an optimal policy.

Contributions: Our algorithmic pipeline learns both a task specification as a TA (encoded as a deterministic finite automaton (DFA)) and a model of the MDP, from episodes of agent experience within an environment with unknown dynamics and sparse, non-Markovian reward. This improves sample efficiency in three ways: a TA exposes a task specification’s separability, so it allows for solving sub-tasks independently; our model-based approach helps the agent learn an optimal policy more efficiently; and we remove environmental bias from the learned TA for better transfer learning [29].

We learn the TA via an intermediate product MDP structure, composed of the environment’s ‘*spatial*’ MDP and the task specification’s TA. The product MDP is partially observable – the agent only observes its state in the spatial MDP and its reward, not the TA-state (its progress through the unknown task). The product MDP is learnt using the Baum-Welch algorithm [7] by first learning an estimate of the spatial MDP using a uniform prior. Then, this learnt spatial MDP is used as an inductive bias for learning the full product MDP. In Section 4, we show that our approach is more efficient than existing SAT-based approaches to learning TAs [1, 9, 16, 34, 36] by comparing with Biermann and Feldman’s [8] SAT-based algorithm.

Once the product MDP is learnt, distilling the TA via our ‘Cone Lumping’ method is computationally cheap. This makes the product MDP useful for transfer learning. If the environment changes, the agent only needs to update the affected part of the product MDP before efficiently re-distilling the TA.

This paper also addresses reward misspecification [3] by using a task specification not a scalar reward function to encode a complex, non-Markovian task. Scalar reward functions are brittle with respect to small environmental changes [33], do not signal a task’s separability, and can be difficult for a human to interpret. In contrast, our learnt TA reveals the full task specification. Since our pipeline works in unknown environments where many high-level tasks may exist, agents using our pipeline can learn multiple tasks in one TA. Although we focus on deterministic TAs in this paper (i.e., DFAs), our approach can be straightforwardly extended to learning probabilistic automata [11, 24]. This paper is self-contained, but see [2] for proposition proofs and experimental details.

2 Setup

An agent’s interaction with its environment can be viewed at many levels of abstraction, amongst which a low and fine-grained description is often modelled by a *Markov Decision Process (MDP)*. However, agents can also recognise higher-level features of the environment (e.g., chairs or tables). These are included in a *labelled MDP* (in Section 2.1) using a set of (atomic) propositional variables from an alphabet \mathcal{AP} , which are assigned truth values at every MDP state via a labelling function L . Task specifications refer to sequential compositions of these high-level features and can be encoded as TAs (in Section 2.2), which graphically represent the structure of tasks that an agent can get reward from completing. A TA and the MDP can be combined into a single structure, a product MDP (in Section 2.3) to provide a holistic representation of the agent’s exploration.

2.1 MDP Environment

Definition 1. A *Markov Decision Process (MDP)* is a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, s_0, P, R, \gamma)$, where \mathcal{S} and \mathcal{A} are finite sets of states and actions, s_0 is the initial state (or distribution over states), $P: \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the transition probability distribution, $R: (\mathcal{S} \times \mathcal{A})^+ \times \mathcal{S} \rightarrow \mathbb{R}$ is a (non-Markovian) reward function¹, and $\gamma \in [0, 1)$ is the discount factor. A *labelled MDP* is a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, s_0, P, R, \gamma, \mathcal{AP}, L)$; an MDP enriched with \mathcal{AP} , a finite alphabet of atomic propositions (labels), and $L: \mathcal{S} \rightarrow 2^{\mathcal{AP}}$, a labelling function.

A sequence of states and actions $s_0, a_0, s_1, \dots, a_{n-1}, s_n$ is called a *trajectory*. A *state or action trajectory* is the restriction to just the sequence of states $(s_t)_{t=0}^n \in \mathcal{S}^+$ or actions $(a_t)_{t=0}^n \in \mathcal{A}^+$. A state trajectory is *attainable* if $\prod_{t=0}^{n-1} P(s_{t+1} | s_t, a_t) > 0$ for some action trajectory $(a_t)_{t=0}^n$, and *impossible* otherwise. A trajectory is accompanied by a *reward sequence* $(r_t)_{t=0}^n$ and a sequence of labels, called a *trace*, $(\ell_t)_{t=0}^n \in (2^{\mathcal{AP}})^*$, where $L(s_t) = \ell_t$ for all $t \leq n$. A (*history-based*) *policy* $\pi: (\mathcal{S} \times \mathcal{A})^+ \times \mathcal{A} \rightarrow [0, 1]$ is a function mapping trajectories to probability distributions over actions, where $\pi(a_t | s_0, a_0, \dots, s_t)$ gives the probability of the agent choosing action a_t at state s_t given the (historical) trajectory s_0, a_0, \dots, s_t . *Memoryless policies* are a special case where the distribution over actions only depends on the last state (standard in RL because storing entire histories is intractable). The agent starts at s_0 in each exploration episode and, at each time step t , in state s_t , it selects action a_t with probability $\pi(a_t | s_t)$, before transitioning to a new state s_{t+1} with probability $P(s_{t+1} | s_t, a_t)$ whilst receiving a reward $r_t = R(s_0, a_0, \dots, s_t, a_t, s_{t+1})$. The agent’s task is to find a policy π that maximises the expected sum of discounted reward $\mathbb{E}_\pi [\sum_{t=0}^n \gamma^t r_t]$. A policy π is *fully mixed* if every action at every state is chosen with non-zero probability (i.e., $\pi(a | s) > 0$ for all $s \in \mathcal{S}$ and $a \in \mathcal{A}$). We assume the agent observes state trajectories, traces, and reward sequences, as it explores the labelled MDP.

Given the initial state s_0 , fixing the agent’s policy π in an MDP \mathcal{M} induces a Markov chain (MC) \mathcal{M}_π if for all state transitions, the reward function’s output is independent of the last action (i.e., $R(\zeta, a_i, s_{t+1}) = R(\zeta, a_j, s_{t+1})$ for all $a_i, a_j \in \mathcal{A}$, where $\zeta \in (\mathcal{S} \times \mathcal{A})^+$). Definition 1’s reward function is *non-Markovian* because it depends upon the full trajectory up until that point. A *Markovian reward function* $R: \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ only depends on the last transition. State-based rewards could also be handled straightforwardly.

¹ This must encompass the initial transition, as evidenced by the use of the non-empty repetition operator $+$.

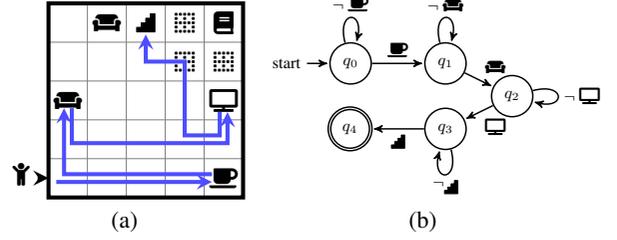


Figure 1: (a) The labelled MDP environment and (b) 5 state TA representing a task specification for Example 1.

Example 1. An RL agent is in Figure 1a’s environment. The task is to collect coffee ☕ for the guest on a couch 🛋️ before turning on the TV 📺 and then ascending the stairs 🧶 (Figure 1b). The carpet 🧶 and book 📖 do not affect whether the agent can achieve its goal. The arrow next to the silhouette indicates the initial state.

