

Finite-State Controllers for (Hidden-Model) POMDPs using Deep Reinforcement Learning

David Hudák
Brno University of Technology
Brno, Czech Republic
ihudak@fit.vutbr.cz

Maris F. L. Galesloot
Radboud University
Nijmegen, The Netherlands
maris.galesloot@ru.nl

Martin Tappler
TU Wien
Vienna, Austria
martin.tappler@tuwien.ac.at

Martin Kurečka
Masaryk University
Brno, Czech Republic
martin.kurecka@fi.muni.cz

Nils Jansen
Ruhr-Universität Bochum &
Radboud University
Bochum, Germany
n.jansen@rub.de

Milan Češka
Brno University of Technology
Brno, Czech Republic
ceskam@fit.vutbr.cz

ABSTRACT

Solving partially observable Markov decision processes (POMDPs) requires computing policies under imperfect state information. Despite recent advances, the scalability of existing POMDP solvers remains limited. Moreover, many settings require a policy that is *robust* across multiple POMDPs, further aggravating the scalability issue. We propose the LEXPOP framework for POMDP solving. LEXPOP (1) employs deep reinforcement learning to train a *neural policy*, represented by a *recurrent neural network*, and (2) constructs a *finite-state controller* mimicking the neural policy through efficient extraction methods. Crucially, unlike neural policies, such controllers can be formally evaluated, providing performance guarantees. We extend LEXPOP to compute robust policies for *hidden-model* POMDPs (HM-POMDPs), which describe finite sets of POMDPs. We associate every extracted controller with its worst-case POMDP. Using a set of such POMDPs, we iteratively train a robust neural policy and consequently extract a robust controller. Our experiments show that on problems with large state spaces, LEXPOP outperforms state-of-the-art solvers for POMDPs as well as HM-POMDPs.

KEYWORDS

POMDPs; Deep Reinforcement Learning; Planning; Finite-State Controllers; Robustness; Model Uncertainty; Synthesis; Verification

ACM Reference Format:

David Hudák, Maris F. L. Galesloot, Martin Tappler, Martin Kurečka, Nils Jansen, and Milan Češka. 2026. Finite-State Controllers for (Hidden-Model) POMDPs using Deep Reinforcement Learning. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 10 pages. <https://doi.org/10.65109/TBFQ5922>

1 INTRODUCTION

Partially observable Markov decision processes (POMDPs) [37] are a ubiquitous model for automated sequential decision-making under imperfect state information. Applications for POMDPs range from

robotics [42, 43], over network protocols [20], to healthcare [80]. Solving a given discrete POMDP, i.e., computing policies that maximize expected rewards, is intractable in general [48], as optimal policies may require an infinite memory. State-of-the-art *model-based* POMDP solvers thus leverage various approximation techniques or heuristics. Here, the primary focus is on (1) the analysis of belief states representing probability distributions over the unobservable states [29, 43, 65] and (2) symbolic search techniques over finite-state policies [4, 41]. Furthermore, since exact environment models are rarely available, recent research has considered computing *robust* policies that provide sufficient performance over a given set of POMDPs [18, 22, 23]. A joint limitation of the above model-based methods is their scalability with respect to the model size. These methods must either (1) unfold a very large belief space or (2) explore a very large space of potential finite-state policies.

In contrast to model-based methods, model-free deep reinforcement learning (DRL) provides a means to mitigate these scalability challenges. DRL for POMDPs uses *recurrent neural networks* (RNNs) to represent memory, allowing policies to generalize behavior for similar states and observations using only interactions with the environment [38, 68]. While DRL scales well to large state spaces, there are two crucial obstacles. First, aligning DRL with our goal of solving a given (set of) POMDPs presents a significant technical challenge, as DRL must be set up correctly to achieve well-performing policies, requiring a highly efficient pipeline, combined with insights from the literature on addressing issues of sparse rewards, robustness, and partial observability [19, 56]. To the best of our knowledge, such a pipeline using DRL for solving discrete POMDPs has not yet been achieved. Second, the exact evaluation of DRL policies is not straightforward. The resulting lack of guarantees regarding the performance of a policy restricts the utilization of DRL to games [76] or problems where soft guarantees are sufficient, such as language models [59]. Using POMDPs in safety-critical domains requires means to exactly evaluate the resulting policies.

Contributions. Our main contribution is LEXPOP – Learning-based policy EXtraction for POMDP Planning – that yields high-quality finite-state policies with verified performance. LEXPOP builds on two key insights: (1) DRL, when engineered correctly, scales up to environments that are beyond reach of existing POMDP solvers. (2) Finite-state approximations of RNN-based policies often achieve



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/TBFQ5922>

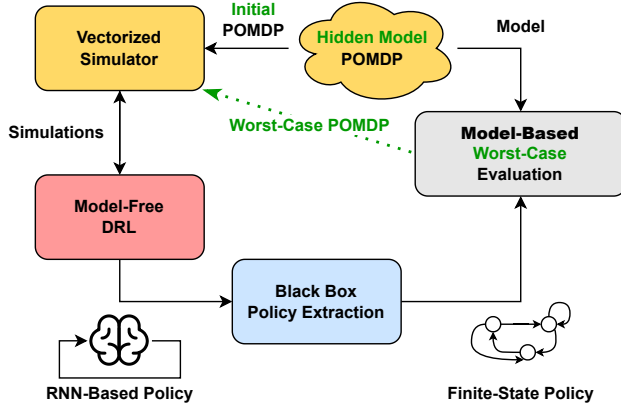


Figure 1: The overall idea of LEXPOP, our RL-based FSC extraction. The robust extensions are depicted in green.

comparable performance while enabling computationally tractable model-based evaluation, providing the required guarantees. We combine the strengths of these two directions.

Figure 1 depicts the three steps of LEXPOP: (1) Training an RNN-based neural policy using DRL equipped with a vectorized simulator that provides highly efficient access to the known POMDP model. (2) Extracting a stochastic finite-state controller (FSC) that mimics the behavior of the neural policy. (3) Model-based verification of the extracted FSC. For extraction, we use two sampling-based methods: automata learning based on ALERGIA [49] and a novel, scalable method that learns surrogate, self-interpretable networks, providing flexible control over FSC sizes. Both methods are policy-architecture agnostic and treat policies as black boxes. We discuss prior work on using RNNs to find FSCs [11, 23] in the related work section.

We extend our approach towards *robust* FSCs for *hidden-model* POMDPs (HM-POMDPs) [7, 22], capturing different environments, i.e., sets of POMDPs with a shared action and observation space. A robustified version of LEXPOP searches for an FSC that achieves the best worst-case performance for a given HM-POMDP. As depicted by the green parts of Figure 1, for the candidate FSCs, we identify the worst-case POMDPs that iteratively enter the learning step. Similar to Galesloot et al. [22], a deductive verification technique [5] finds the worst-case POMDPs. We sample trajectories from these worst-case POMDPs to train a robust neural policy and, subsequently, extract a robust FSC.

In summary, our main contributions are:

- (1) LEXPOP, a scalable learning-based approach for computing high-quality FSCs for large POMDPs,
- (2) an extension of LEXPOP computing robust FSCs for large HM-POMDPs via iterative learning from worst-case POMDPs,
- (3) a scalable extraction method that approximates complex neural policies via stochastic FSCs.

In our experimental evaluation, we focus on planning problems with large state spaces that go beyond the reach of existing model-based approaches. We demonstrate a clear advantage in terms of

scalability when comparing LEXPOP to state-of-the-art approaches in both the single POMDP and HM-POMDP settings.

2 PROBLEM FORMULATION

In this section, we begin with notation and preliminaries, then formulate the central problems addressed in the paper.

The set of all distributions over a countable set X is $\Delta(X)$. For $f: X \rightarrow \Delta(Y)$, we may write $f(y | x)$ to denote $f(x)(y)$. Below, we introduce (partially observable) *Markov decision processes* [62].

2.1 MDPs, POMDPs and Policies

DEFINITION 1 (MDP). A *Markov decision process (MDP)* is a tuple $M = \langle S, s_0, A, T, R \rangle$, where S represents a countable set of states, $s_0 \in S$ is the initial state, A is a finite set of actions, $T: S \times A \rightarrow \Delta(S)$ is a transition function, and $R: S \times A \rightarrow \mathbb{R}$ is a reward function.

A *Markov chain (MC)* with rewards is an MDP without actions. To handle environments that do not provide precise state information, MDPs are extended to model partially observable states.

DEFINITION 2 (POMDP). A *partially observable Markov decision process (POMDP)* [37] is a tuple $\mathcal{M} = (M, Z, O)$, where M is an underlying MDP, Z is a finite set of observations, and $O: S \rightarrow Z$ is, w.l.o.g. [15], a deterministic and state-based observation function.