Example 1’s MDP environment is shown in Figure 1a. The agent’s possible actions \mathcal{A} are to move in the adjacent up, down, left, and right directions; the MDP’s state space cardinality $|\mathcal{S}|$ is 25; and the transition dynamics are initially unknown (as discussed later, stochasticity in the dynamics derives from randomised action execution). The set of (atomic) propositional variables $\mathcal{AP} \equiv \{\text{☕}, \text{🛋️}, \text{📺}, \text{🧶}, \text{📖}\}$ represent the high-level features of the environment. Empty cells are implicitly labelled $L(s) = \emptyset$ (i.e., the MDP state $s \in \mathcal{S}$ has no high-level environmental feature of interest).

If the agent only gets a positive reward upon completing the full task specification, this cannot be expressed with a Markovian reward function with respect to the MDP’s state space \mathcal{S} because the agent cannot go between states labelled ☕ and 🛋️ in one timestep, let alone complete all four subtasks. So, the agent must learn an optimal policy in an environment with a non-Markovian reward function that provides sparse feedback. This challenge of planning using full *trajectories* can be surmounted by recognising that the agent also observes *traces* of high-level features in \mathcal{AP} . By maintaining a memory of complete subtasks (e.g., coffee collected) at this higher-level of abstraction (with smaller dimension), the agent can overcome the ‘curse of dimensionality’ and synthesise an optimal policy more efficiently by exploiting the compositionality of many tasks.

2.2 Task Specification

Linear temporal logic (LTL) or *Linear Dynamic Logic (LDL)*, evaluated with respect to finite trajectories (LTL_f/ LDL_f), is often used to express a task specification as temporally extended goals, or constraints on plans because this mitigates the sample inefficiency of RL in non-Markovian environments [15]. In fact, Littman et al [22] argue that task specifications are superior to reward functions for specifying behaviour, as they are environment independent. We exploit the fact that the expressive power of both LDL_f and LTL_f is subsumed by the class of regular expressions, so they can be represented by deterministic or nondeterministic finite automata (DFAs/NFAs) [6].

Definition 2. A *Non-Deterministic Finite Automaton (NFA)* is a tuple $\mathcal{N} = (\mathcal{Q}, q_0, \Sigma, \delta, \mathcal{F})$, where \mathcal{Q} is a finite set of states, $q_0 \in \mathcal{Q}$ is the initial state, Σ is a finite alphabet, $\delta: \mathcal{Q} \times \Sigma \rightarrow 2^{\mathcal{Q}}$ is a non-deterministic transition function, and $\mathcal{F} \subseteq \mathcal{Q}$ is the set of accepting states. If $|\delta(q, \alpha)| = 1$ for all $q \in \mathcal{Q}$ and $\alpha \in \Sigma$, then \mathcal{N} is deterministic and is called a *Deterministic Finite Automaton (DFA)*.

Let Σ^* be the set of all finite words over Σ . A finite word $w = \alpha_1 \alpha_2 \dots \alpha_n \in \Sigma^*$ is accepted by an NFA \mathcal{N} if there exists a finite run $u_0, \dots, u_n \in \mathcal{Q}^*$ where $u_0 = q_0, u_{i+1} \in \delta(u_i, \alpha_{i+1})$ for

$0 \leq i < n$, and $u_n \in \mathcal{F}$. In a DFA, only one path follows w , as each transition is deterministic, so the final state must be *accepting*, otherwise the word is *rejected*. The set of all accepted words is the language recognised by \mathcal{N} , denoted by $\mathcal{L}(\mathcal{N})$. A DFA is trivially an NFA and any NFA can be converted into an equivalent DFA using Rabin and Scott [25]’s subset construction method, so NFAs and DFAs both recognise the class of regular languages.

Definition 3. A *Task Automaton (TA)* \mathcal{A} is a DFA $(\mathcal{Q}, q_0, \Sigma, \delta, \mathcal{F})$ with $\Sigma = 2^{\mathcal{AP}}$, where \mathcal{AP} is a set of propositional variables inherited from the labelled MDP \mathcal{M} .

We define a TA as a DFA to make our presentation as clear as possible. However, other types of automaton may also be used, such as Icarte et al [19]’s reward machines which, rather than having a set of accepting states (as in DFAs), outputs a reward function at each state $\delta_r : \mathcal{Q} \rightarrow [S \times \mathcal{A} \times S \rightarrow \mathbb{R}]$. Our statistical approach to learning TAs also allows for probabilistic TAs [24], where rewards exhibit both non-Markovian and stochastic dynamics.

The TA \mathcal{A} keeps a ‘memory’ of traces $(\ell_t)_{t=0}^n \in (2^{\mathcal{AP}})^+$ that the exploring agent observes by transitioning according to the current truth assignment of the propositional variables. A trace is *MDP-attainable* if there exists a state trajectory $(s_t)_{t=0}^n$ such that $L(s_t) = \ell_t$ for all t . The TA’s accepting states \mathcal{F} mark the satisfaction of some task specification. In other words, the agent completes a task via trace $(\ell_t)_{t=0}^n$ if, correspondingly, the TA starts in state $u_0 = q_0$, transitions according to $u_{t+1} = \delta(u_t, L(s_{t+1}))$ for $t \geq 0$, and $u_n \in \mathcal{F}$.

The TA in Figure 1b represents the task specification for Example 1 in terms of the high-level environmental features. Each move in the MDP is matched by a transition in the TA (recall that empty squares in Figure 1a are labelled by \emptyset). For example, the (blue) marked trajectory in Figure 1a represents an MDP-attainable accepting path in the TA and each ‘sub-path’ denotes a TA ‘sub-task’. The sub-path $(\emptyset, \emptyset, \emptyset, \emptyset, \blacksquare)$ denotes the path in the MDP that corresponds to traversing the $q_0 \rightarrow q_1$ edge in the TA. Note that the labels on self-loops in Figure 1b (e.g., $\neg \blacksquare$) are syntactic sugar to denote all labels not associated with outgoing edges from that TA state. The TA is environment-agnostic because it applies to any labelled MDP that includes the set of all high-level features $\{\blacksquare, \blacksquare, \blacksquare, \blacksquare\}$ necessary for completing the corresponding task specification. Therefore, learning the TA can not only expedite policy synthesis within this environment, but it also helps transfer learning, where an agent wants to use what it knows about attaining reward in one environment to obtain reward in different environments – e.g., if an agent learns it is necessary to pick up a key before opening a door in one house, then this is likely applicable to other houses.