We focus on a general form of *reachability* over an infinite horizon with rewards defined w.r.t. a set of *target states* $G \subseteq S$ that are fully observable and absorbing [31]. It encompasses (1) *reachability*; maximizing reachability probabilities of target states G , and (2) *reachability reward*; maximizing cumulative reward before reaching target states G with probability one. For (1), the reward function is one for states $g \in G$ and zero everywhere else. For (2), the states $g \in G$ have no reward and are reachable from all states. Note that the reachability reward objective (2) generalizes the discounted reward setting [6]. For a sequence of states (s_0, s_1, \dots, s_t) its observable fragment is the *observation history* $o_t = (O(s_0), O(s_1), \dots, O(s_t))$. A *trajectory* is a sequence containing observations and actions. A *policy* $\pi \in \Pi$ maps trajectories to distributions over actions. We search for policies represented by finite-state controllers [26, 52].

DEFINITION 3 (FSC). A *finite-state controller (FSC)* is a tuple $\pi_F = (N, n_0, \delta, \eta)$, where N is a set of memory nodes (states), $n_0 \in N$ is an initial node, $\delta: N \times Z \rightarrow \Delta(A)$ is a (stochastic) action function, and $\eta: N \times Z \rightarrow \Delta(N)$ is a (stochastic) memory update function.

We employ stochastic Mealy-like FSCs [2] that select actions based on memory and observations and thus require fewer memory nodes compared to Moore-like FSCs as used, e.g., in You et al. [77].

The (uncountable) set of all (stochastic) FSCs for a POMDP \mathcal{M} is $\mathcal{F}_{\mathcal{M}}$. A POMDP \mathcal{M} together with an FSC π_F constitutes the Markov chain $\mathcal{M}^{\pi_F} = \langle S^{\pi_F}, \langle s_0, n_0 \rangle, T^{\pi_F}, R^{\pi_F} \rangle$ over the product state-space $S^{\pi_F} = S \times N$, with, for all $\langle s, n \rangle \in S^{\pi_F}$, using $z = O(s)$:

$$T^{\pi_F}(\langle s', n' \rangle | \langle s, n \rangle) = \eta(n' | n, z) \sum_{a \in A} \delta(a | n, z) T(s' | s, a),$$

$$R^{\pi_F}(\langle s, n \rangle) = \sum_{a \in A} \delta(a | n, z) R(s, a).$$

Consequently, the *value-function* of an FSC π_F in a POMDP \mathcal{M} is the value-function $V^{\pi_F}: S \times N \rightarrow \mathbb{R}$ of the MC \mathcal{M}^{π_F} , given by the fixed point of the following equation [52], again using $z = O(s)$:

$$V^{\pi_F}(\langle s, n \rangle) = R^{\pi_F}(\langle s, n \rangle) + \sum_{\langle s', n' \rangle} T^{\pi_F}(\langle s', n' \rangle | \langle s, n \rangle) V^{\pi_F}(\langle s', n' \rangle).$$

Then, the value $J^{\pi_F} = V^{\pi_F}(\langle s_0, n_0 \rangle)$ of the FSC π_F is determined from the value function on the initial state of the MC \mathcal{M}^{π_F} .

Neural policies. Recurrent neural networks (RNN), similarly to FSCs, represent state machines. Nowadays, RNNs are a standard architecture implemented by gated recurrent units (GRU) [16], or long-short term memory (LSTM) [30] layers that give DRL the capability of memory [27, 56] and have been used to find FSCs for planning in (robust) POMDPs [11, 23]. In contrast to FSCs, RNNs have an infinite space of memory, represented by the *hidden state* $h \in \mathbb{R}^d$, with size $d \in \mathbb{N}$. RNN policies typically consist of a parameterized (deterministic) memory update $\eta_\theta: \mathbb{R}^d \times Z \rightarrow \mathbb{R}^d$ and a parameterized (stochastic) action function $\delta_\theta: \mathbb{R}^d \rightarrow \Delta(A)$. Parameters θ are optimized by back-propagating the gradients of DRL losses through time [72, 73]. Just like an FSC, an RNN-based policy maps trajectories to actions and induces a value $J^{\pi_{\text{RNN}\theta}}$ representing the expected reward. It is, in general, intractable to compute it in closed form; thus, the value $J^{\pi_{\text{RNN}\theta}}$ is typically approximated empirically by simulating the policy $\pi_{\text{RNN}\theta}$ in the POMDP \mathcal{M} .

2.2 Hidden-model POMDPs

Hidden-model POMDPs (HM-POMDPs) model a family of POMDPs that share the same state, observation, and action spaces but differ in their reward and transition functions. As such, HM-POMDPs capture uncertainty about the exact dynamics of the environments. Here, we largely follow Galesloot et al. [22].

DEFINITION 4 (HM-POMDP). A hidden-model POMDP (HM-POMDP) [22] is a tuple $\mathcal{M} = \langle S, s_0, A, \{T_i\}_{i \in I}, \{R_i\}_{i \in I}, Z, O \rangle$, where I is a finite set of indices, S, s_0, A, Z , and O are defined as for standard POMDPs, $\{T_i\}_{i \in I}$ and $\{R_i\}_{i \in I}$ are indexed sets of transition and reward functions, respectively, such that, for every $i \in I$, $\mathcal{M}_i = \langle S, s_0, A, T_i, R_i, Z, O \rangle$ is a POMDP. Consequently, an HM-POMDP \mathcal{M} describes a finite set (a family) of POMDPs $\{\mathcal{M}_i\}_{i \in I}$.

Due to the shared action and observation spaces, all POMDPs in the HM-POMDP share the same set of FSCs $\mathcal{F}_{\mathcal{M}}$. Additionally, structural similarities, i.e., transitions with equivalent outcomes, can be merged in the construction of a *quotient POMDP* [22]. Given a set of POMDPs, an FSC π_F induces a set (family) of MCs $\{\mathcal{M}_i^{\pi_F}\}_{i \in I}$ [14] and an associated set of value-functions $\{V_i^{\pi_F}\}_{i \in I}$. We are interested in the robust value, i.e., the worst-case value among the models.

DEFINITION 5 (ROBUST VALUE). Given an HM-POMDP \mathcal{M} and a policy π_F , the robust value is the worst-case value among the set:

$$\mathcal{J}^{\pi_F} = \min_{i \in I} V_i^{\pi_F}(\langle s_0, n_0 \rangle). \quad (1)$$

As explained in Galesloot et al. [22], the robust value can be found efficiently through deductive verification techniques [5].

2.3 Problem Statement

Now, we are ready to define the main problems: maximizing reachability rewards (or probabilities) for either a single POMDP or an HM-POMDP. First, we define our problem for individual POMDPs.

PROBLEM 1. Given a POMDP \mathcal{M} , find an FSC policy π_F that maximizes the value $V^{\pi_F}(\langle s_0, n_0 \rangle)$. Formally:

$$\pi_F^* \in \sup_{\pi_F \in \mathcal{F}_{\mathcal{M}}} V^{\pi_F}(\langle s_0, n_0 \rangle). \quad (2)$$

Similarly, we define the robust problem for HM-POMDPs:

PROBLEM 2. Given an HM-POMDP \mathcal{M} , find an FSC policy π_F that maximizes the robust value \mathcal{J}^{π_F} as in Definition 5. Formally:

$$\pi_F^* \in \sup_{\pi_F \in \mathcal{F}_{\mathcal{M}}} \mathcal{J}^{\pi_F} = \sup_{\pi_F \in \mathcal{F}_{\mathcal{M}}} \min_{i \in I} V_i^{\pi_F}(\langle s_0, n_0 \rangle). \quad (3)$$

Note that Problem 2 resembles two-player zero-sum Stackelberg games [81], which suggests that stochastic FSCs are necessary [22]. Similarly to POMDP planning with undiscounted infinite-horizon objectives, the decision variants of both Problem 1 and Problem 2 are undecidable in general [47]. Although there exist convergence guarantees for a variant of an infinite-horizon objective called *Goal-POMDPs* [31], which were recently also extended to a game setting [69], the practical performance of belief-based methods is limited [4], and the ideas cannot immediately be used for Problem 2. In this work, we aim to develop a scalable and sound algorithm to find the best possible FSC policy within a reasonable timeframe.

To that end, we compute (stochastic) FSCs using DRL, specifically using RNN-based agents. In contrast to FSCs, the value of an RNN-based policy $\pi_{\text{RNN}\theta}$ cannot be found analytically, which is why we are ultimately interested in finding an FSC representation of $\pi_{\text{RNN}\theta}$.

3 ON USING DRL TO SOLVE POMDPs

In the following, we motivate the use of DRL with RNNs for solving POMDPs, even when the model is fully known.

3.1 Challenges of Solving POMDPs

In POMDPs, it is well-known that acting near-optimally requires a succinct representation of (past) observation-action sequences [64]. These sequences can be sufficiently compressed into *beliefs*, i.e., probability distributions over the underlying unobserved states, and a near-optimal policy can be found by constructing and solving a *belief-MDP* [37]. Unfortunately, the complexity of updating a belief is quadratic in the size of the state-space [66], and the number of possible belief states grows exponentially in the horizon, rendering full exploration of belief-space intractable in general [47].