2.3 Product MDP

Definition 4. Given an MDP without reward $\mathcal{M} \setminus R = (\mathcal{S}, \mathcal{A}, s_0, P)$ and a TA $\mathcal{A} = (\mathcal{Q}, q_0, \Sigma, \delta, \mathcal{F})$, the *product MDP* is the structure $\mathcal{M} \otimes \mathcal{A} = (\mathcal{S}^\otimes, \mathcal{A}, s_0^\otimes, P^\otimes, R^\otimes, \gamma)$, where $\mathcal{S}^\otimes = \mathcal{S} \times \mathcal{Q}$, $s_0^\otimes = \langle s_0, q_0 \rangle$ is the initial state pair; $P^\otimes : \mathcal{S}^\otimes \times \mathcal{A} \times \mathcal{S}^\otimes \rightarrow [0, 1]$ is the transition probability function, where $P^\otimes(\langle s_j, q_j \rangle | \langle s_i, q_i \rangle, a_i) = P(s_j | s_i, a_i)$ if $q_j = \delta(q_i, L(s_j))$ and 0 otherwise; and $R^\otimes : \mathcal{S}^\otimes \times \mathcal{A} \times \mathcal{S}^\otimes \rightarrow \{0, 1\}$ is a reward function such that $R^\otimes(\langle s_t, q_t \rangle, a_t, \langle s_{t+1}, q_{t+1} \rangle) = \chi_{\mathcal{F}}(q_{t+1})$, where $\chi_{\mathcal{F}}$ is the indicator function for the TA’s set of accepting states. Non-zero reward is obtained iff the agent transitions to an *accepting state* of the product MDP (those related to an accepting TA state). Fixing any policy π in the product MDP $\mathcal{M} \otimes \mathcal{A}$, induces a *product MC* $\mathcal{M}_\pi \otimes \mathcal{A}$.

A product MDP $\mathcal{M} \otimes \mathcal{A}$ more holistically represents the agent’s exploration progress towards achieving a task because both the MDP \mathcal{M} and the TA \mathcal{A} transition: in \mathcal{M} , these are the low-level dynamics over states in \mathcal{S} ; in \mathcal{A} , these represent the agent’s progress towards completing a task specification. $\mathcal{M} \otimes \mathcal{A}$ ’s reward function R^\otimes is always Markovian because although a task may be non-Markovian with respect to \mathcal{M} , $\mathcal{M} \otimes \mathcal{A}$ keeps track of what is happening at both \mathcal{M} ’s and \mathcal{A} ’s level of abstraction. In other words, if \mathcal{M} has a non-Markovian reward function R representing a task encodable by \mathcal{A} , then the non-Markovianity in R is resolved by synchronising $\mathcal{M} \otimes \mathcal{A}$. The transitions in $\mathcal{M} \otimes \mathcal{A}$ are executed ‘under the hood’ as the agent explores (i.e., no a priori knowledge about $\mathcal{M} \otimes \mathcal{A}$ is needed). Moreover, $\mathcal{M} \otimes \mathcal{A}$ does not increase the agent’s possible policy space because \mathcal{A} has deterministic transitions, so the trajectories (and so also policies) in \mathcal{M} and $\mathcal{M} \otimes \mathcal{A}$ are in a bijection.

3 Task Automaton Learning

We now give our three-step algorithmic pipeline for learning an MDP and a TA (i.e., the task specification) in unknown environments.

Algorithm 1 Learning a TA in an unknown labelled MDP

Input: put agent into an (unknown) labelled MDP \mathcal{M}

Output: TA \mathcal{A}^* that represents a task specification

- 1: OBSSEQ \leftarrow collect episodes of corresponding trajectories, traces, and reward sequences.
 - 2: **function** LEARNPRODUCTMDP(OBSSEQ) \triangleright **Step 1**
 - 3: Use an HMM/POMDP learning algorithm to learn estimates of
 - (i) the spatial MDP $\hat{\mathcal{M}}$ ’s transition probability distribution \hat{P}
 - (ii) the product MDP $\hat{\mathcal{M}} \otimes \hat{\mathcal{A}}$ ’s transition probability distribution \hat{P}^\otimes .
 - 4: **function** DISTILTA($\hat{\mathcal{M}} \otimes \hat{\mathcal{A}}$, OBSSEQ) \triangleright **Step 2**
 - 5: Determinise $\hat{\mathcal{M}} \otimes \hat{\mathcal{A}}$ using Cone Lumping to return the MDP-restricted TA $\hat{\mathcal{A}}_{\hat{\mathcal{M}}}$
 - 6: **function** POSTPROCESS($\hat{\mathcal{A}}_{\hat{\mathcal{M}}}$, OBSSEQ) \triangleright **Step 3**
 - 7: Remove environmental bias and minimise $\hat{\mathcal{A}}_{\hat{\mathcal{M}}}$
-

3.1 Step 1: Learn the Product MDP

Although the agent’s interactions with its environment are fully modelled by the product MDP, the product MDP is only partially observable. The agent observes MDP states as well as the associated labels and rewards but does not observe the TA states it visits. So, Step 1 involves learning the product MDP’s transition function P^\otimes by viewing the agent’s interaction with the (product MDP) environment as a partially-observable Markov Decision Process (POMDP).

Definition 5. A *Partially-Observable Markov Decision Process (POMDP)* is a tuple $\mathcal{P} = (\mathcal{S}, \mathcal{A}, s_0, P, R, \gamma, \mathcal{O}, Z)$ where $(\mathcal{S}, \mathcal{A}, s_0, P, R, \gamma)$ is an MDP, \mathcal{O} is an observation set, and $Z : \mathcal{O} \times \mathcal{S} \rightarrow [0, 1]$ is an observation probability function, where $Z(o | s)$ is the probability that observation $o \in \mathcal{O}$ is seen when the agent is at the *hidden state* $s \in \mathcal{S}$.

Remark 1. $\mathcal{P}^\otimes := (\mathcal{S}^\otimes, \mathcal{A}, s_0^\otimes, P^\otimes, R^\otimes, \gamma, \mathcal{O}, Z)$ is a POMDP that models an agent interacting with a product MDP $\mathcal{M} \otimes \mathcal{A} = (\mathcal{S}^\otimes, \mathcal{A}, s_0^\otimes, P^\otimes, R^\otimes, \gamma)$, where the agent makes observations from the set $\mathcal{O} = \mathcal{S} \times \{0, 1\}$ according to the deterministic observation function $Z : \mathcal{O} \times \mathcal{S}^\otimes \rightarrow \{0, 1\}$.

Remark 1 follows from Definitions 4 and 5. The agent’s observations in the POMDP \mathcal{P}^\otimes are from $\mathcal{S} \times \{0, 1\}$, where the latter component is a binary value corresponding to the co-domain of the reward function R^\otimes . At each time step t , the agent transitions from a hidden state $s_t^\otimes = \langle s_t, q_t \rangle$ to $s_{t+1}^\otimes = \langle s_{t+1}, q_{t+1} \rangle$ in the product MDP, but observes just $\langle s_{t+1}, \chi_{\mathcal{F}}(q_{t+1}) \rangle$ with probability 1. The product MDP state s_{t+1}^\otimes is *partially obscured* from the agent. Recall $\chi_{\mathcal{F}}$ is the indicator function on the set of accepting states \mathcal{F} in the TA \mathcal{A} . If a policy π is chosen by the agent interacting with the POMDP \mathcal{P}^\otimes , then the POMDP becomes a partially observable Markov chain \mathcal{P}_π^\otimes , known as a *Hidden Markov Model (HMM)* [26].

To learn the product MDP $\mathcal{M} \otimes \mathcal{A}$, we can employ any algorithm for learning HMMs or POMDPs and apply it to \mathcal{P}^\otimes or \mathcal{P}_π^\otimes . We use the Baum-Welch algorithm (BW) [7], due to its popularity and the ability to steer optimisation towards a desired local optimum (see Section 4 and, for further details, [2, Appendices C and D]). We learn \mathcal{P}_π^\otimes in two stages. (i) Learn an estimate \hat{P} of the true spatial MDP \mathcal{M} ’s transition probability distribution P using a fully mixed exploration policy π . This can be any fully mixed policy, but we chose a uniform random policy to assume no initial knowledge about the underlying model. Since we know π , we can always recover P from P_π . (ii) Use this estimate \hat{P} as an inductive bias for learning the transition probability distribution \hat{P}^\otimes of the full product MDP $\mathcal{M} \otimes \mathcal{A}$.