Searching for FSCs. Exploration of FSCs offers an alternative method to belief-based policies with benefits in efficient verification and deployment. One direction is to optimize the parameters of an FSC, where the size and structure of the memory must be established beforehand [2, 36, 61]. An alternative direction is to build the FSC incrementally through various symbolic search techniques [4, 41]. However, both directions suffer from an enormous design space.

Robust policies. For many situations, the assumption that a single POMDP model describes all cases is too strong. Therefore, recent efforts focus on computing a robust policy for sets of POMDPs [18, 22, 23, 58]. In a robust setting, the challenge is to find a single policy that works for (large) sets of POMDPs. Thus, the scalability challenges and the importance of generalization are further increased, and existing approaches for POMDPs do not readily extend.

Algorithm 1: Overview of LEXPOP and Robust LEXPOP

```

1: function TRAINPOLICY( $\pi_{\text{RNN}\theta}$ ,  $\{\mathcal{M}_i\}_{i \in I'}$ )
2:   VecSimulator  $\leftarrow$  VECSTORM( $\{\mathcal{M}_i\}_{i \in I'}$ )
3:    $\pi'_{\text{RNN}\theta} \leftarrow$  DRL( $\pi_{\text{RNN}\theta}$ , VecSimulator)  $\triangleright$  Using PPO
4:   return  $\pi'_{\text{RNN}\theta}$ 
5: function LEXPOP( $\mathcal{M}$ : POMDP,  $\pi_{\text{RNN}\theta}$  : Policy)
6:    $\pi_{\text{RNN}\theta} \leftarrow$  TRAINPOLICY( $\pi_{\text{RNN}\theta}$ ,  $\{\mathcal{M}\}$ )
7:    $\pi_F \leftarrow$  EXTRACTFSC( $\pi_{\text{RNN}\theta}$ )  $\triangleright$  Using Alergia or SIG
8:    $J^{\pi_F} \leftarrow$  EVALUATE( $\pi_F$ ,  $\mathcal{M}$ )  $\triangleright$  Using STORM
9:   return  $\pi_F, J^{\pi_F}$ 
10: function ROBUSTLEXPOP( $\mathcal{M}$ : HM-POMDP,  $\pi_{\text{RNN}\theta}$  : Policy)
11:    $B \leftarrow$  INITIALIZE()  $\triangleright B$  is a buffer of POMDPs
12:   while  $\neg$  TIMEOUT() do
13:      $\pi_{\text{RNN}\theta} \leftarrow$  TRAINPOLICY( $\pi_{\text{RNN}\theta}$ ,  $B$ )  $\triangleright$  See App. B in [34]
14:      $\pi_F \leftarrow$  EXTRACTFSC( $\pi_{\text{RNN}\theta}$ )  $\triangleright$  Using Alergia or SIG
15:      $\mathcal{M}, \mathcal{J}^{\pi_F} \leftarrow$  WORSTPOMDP( $\pi_F$ ,  $\mathcal{M}$ )  $\triangleright$  Using PAYNT
16:      $B \leftarrow B \cup \{\mathcal{M}\}$ 
17:   return  $\pi_F, \mathcal{J}^{\pi_F}$ 

```

3.2 Opportunities and Challenges of DRL

To address the aforementioned challenges, we leverage model-free DRL, which does not rely on a model and instead optimizes over its simulations, enabling scaling to large state spaces. DRL handles partial observability, by using RNNs as memory representations for policies [27, 56, 73], due to their capability of generalization [56] and learning sufficient statistics of observation sequences [45].

DRL for solving (HM-)POMDPs. In our setting, we are given a fully specified (HM-)POMDP model. As explained in the introduction, utilizing DRL to solve discrete (HM-)POMDPs presents its own challenges and remains an unsolved problem. We must compute a policy in symbolic form, specifically an FSC, such that we can evaluate its performance analytically. Thus, while we neglect the POMDP model in searching for the FSC, we still use it to provide guarantees on (i.e., verify) the policy. We observe that using RNNs to search for FSCs for a given POMDP has been very successful [11, 23]. However, so far, training the RNN has relied on model-based information, which suffers from the fact that (1) it does not scale and (2) relies on suboptimal approximations. LEXPOP fills this gap.

4 OVERVIEW OF OUR APPROACH

Algorithm 1 provides an overview of LEXPOP.

(Robust) policy training. LEXPOP is agnostic to the DRL algorithm used to train the neural policy. We chose PPO [63], due to its stable performance on our challenging POMDP benchmarks. For efficient sampling, we developed a vectorized simulation framework on top of the STORM model-checker, enabling parallel simulations across single or multiple POMDPs (see Line 2). In Line 3, our agent, represented by an RNN $\pi_{\text{RNN}\theta}$ with parameters θ , is optimized on a (new) set of POMDPs through multiple training iterations. DRL for the considered (HM-)POMDPs poses significant technical challenges (cf. [19]), including poorly defined and sparse rewards,

partial observability, and dynamic action spaces. We provide a description of our DRL pipeline, including important design choices and alternative learning methods in Appendix B in [34].

LEXPOP for Problem 1. Given a POMDP \mathcal{M} , LEXPOP trains a neural policy $\pi_{\text{RNN}\theta}$ starting from a randomly initialized policy $\pi_{\text{RNN}\theta}$ (Line 6). The key step is the extraction of the FSC π_F from $\pi_{\text{RNN}\theta}$ (Line 7). We consider two methods discussed in Section 5. Finally, π_F is analytically evaluated by constructing and analyzing the Markov chain \mathcal{M}^{π_F} in STORM [28], yielding its value J^{π_F} (Line 8).

LEXPOP for Problem 2. Given an HM-POMDP \mathcal{M} , ROBUSTLEXPOP iteratively trains a robust neural policy and consequently extracts a robust FSC policy that performs well on the given set of POMDPs. ROBUSTLEXPOP maintains a candidate neural policy $\pi_{\text{RNN}\theta}$ and its finite-state approximation π_F . Every candidate π_F is associated with its worst-case POMDP \mathcal{M} from \mathcal{M} . These worst-case POMDPs are iteratively stored in a buffer B . We start with a randomly initialized $\pi_{\text{RNN}\theta}$ and a randomly selected \mathcal{M} (Line 11). In every iteration, we improve $\pi_{\text{RNN}\theta}$ through robust policy training, using the trajectories collected from all POMDPs stored in the buffer B (Line 13). Appendix B.7 in [34] provides more details. Then, we extract an improved π_F (Line 14) using the same extraction methods as in the single-POMDP setting (see Section 5). Finally, we analytically compute the robust value \mathcal{J}^{π_F} associated with π_F and the corresponding worst-case POMDP \mathcal{M} (Line 15) which we add to the buffer B (Line 16). We employ PAYNT [5] to efficiently search through the family of Markov chains $\{\mathcal{M}_i^{\pi_F}\}_{i \in I}$ induced by \mathcal{M} and π_F .

5 POLICY EXTRACTION

The extraction of FSCs from RNN-based policies connects model-free policy learning and model-based policy verification. It relies on two insights: 1) for RNN policies, there often exists a smaller finite-state representation [39] and 2) for challenging POMDPs, there often exists a small FSC achieving a close-to-optimal value [4].

To retrieve an FSC, we propose two extraction methods with a common working principle: Provided an RNN policy $\pi_{\text{RNN}\theta}$, we synthesize an FSC policy π_F whose decisions match $\pi_{\text{RNN}\theta}$ on sampled training trajectories as closely as possible. Both methods learn stochastic FSCs from a sampled dataset of trajectories \mathcal{D} from a subset of POMDPs $\{\mathcal{M}_i\}_{i \in I'}$ using policy $\pi_{\text{RNN}\theta}$ trained on that subset. The data set contains sequences of tuples (z_t, a_t) of observations and actions selected by $\pi_{\text{RNN}\theta}$ at time t . The first extraction method is a direct application of the automata learning algorithm ALERGIA [13], whereas the second method introduces a novel RNN training approach with a hidden-state space encoding.

5.1 Automata Learning

We learn stochastic Mealy machines [54] using a variant of ALERGIA proposed for MDPs [49] and implemented in AALpy [55]. Therefore, we slightly adjust the FSC definition to $\pi_F = (N, n_0, \sigma)$, combining the memory update and action function into $\sigma : N \times Z \rightarrow \Delta(N \times A)$.