BW requires an initial guess on the number of hidden states ($|\mathcal{S}^\otimes| = |\mathcal{S}| \times |\mathcal{Q}|$). Since the agent can observe the MDP states, we only need guess the number of TA states \hat{k} . Furthermore, this guess only needs to be at least as high as the true k (i.e., $\hat{k} \geq \mathcal{Q}$). If \hat{k} is too high, it generates duplication in the learnt matrix, which is easily identified. For a sufficient upper bound for \hat{k} one can: use the label set cardinality, use a TA simplicity assumption, or identify the minimum number of distinct labels needed for task completion using successful exploration episodes. We did the latter. Further details about the convergence criterion and other hyper-parameters (such as the episode length and the number of episodes) are in Section 4.

To encode the learnt spatial MDP’s transition probability distribution \hat{P} as an inductive bias for learning \hat{P}^\otimes , we take the Kronecker product between a $\hat{k} \times \hat{k}$ identity matrix and \hat{P}_π (the matrix induced from \hat{P} and π) to form the basis of the initial estimate of the product MDP’s transition matrix. For example, with a three-state TA:

$$I_{3 \times 3} \otimes_K \hat{P}_\pi = \begin{pmatrix} \hat{P}_\pi & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \hat{P}_\pi & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \hat{P}_\pi \end{pmatrix}.$$

Then, a small positive probability $\varepsilon > 0$ is added to each zero entry in the above construction to ensure that all expressions in the HMM learning procedure are well-defined. The probability ε can take any positive value, but we chose a value inversely proportional to $|\mathcal{S}^\otimes|$ for consistency across experiments. Finally, each row is normalised so that the resulting initial estimate of the product MDP’s transition matrix \hat{P}_π^\otimes is a valid transition probability matrix.

Concentrating on learning \hat{P}^\otimes (the process is the same for learning \hat{P}), BW begins with initial estimates for the hidden transition distribution P_π^\otimes (an $|\mathcal{S}^\otimes| \times |\mathcal{S}^\otimes|$ matrix constructed from the inductive bias, as explained), the observation probability function Z (an $|\mathcal{S}^\otimes| \times |\mathcal{O}|$ matrix), and the initial state distribution (a vector $\rho \in [0, 1]^{|\mathcal{S}^\otimes|}$ defined as $\rho_j \equiv \Pr[s_0^\otimes = s_j^\otimes]$). It then finds the P_π^\otimes and Z that maximises the likelihood of obtaining the set of observation sequences $O = (o_1, \dots, o_T) \in (\mathcal{O}^+)^T$:

$$(\hat{P}_\pi^\otimes, \hat{Z}) = \underset{(P^\otimes, Z)}{\operatorname{argmax}} \Pr(O \mid \langle P_\pi^\otimes, Z \rangle).$$

Definition 6. Two product MDPs $\mathcal{M}_1 \otimes \mathcal{A}_1$ and $\mathcal{M}_2 \otimes \mathcal{A}_2$ with the same action spaces $\mathcal{A}_1 = \mathcal{A}_2$ are *observationally equivalent*, denoted $\mathcal{M}_1 \otimes \mathcal{A}_1 \cong \mathcal{M}_2 \otimes \mathcal{A}_2$, if for any observation sequence, any action trajectory, and all $\tau \in \mathbb{N}$, it holds that $\Pr^{\mathcal{P}_1^\otimes} [(o_t)_{t=0}^\tau \mid (a_t)_{t=1}^\tau] = \Pr^{\mathcal{P}_2^\otimes} [(o_t)_{t=0}^\tau \mid (a_t)_{t=1}^\tau]$, where $\mathcal{P}_i^\otimes = (\mathcal{M}_i \otimes \mathcal{A}_i, \mathcal{O}, Z)$. The observation spaces need not be identical for this to hold.

Definition 6 says that the learnt product MDP is, in general, not unique because several can have the same observational properties, i.e., no gathered experiment (episode) can distinguish between them.

Proposition 1. *The observational equivalence relation \cong is an equivalence relation. It is straightforward to show that it is reflexive, symmetric, and transitive.*

Proposition 2. *Let \mathcal{M} be a labelled MDP. For any two TAs $\mathcal{A}, \mathcal{A}'$, if $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}')$, then $\mathcal{M} \otimes \mathcal{A} \cong \mathcal{M} \otimes \mathcal{A}'$.*

Proposition 1 says we can find the quotient set of all product MDPs with the same action spaces under the equivalence relation \cong . The equivalence class for the product MDP $\mathcal{M} \otimes \mathcal{A}$ is denoted by $[\mathcal{M} \otimes \mathcal{A}]$. Proposition 2 enables us to formalise Step 1 of Algorithm 1.

Step 1: Given state trajectories and reward sequences from interacting with the initially unknown product MDP $\mathcal{M} \otimes \mathcal{A}$, learn an estimate $\hat{\mathcal{M}} \otimes \hat{\mathcal{A}} \in [\mathcal{M} \otimes \mathcal{A}]$.

Definition 7. The digraph $G = (E, V)$ underlying an MDP \mathcal{M} with state-space \mathcal{S} and transition probability function P has $V = \mathcal{S}$ and $E = \{s_i \rightarrow s_j : \exists a_k \in \mathcal{A} \text{ s.t. } P(s_j \mid s_i, a_k) > 0\}$. \mathcal{M} is *structurally correct* with respect to another MDP \mathcal{M}' if the digraph underlying \mathcal{M} is isomorphic to \mathcal{M}' .

Any algorithm which learns a maximum likelihood estimate (MLE) for the HMM parameters is asymptotically structurally correct, provided it converges to the true MLE and belongs in $[\mathcal{M} \otimes \mathcal{A}]$. So, every hidden state of the learned model must correspond to a state $s^\otimes \in \mathcal{S}^\otimes$ in the true product MDP $\mathcal{M} \otimes \mathcal{A}$ [38].

3.2 Step 2: Distil TA from Product MDP

The TA component q of each product MDP state $s^\otimes = \langle s, q \rangle$ is hidden, but even when visible, the agent can only recover the TA fragment accessible using paths in the MDP.

Definition 8. Given an MDP \mathcal{M} and a TA \mathcal{A} , the *MDP-restricted TA* $\mathcal{A}_{\mathcal{M}}$ is the sub-graph of \mathcal{A} covered by all MDP-attainable traces. We refer to paths within/outside of the MDP-restricted TA sub-graph as the *MDP-attainable/MDP-non-attainable* paths of the TA.

Proposition 3. *Suppose an agent is given a structurally correct product MDP $\hat{\mathcal{M}} \otimes \hat{\mathcal{A}}$ belonging to $[\mathcal{M} \otimes \mathcal{A}]$, where \mathcal{M} is the true MDP underlying the product MDP, and the agent can observe s and q for all $\langle s, q \rangle \in \mathcal{S}^\otimes$. Then, the agent can learn the MDP-restricted TA $\hat{\mathcal{A}}_{\hat{\mathcal{M}}}$, but may not be able to learn the full TA $\hat{\mathcal{A}}$.*

From Step 1, the agent knows the digraph of $\mathcal{M} \otimes \mathcal{A}$ and an estimate \hat{P} for the transition probability function of the true MDP \mathcal{M} . Therefore, we first compute a product MC $\hat{\mathcal{M}}_\pi \otimes \hat{\mathcal{A}}$ by assigning to the agent a fully mixed policy π . We now recognise that when the edges of the digraph underlying $\hat{\mathcal{M}}_\pi \otimes \hat{\mathcal{A}}$ are labelled according to each edge’s target state (recall that the agent observes traces), it results in an NFA whose accepting states are those in which reward 1 is obtained (see [2, Appendix B] for a worked example).

Proposition 4. Given a labelled product MC $\mathcal{M}_\pi \otimes \mathcal{A}$ with binary reward, labelling every edge $\langle s_i, q \rangle \rightarrow \langle s_j, q_j \rangle$ according to the label of its target state’s MDP component (i.e., in this case, $L(s_j) = \ell_j$) results in an NFA.

This NFA can be converted into a DFA with language equal to the MDP-restricted TA $\mathcal{A}_\mathcal{M}$ using the subset construction method [25].

Proposition 5. Let \mathcal{N} be the NFA underlying the learnt product MC $\mathcal{M}_\pi \otimes \mathcal{A}$. Then $\mathcal{L}(\mathcal{N}) = \mathcal{L}(\mathcal{A}_\mathcal{M})$, where $\mathcal{A}_\mathcal{M}$ is the MDP-restricted true TA.

However, the subset-construction algorithm has worst-case exponential time complexity in the size of the NFA (i.e., $\mathcal{O}(2^{|\mathcal{S}^\otimes|})$). We therefore introduce our more efficient *Cone Lumping* method for determining NFAs that underlie product MDPs, which explicitly leverages their product structure. ‘Lumping’ aligns with related literature (e.g., [10]). Our key insight is that out-going edges from a product state $\langle s, q \rangle \in \mathcal{S}^\otimes$ of a product MC $\mathcal{M}_\pi \otimes \mathcal{A}$ that have the same label *must* transition to the same next TA state. Every time this occurs, the product states are merged. When representing state transitions sequentially from left to right, this merging criterion looks for a *cone* structure to merge all states on the right-hand side (Figure 2).

Lemma 1 (Cone Lumping). Given a labelled product MC $\mathcal{M}_\pi \otimes \mathcal{A}$, let $E_\ell^{(s,q)} = \{\langle s, q \rangle \xrightarrow{L(s_i)} \langle s_i, q_j \rangle : L(s_i) = \ell\}$ be the set of all out-going edges from state $\langle s, q \rangle$ in the digraph underlying $\mathcal{M}_\pi \otimes \mathcal{A}$ with label ℓ . Since any TA transition $q \xrightarrow{\ell} q_j$ is deterministic, then $q_j = q_k$ for any pair $q_j, q_k \in \{q' : \langle s, q \rangle \xrightarrow{\ell} \langle s_i, q' \rangle\}$. Thus, for a product state $\langle s, q \rangle$, $|\{q' : \langle s, q \rangle \xrightarrow{L(s_i)} \langle s_i, q' \rangle \wedge L(s_i) = \ell\}| = 1$.

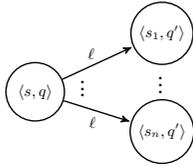


Figure 2: Schematics of ‘Cone Lumping’.

Performing Cone Lumping for all $\langle s, q \rangle \in \mathcal{S}^\otimes$ until a fixed point is reached (i.e., no more states can be merged), obtains a TA accepting (at minimum) $\mathcal{L}(\mathcal{A}_\mathcal{M})$, and at most $\mathcal{L}(\mathcal{A})$ (by construction of the product MDP). This is because all non-determinism in the NFA underlying $\mathcal{M} \otimes \mathcal{A}$ comes from the possible existence of multiple out-going edges with the same label from a product state.

Proposition 6. Cone Lumping determines any product MDP in $\mathcal{O}(|\mathcal{S}^\otimes|^3)$ time, exponentially faster than the $\mathcal{O}(2^{|\mathcal{S}^\otimes|})$ algorithm of Rabin and Scott [25]’s subset-construction.

Cone Lumping therefore efficiently converts an NFA underlying a labelled MC into a DFA (our TA in this work, Section 4) and it assists in transfer learning because the hardness is confined to learning the product MDP. If the MDP or TA changes slightly, the agent only has to relearn the affected part of the product MDP, before using Cone Lumping to efficiently distil the new TA. In practice, this uses the transition probability distribution \hat{P}_π^\otimes of the old MDP and TA’s product MDP as an initial estimate for the new product MDP. The parts of the MDP and TA that remain the same will not need updating, so Baum-Welch only needs to update the parts that have changed. This speeds up convergence significantly [38]. Finally, this process can also be useful outside of our work, e.g., for (bi)simulation or model reduction studies [20].

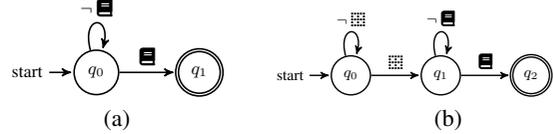


Figure 3: Two TAs for Example 1’s environment in Figure 1a to show that the learned TA can be environmentally biased.

Step 2: Applying Cone Lumping to a labelled product MC $\mathcal{M}_\pi \otimes \mathcal{A}$ for all $\langle s, q \rangle \in \mathcal{S}^\otimes$ until a fixed point is reached returns a TA $\mathcal{A}'_\mathcal{M}$, where $\mathcal{L}(\mathcal{A}'_\mathcal{M}) = \mathcal{L}(\mathcal{A}_\mathcal{M})$.

3.3 Step 3: Remove environmental bias and minimise the learnt TA

The TA distilled from the product MDP in Step 2 is in general not unique (cf. Section 3.1). If the TA is a DFA, as has been this paper’s focus, it can be minimised using, for instance, Hopcroft [17]’s algorithm in $\mathcal{O}(n \log n)$. However, the TA might also not be minimal due to environmental bias. If a certain label *must* be traversed by the agent to reach a necessary label for achieving the underlying task specification, automaton-learning algorithms cannot differentiate the importance of the two labels and will therefore assume that *both* are necessary TA transitions, even when the former is unnecessary.

Theorem 1. Given an arbitrary product MDP $\mathcal{M} \otimes \mathcal{A}$, the TA distilled from $\mathcal{M} \otimes \mathcal{A}$ may not be unique because there can exist $\mathcal{A}' \neq \mathcal{A}$, with or without $\mathcal{L}(\mathcal{A}') \neq \mathcal{L}(\mathcal{A})$, such that $\mathcal{M} \otimes \mathcal{A} \cong \mathcal{M} \otimes \mathcal{A}'$.