Given the dataset \mathcal{D} of trajectories, ALERGIA first transforms them into a prefix tree by merging common prefixes. Then, it merges tree nodes having compatible subtrees, i.e., nodes whose observed future behavior is similar. Compatibility checks recursively determine if the action distributions conditioned on observation-action




Figure 2: High-level architecture of the self-interpretable Gumbel softmax network.

histories are not statistically different. Therefore, merging compatible nodes creates an FSC memory structure that mimics $\pi_{\text{RNN}\theta}$. After processing all possible merges, ALERGIA estimates action and update probability distributions. Instead of commonly used compatibility checks based on Hoeffding bounds [13, 49], we use χ^2 homogeneity tests. This better distinguishes memory nodes, at the cost of slightly larger FSCs and, thus, longer verification times. In the limit, ALERGIA [49] converges to an FSC that is isomorphic to the FSC underlying the RNN policy, if such underlying FSC exists.

5.2 Self-Interpretable Networks

As automata learning typically struggles with large action and observation spaces, we additionally propose an approach that learns interpretable neural surrogate models. The core idea is to approximate the behavior of the policy $\pi_{\text{RNN}\theta}$ using a neural network $\pi_{\text{FSC}\phi}$, which is explicitly designed to emulate an FSC. The idea of using a discrete memory structure is inspired by *quantized bottleneck layers* proposed by Koul et al. [39], which employs a hyperbolic tangent (tanh) layer followed by rounding to quantize the RNN feedback. Training with quantization requires a straight-through gradient estimator. Our extraction method, *Self-Interpretable Gumbel Softmax Network* (SIG), mitigates three significant limitations arising from this design: a) an exponential blow-up of the number of memory nodes caused by the flat quantization of the activation values; b) the lack of support for stochastic memory updates; and c) the bias of the straight-through gradient estimates. To combat these issues, we propose a bottleneck architecture that directly models the transition distribution to one of $|N|$ memory states.

Architecture and Training. Figure 2 depicts the architecture for training the surrogate policy $\pi_{\text{FSC}\phi}$. The recurrent part transforms hidden states $h_t \in \mathbb{R}^d$ into a space of vectors $m_{t+1} \in \{0, 1\}^{|N|}$ that resembles categorical distributions over one-hot encoded memory nodes. The size $|N|$ is a parameter of extraction. The forward pass first computes the vector $p_t \in \mathbb{R}^{|N|}$ of non-normalized log-probabilities of the memory update. While the distilled FSC directly samples $m_{t+1} \sim \text{Softmax}(p_t)$ to update the memory state, this approach is not differentiable. Therefore, during training, we use a Gumbel Softmax layer [35] which is designed to sample m_{t+1} in a differentiable way. We pay the cost that m_{t+1} is no longer a one-hot vector, though its distribution is close to $\text{Softmax}(p_t)$ (in the Wasserstein distance); the distance is controlled by a temperature parameter and is reduced during training. To train the SIG network, we minimize the empirical cross-entropy between the actions from the dataset \mathcal{D} and the computed actions.

FSC Inference. To construct an FSC, we infer the actions and memory for each combination of memory node and input observation [23]. Given a trained SIG with a bottleneck layer of size $|N|$, we create a placeholder FSC π_F with memory nodes N . Then, for every node $n \in N$ and observation $z \in Z$, we perform these steps:

- (1) Create a one-hot encoding $m \in \{0, 1\}^{|N|}$ of n .
- (2) Process z and m through the RNN to yield an action distribution $\sigma_a \in \Delta(A)$ and hidden state $h \in \mathbb{R}^d$.
- (3) Process h with the Gumbel encoder to get a distribution $\sigma_{n'} \in \Delta(N)$ over the next memory node $n' \in N$.
- (4) Add $(n, z) \mapsto \sigma_a$ to δ and $(n, z) \mapsto \sigma_{n'}$ to η of the FSC π_F .

Then, the above yields a fully specified FSC π_F . We prune nodes that are not reachable from the initial observation-memory tuple. While SIG does not provide convergence guarantees like ALERGIA, we find that SIG works well and often even better in practice.

6 EXPERIMENTS

We derive two main research questions from the problems formulated in Section 2.3, and further refine them into sub-questions, considering the individual components of LEXPOP:

- RQ1.** Can LEXPOP construct FSCs with a higher value than FSCs synthesized by a state-of-the-art POMDP solver SAYNT [4]?
- RQ1a** Can the PPO algorithm train RNN-based neural policies that empirically outperform the value of FSCs from SAYNT?
- RQ1b** Does the extraction of FSCs from the neural policies preserve the value?
- RQ1c** Does the self-interpretable extraction improve the fidelity over stochastic automata learning using ALERGIA [49]?
- RQ2.** For the given HM-POMDPs, can LEXPOP construct robust FSC policies with a higher robust value than the FSCs produced by the state-of-the-art model-based algorithm rFPG [22]?
- RQ2a** Does the FSC extraction preserve the robust value of the robustified neural policy on the given HM-POMDPs?
- RQ2b** Can the robustified PPO algorithm, combined with FSC extraction, outperform the FSCs produced by rFPG?
- RQ2c** Is the worst-case POMDP selection important for solving HM-POMDPs?