Proof sketch by (counter)example. Let \mathcal{M} be the labelled MDP defined by the environment in Figure 1a and \mathcal{A} and \mathcal{A}' be the TAs in Figures 3a and 3b respectively, where \mathcal{A} is the true TA. Then, \mathcal{A}' satisfies $\mathcal{M} \otimes \mathcal{A}' \cong \mathcal{M} \otimes \mathcal{A}$. In Figure 1, a \blacksquare -label must be traversed immediately prior to reaching the \blacksquare -label, no matter which MDP path to \blacksquare is taken. Because reward is only obtained upon reaching \blacksquare , there is no way to identify from the reward signal if the sequence (\blacksquare , \blacksquare) is necessary for reward, or just reaching \blacksquare suffices. The agent does not know which of the two task specifications, expressed by the TAs \mathcal{A} and \mathcal{A}' , is sufficient. \square

Corollary 1. In general, it is impossible to verify that a TA \mathcal{A}' is the TA that underpins the product MDP $\mathcal{M} \otimes \mathcal{A}$ (i.e., that $\mathcal{A}' = \mathcal{A}$), even given knowledge of \mathcal{M} and $\mathcal{M} \otimes \mathcal{A}$.

Thus, to *minimise* environmental bias, we must use counterfactual reasoning in our *post-processing* step to select among the set of TAs $\{\mathcal{A}' : \mathcal{M} \otimes \mathcal{A}' \cong \mathcal{M} \otimes \mathcal{A}\}$ the TA that requires the smallest alphabet to explain its behaviour.

Step 3: For every label in $L(\mathcal{S})$, guess that it has no effect in the supposed true TA \mathcal{A} (i.e., the agent self-loops on every TA state). To check this guess, first merge any states of the learned-automaton \mathcal{A}' that have a non-loop edge between them with this label, then check if the resulting TA outputs the correct acceptance/rejection on every word in the set of observed sequences.

This post-processing only removes labels that are never meaningful to \mathcal{A} ; it does not remove labels that *are* necessary somewhere, but may still exist ‘wrongfully’ in \mathcal{A}' due to environmental bias. Further post-processing can address this by looking for neighbourhoods of (possibly same-labelled) states that entirely surround another label, then checking, as above, whether these are necessary in \mathcal{A} .

4 Experimental Results

Although the main contributions of this paper are theoretical, our experiments (available at <https://github.com/dkhyland/TALearner>) evaluate the efficacy of Algorithm 1. To our knowledge, we are the first to focus on learning both the TA and spatial MDP starting from no a-priori knowledge. Our experiments assess our pipeline’s ability to scale up to 5x5 MDP grid-worlds and 5-state TAs and benchmark, where possible, to related work. These sizes are comparable to or beyond those handled in all related work (see Section 5).

Figure 4 illustrates the three grid worlds and the three-, four-, and five-state TAs used for our experiments. The grid worlds all share a common set of high-level environmental features $\mathcal{AP} \equiv \{\text{TV, coffee, couch, stairs}\}$, which have been distributed randomly across the grid. The agent can visit every action and state within the grid world.

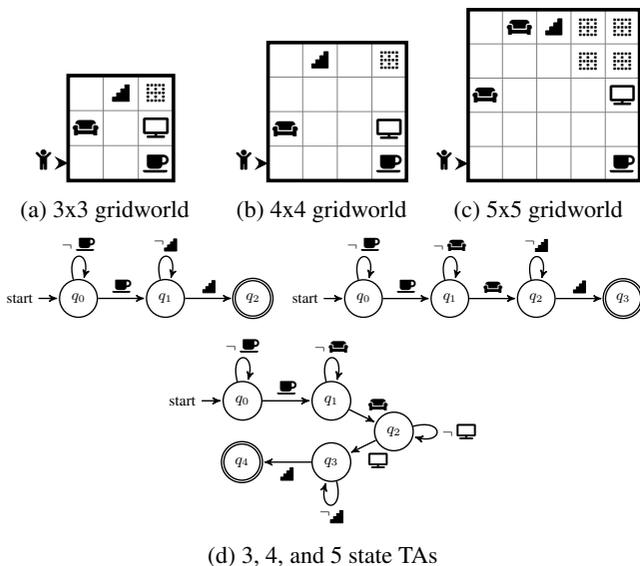


Figure 4: (a), (b), and (c) show the grid worlds used in all experiments. The agent starts in the bottom left-hand corner of the grid and follows a random exploration policy over all episodes. (d) gives the TAs used for their respective task specifications: get coffee, then go upstairs; get coffee, serve to guest on a couch, then go upstairs; and get coffee, serve to guest on a couch, turn the TV on, then go upstairs.

For all experiments, the convergence criterion was taken to be the point when the difference between the transition matrix estimates over one pass through the training data achieves a maximum absolute row sum less than 10^{-6} . In practice, this is sufficient to ensure that a structurally correct product MDP is learned.

The episode length and number of episodes used in the training data are both hyperparameters that can affect the convergence time and the structural correctness of the learnt transition matrix. In practice, choosing the episode length so that the agent achieves (on average) the task in 20 – 50% of the episodes ensures sufficient exploration of all states in the product MDP. Additionally, a trend was observed that for larger product MDPs, more training episodes were required to obtain a structurally correct estimate of the transition matrix. The number of episodes sufficient for convergence in our experiments was less in our model-based approach than in model-free approaches in related work (Section 5). Exact hyperparameters used and full results are presented in [2, Table 1, Appendix E].

Step 1 of our algorithm (learning the product MDP) is implemen-

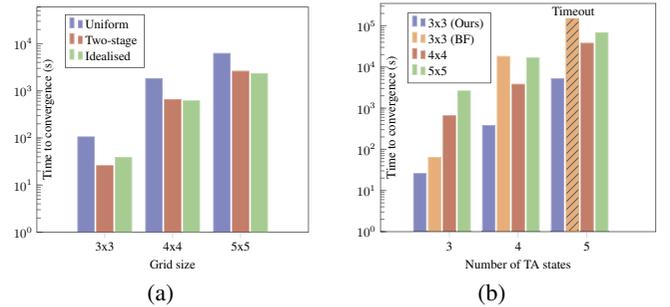


Figure 5: Step 1’s convergence times (a) comparing our pipeline (red) with two baseline approaches, a uniform prior (blue) and an idealised prior in which the spatial MDP is known (green) (b) for varying MDP and TA sizes.

ted in C++ (see [2, Appendices D and E] for more details), and the remaining steps are implemented in Python. All experiments are the average of 3 runs (with negligible standard errors) and were conducted on an Intel Xeon Gold 6248 (2.50GHz) CPU. Cone Lumping (Step 2) determines the NFA underlying the product MDP for the 5x5 grid and 5-state TA in Figure 1 in 0.43s, and removing environmental bias (Step 3) takes 0.8ms. This also confirms the usefulness of these steps for other purposes (argued in Sections 3.2 and 3.3).

We, therefore, turn to Step 1, Algorithm 1’s bottleneck. Although, note that the question of whether there is an automaton with n states that agrees with a finite set of data is NP-complete [14]. Recall from Section 3.1 that we learn our product MDP estimate $\hat{\mathcal{M}} \otimes \hat{\mathcal{A}}$ using the Baum-Welch algorithm in two stages. First, we learn an estimate of the spatial MDP \mathcal{M} ’s transition probability distribution \hat{P} . Then, this is used as an inductive bias for learning the transition probability distribution \hat{P}^{\otimes} of the full product MDP. Figure 5a demonstrates the merits of this *two-stage* approach using two baselines: i) a naive *uniform* approach, which learns a product MDP estimate using a uniform initialisation (each row in the transition matrix estimate was initialised to a uniform distribution), ii) an *idealised* approach which assumes that the agent has full a priori knowledge of the true spatial MDP \mathcal{M} ; this is encoded in the initial estimate of the product MDP’s transition function. Run time was chosen to measure performance instead of the number of iterations because the latter depends on the number of training steps per episode and the number of episodes generated, which varies across experiments.

Figure 5a demonstrates our approach’s superiority over the uniform baseline. We plot the convergence times of Step 1 for our approach (in red) against the two baselines for the 3 state TA and 3 different grid sizes. As the size of the grid increases, the time taken to learn an accurate estimate of the product MDP trends towards the idealised baseline. This is because the time taken to learn \hat{P} in the first stage becomes negligible compared with learning the TA. Thus, learning the TA via a larger (and therefore ostensibly more complex) structure first, the product MDP, does not noticeably increase the time complexity of the process. This also demonstrates why our process is useful for transfer learning (Section 3.2).

Figure 5b shows the dependence of the convergence time on the sizes of the MDP environment and TA. Increasing the TA size increases the number of sub-tasks the agent must complete before receiving reward. This Figure also shows how our algorithm’s efficiency compares with related work. We presented the codebases of five related works (see [2, Appendix F] for details) the same sequences used in our smallest experiment (3x3 grid with a 3-state TA), but none of these could successfully learn the TA (or threw a recursion error). We were, however, able to perform some bench-

marks against an implementation of Biermann and Feldman [8]’s SAT-based algorithm in the libalf library to infer a hypothesis reward machine consistent with the data for the 3- and 4-state TAs. We used the same set of sequences that were used in our algorithm for comparison, which were labelled as positive or negative sequences depending on whether the agent received a reward. With a 4-state TA, ours had a lower convergence time in a 4x4 gridworld than BF managed in even a 3x3 gridworld, and BF timed out at 150,000s when trying to learn the 5-state TA in the 3x3 gridworld.

5 Related Research

SAT-based approaches [1, 9, 16, 34, 36] generate a hypothesis automaton (equivalent to a temporal specification) and then verify whether this agrees with the agent’s observations during exploration. We showed that these are less efficient than our approach by benchmarking against the only available codebase that learnt an automaton from our traces. Meanwhile, Icarte et al [32] use *Tabu search*, which relies on already knowing a partial model, and Furelos-Blanco et al [12] use an *Answer Set Programming* algorithm, which assumes a known upper bound for the maximum finite distance between TA states. Our assumptions are weaker than these as we are only required to guess an upper bound on the number of TA states (in Section 4, we explained why this is easy to do) and we do not require any a priori knowledge about the spatial MDP. Also, all of these approaches are local, whereas ours is complete in the sense that our TA captures all sequential behaviour seen in the traces (high-level features).

[21, 28, 33] all consider probabilistic approaches to learning specifications. Their maximum entropy approaches (inspired by Ziebart et al [39]’s variant of inverse reinforcement learning) start from an initial pool of candidate specifications. In contrast, our maximum-likelihood approach does not a priori require any structure of the specification or the spatial MDP environment. Meanwhile, [11, 13, 27, 31, 37] use Angluin [5]’s L^* algorithm to learn a TA, relying on an *oracle* for equivalence and membership queries. We assume that the agent cannot access an oracle and must learn the TA fully autonomously, which aligns with the standard setup of model-free RL (note that L^* was not originally developed for RL applications). Because this means our training data cannot be provided as input to L^* , we cannot benchmark fairly against these approaches.

Learning an automaton from data is NP-hard [14], so under standard complexity assumptions, the time taken to solve this problem will necessarily grow exponentially as a function of the size of the TA. The most appropriate method will depend on the use-case as different approaches choose to relax different assumptions. For example, Corazza et al [9] and Dohmen et al [11] extend the SAT-based approach to noisy rewards and the L^* approach to learning probabilistic automata, respectively.

Our method is most appropriate in the following situations. First, when one does not assume the presence of an oracle (required for L^*). Second, if the reward environment is sparse, i.e. where the agent must follow a sequence of non-rewarding steps before getting a reward if and only it completes the full task. Third, when sample-efficiency is critical (as in many real-world scenarios) since our model-based approach requires fewer episodes to learn the TA. Fourth, for transfer learning: our approach can readily handle the underlying MDP or TA changing slightly. For example, suppose the initial task specification is to get coffee for the guest on the couch. Then, the furniture is rearranged and the task changes from collecting coffee to collecting tea. Only the affected part of the product MDP needs to be re-learned before performing the subsequent computation-

ally cheap steps (2 and 3) in our algorithmic pipeline. Other methods would need to start learning the new TA from scratch each time. Finally, the stochastic nature of our approach has advantages over the discrete approaches of related work; e.g. we can easily incorporate active-learning to more intelligently guide the agent’s exploration [4]. Furthermore, our stochastic approach allows for entropy-based uncertainty quantification and it will always generate an estimate of the solution even with few exploration episodes.

6 Conclusions and Future work

We have introduced an unsupervised pipeline for an RL agent to learn task specifications, encoded as deterministic, finite automata known as ‘task automata’, in unknown environments with sparse and non-Markovian rewards. The agent first learns a product MDP – a composition of the task automaton with the environment’s MDP before distilling out a task automaton using our efficient ‘Cone Lumping’ method. Finally, our post-processing routine, which is useful for any method that learns automata in RL, minimises environmental bias and simplifies the learnt task automaton.

As future work, we are extending to environments with multiple tasks (i.e., cases where an agent can achieve reward by satisfying multiple different task specifications); introducing active learning to further improve the efficiency of the rate-limiting Step 1; and trying different techniques for learning HMMs/POMDPs (e.g., spectral learning [18]). We are also interested in multi-agent settings, where agents get rewarded according to different tasks or have to optimally share information to learn a common task. Finally, we intend to extend to cases where the task automata encode ω -regular LTL properties, which require evaluation over infinite trajectories.

References

- [1] Eden Abadi and Ronen I Brafman, ‘Learning and solving regular decision processes’, in *29th International Joint Conference on Artificial Intelligence, IJCAI 2020*, pp. 1948–1954. International Joint Conferences on Artificial Intelligence, (2020).
- [2] Alessandro Abate, Yousif Almulla, James Fox, David Hyland, and Michael Wooldridge, ‘Learning task automata for reinforcement learning using hidden markov models’, *arXiv preprint arXiv:2208.11838*, (2022).
- [3] Dario Amodei, Chris Olah, Jacob Steinhardt, Paul Christiano, John Schulman, and Dan Mané, ‘Concrete problems in ai safety’, *arXiv preprint arXiv:1606.06565*, (2016).
- [4] Brigham Anderson and Andrew Moore, ‘Active learning for hidden markov models: Objective functions and algorithms’, in *Proceedings of the 22nd international conference on Machine learning*, pp. 9–16, (2005).
- [5] Dana Angluin, ‘Learning regular sets from queries and counterexamples’, *Information and computation*, **75**(2), 87–106, (1987).
- [6] Christel Baier and Joost-Pieter Katoen, *Principles of model checking*, MIT press, 2008.
- [7] Leonard E Baum and Ted Petrie, ‘Statistical inference for probabilistic functions of finite state markov chains’, *The annals of mathematical statistics*, **37**(6), 1554–1563, (1966).
- [8] Alan W Biermann and Jerome A Feldman, ‘On the synthesis of finite-state machines from samples of their behavior’, *IEEE transactions on Computers*, **100**(6), 592–597, (1972).
- [9] Jan Corazza, Ivan Gavran, and Daniel Neider, ‘Reinforcement learning with stochastic reward machines’, in *Proceedings of*

- the AAAI Conference on Artificial Intelligence, volume 36, pp. 6429–6436, (2022).
- [10] Salem Derisavi, Holger Hermanns, and William H Sanders, ‘Optimal state-space lumping in markov chains’, *Information Processing Letters*, **87**(6), 309–315, (2003).
- [11] Taylor Dohmen, Noah Topper, George Atia, Andre Beckus, Ashutosh Trivedi, and Alvaro Velasquez, ‘Inferring probabilistic reward machines from non-markovian reward signals for reinforcement learning’, in *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 32, pp. 574–582, (2022).
- [12] Daniel Furelos-Blanco, Mark Law, Anders Jonsson, Krysia Broda, and Alessandra Russo, ‘Induction and exploitation of subgoal automata for reinforcement learning’, *Journal of Artificial Intelligence Research*, **70**, 1031–1116, (2021).
- [13] Maor Gaon and Ronen Brafman, ‘Reinforcement learning with non-markovian rewards’, in *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 3980–3987, (2020).
- [14] E Mark Gold, ‘Complexity of automaton identification from given data’, *Information and control*, **37**(3), 302–320, (1978).
- [15] Mohammadhosein Hasanbeig, Alessandro Abate, and Daniel Kroening, ‘Logically-constrained reinforcement learning’, *arXiv preprint arXiv:1801.08099*, (2018).
- [16] Mohammadhosein Hasanbeig, Natasha Yogananda Jeppu, Alessandro Abate, Tom Melham, and Daniel Kroening, ‘Deep-synth: Automata synthesis for automatic task segmentation in deep reinforcement learning’, in *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 7647–7656, (2021).
- [17] John Hopcroft, ‘An $n \log n$ algorithm for minimizing states in a finite automaton’, in *Theory of machines and computations*, 189–196, Elsevier, (1971).
- [18] Daniel Hsu, Sham M Kakade, and Tong Zhang, ‘A spectral algorithm for learning hidden markov models’, *Journal of Computer and System Sciences*, **78**(5), 1460–1480, (2012).
- [19] Rodrigo Toro Icarte, Toryn Q Klassen, Richard Valenzano, and Sheila A McIlraith, ‘Reward machines: Exploiting reward function structure in reinforcement learning’, *Journal of Artificial Intelligence Research*, **73**, 173–208, (2022).
- [20] K. G. Larsen and A. Skou, ‘Bisimulation through probabilistic testing’, *Information and computation*, **94**(1), 1–28, (1991).
- [21] Niklas Lauffer, Beyazit Yalcinkaya, Marcell Vazquez-Chanlatte, Ameesh Shah, and Sanjit A Seshia, ‘Learning deterministic finite automata decompositions from examples and demonstrations’, in *Conference on Formal Methods in Computer-Aided Design—FMCAD 2022*, p. 325.
- [22] Michael L Littman, Ufuk Topcu, Jie Fu, Charles Isbell, Min Wen, and James MacGlashan, ‘Environment-independent task specifications via GLTL’, *arXiv preprint arXiv:1704.04341*, (2017).
- [23] Shubham Pateria, Budhitama Subagdja, Ah-hwee Tan, and Chai Quek, ‘Hierarchical reinforcement learning: A comprehensive survey’, *ACM Computing Surveys (CSUR)*, **54**(5), 1–35, (2021).
- [24] Michael O Rabin, ‘Probabilistic automata’, *Information and control*, **6**(3), 230–245, (1963).
- [25] Michael O Rabin and Dana Scott, ‘Finite automata and their decision problems’, *IBM journal of research and development*, **3**(2), 114–125, (1959).
- [26] Lawrence R Rabiner, ‘A tutorial on hidden markov models and selected applications in speech recognition’, *Proceedings of the IEEE*, **77**(2), 257–286, (1989).
- [27] Gavin Rens, Jean-François Raskin, Raphaël Reynoaud, and Giuseppe Marra, ‘Online learning of non-markovian reward models’, in *Proceedings of the Thirteenth International Conference on Agents and Artificial Intelligence*, volume 2, pp. 74–66. Scitepress, (2020).
- [28] Raeid Saqur, ‘Reward learning using structural motifs in inverse reinforcement learning’, *arXiv preprint arXiv:2209.13489*, (2022).
- [29] Matthew E Taylor and Peter Stone, ‘Cross-domain transfer for reinforcement learning’, in *Proceedings of the 24th international conference on Machine learning*, pp. 879–886, (2007).
- [30] Sylvie Thiébaux, Charles Gretton, John Slaney, David Price, and Froduald Kabanza, ‘Decision-theoretic planning with non-markovian rewards’, *Journal of Artificial Intelligence Research*, **25**, 17–74, (2006).
- [31] Noah Topper, George Atia, Ashutosh Trivedi, and Alvaro Velasquez, ‘Active grammatical inference for non-markovian planning’, in *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 32, pp. 647–651, (2022).
- [32] Rodrigo Toro Icarte, Ethan Waldie, Toryn Klassen, Rick Valenzano, Margarita Castro, and Sheila McIlraith, ‘Learning reward machines for partially observable reinforcement learning’, *Advances in Neural Information Processing Systems*, **32**, 15523–15534, (2019).
- [33] Marcell Vazquez-Chanlatte, Susmit Jha, Ashish Tiwari, Mark K Ho, and Sanjit Seshia, ‘Learning task specifications from demonstrations’, *Advances in Neural Information Processing Systems*, **31**, 5367–5377, (2018).
- [34] Christos Verginis, Cevahir Koprulu, Sandeep Chinchali, and Ufuk Topcu, ‘Joint learning of reward machines and policies in environments with partially known semantics’, *arXiv preprint arXiv:2204.11833*, (2022).
- [35] Tingwu Wang, Xuchan Bao, Ignasi Clavera, Jerrick Hoang, Yeming Wen, Eric Langlois, Shunshi Zhang, Guodong Zhang, Pieter Abbeel, and Jimmy Ba, ‘Benchmarking model-based reinforcement learning’, *arXiv preprint arXiv:1907.02057*, (2019).
- [36] Zhe Xu, Ivan Gavran, Yousef Ahmad, Rupak Majumdar, Daniel Neider, Ufuk Topcu, and Bo Wu, ‘Joint inference of reward machines and policies for reinforcement learning’, in *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 30, pp. 590–598, (2020).
- [37] Zhe Xu, Bo Wu, Aditya Ojha, Daniel Neider, and Ufuk Topcu, ‘Active finite reward automaton inference and reinforcement learning using queries and counterexamples’, in *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*, pp. 115–135. Springer, (2021).
- [38] Fanny Yang, Sivaraman Balakrishnan, and Martin J Wainwright, ‘Statistical and computational guarantees for the baumwielch algorithm’, in *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pp. 658–665. IEEE, (2015).
- [39] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, Anind K Dey, et al., ‘Maximum entropy inverse reinforcement learning’, in *AAAI*, volume 8, pp. 1433–1438. Chicago, IL, USA, (2008).