Baseline Selection. RQ1 concerns the experimental comparison of DRL-based FSC synthesis and offline model-based planning. As a baseline, we selected one of the state-of-the-art approaches for offline POMDP planning, SAYNT [4]: it has publicly available code, supports both reachability and undiscounted reward specifications, and to our knowledge, works with the largest models. We also considered recent work by Ho et al. [29], which extends the heuristic search value iteration algorithm [65] to reachability specifications, outperforming SARSOP [43]. Our preliminary experiments show that its implementation cannot handle our large reachability models, which are about two orders of magnitude larger than those in [29].

Unfortunately, we were unable to run the approach described in Carr et al. [11], which also employs RNNs to extract FSCs. To our knowledge, the recently proposed rFPG algorithm [22] is the only available method that allows a direct comparison on HM-POMDPs within RQ2. The other approaches discussed in the related work, see Section 7, consider different formulations.

Table 1: Results for the single POMDP setting using 30-minute timeout. In all models, the goal is to maximize the value. For LEXPOP, we report the interquartile mean (IQM) of the empirical values of the neural policies and the IQM of the extracted FSCs, along with their average sizes, using the two extraction methods. The values in brackets show the interquartile range (IQR). We highlight the best value(s) with statistical significance, excluding the empirical value for LEXPOP.

Model	Dimensions			SAYNT Value	LEXPOP (RNN) IQM (IQR)	LEXPOP (Alergia)		LEXPOP (SIG)	
	S	Z	A			IQM (IQR)	FSC Size	IQM (IQR)	FSC Size
maze-10	4.8k	22	4	8.59	8.86 (0.03)	8.52 (0.10)	20.0	8.54 (0.47)	3.0
rocks-16	11k	2.7k	10	-36.91	-48.17 (1.43)	-51.49 (1.41)	6.0	-48.70 (1.28)	3.0
network-3-8-20	17k	2.2k	5	-10.45	-7.36 (0.27)	-8.07 (0.38)	5.0	-7.70 (0.15)	3.0
network-5-10-8	117k	3.7k	8	-16.12	-12.56 (0.20)	-13.76 (0.92)	11.0	-12.71 (0.14)	3.0
intercept-16	130k	21k	6	0.80	1.00 (0.00)	0.98 (0.01)	6.0	0.99 (0.01)	3.0
evade-n17	314k	158k	7	0.58	0.85 (0.02)	0.85 (0.01)	4.0	0.84 (0.01)	3.0
drone-2-8-1	520k	889	6	0.40	0.61 (0.01)	0.58 (0.01)	5.0	0.58 (0.01)	3.0

Table 2: Results for the HM-POMDP setting using 1-hour timeout. The goal is to maximize the robust value. We report the interquartile mean and the best robust value achieved by FSCs produced by rFPG and by the two variants of LEXPOP, together with the average reachable memory of the extracted FSCs. The highlighted value(s) are the best with statistical significance.

Model	Dimensions				rFPG		LEXPOP (Alergia)			LEXPOP (SIG)		
	S	I	Z	A	IQM	Best	IQM	Best	FSC Size	IQM	Best	FSC Size
rover	86	6000	18	5	299.24	299.40	290.17	296.14	7.70	294.68	297.65	2.8
obstacles-8-5	380	12000	25	9	-205.05	-193.93	-225.98	-212.46	2.50	-218.29	-207.63	2.9
network	4k	140	20	6	3.51	3.78	3.76	3.80	4.30	3.78	3.82	3.0
avoid	13k	1600	11	9	-174.23	-139.84	-161.01	-145.87	6.60	-163.15	-140.88	3.0
drone-2-6-1	44k	3200	432	7	0.01	0.04	0.59	0.62	7.50	0.59	0.62	10.0
avoid-large	45k	7056	11	9	-228.42	-166.14	-172.80	-154.02	4.40	-170.53	-141.63	8.8
moving_obstacles	262k	400	14	8	-2113.79	-2036.71	-158.98	-121.98	10.10	-86.86	-76.45	5.2

Benchmark Sets. The experimental evaluation is based on two benchmark sets. For RQ1, we use both original and enlarged POMDP models in the PRISM format [44] from the SAYNT repository [4]. For RQ2, we consider HM-POMDPs used to evaluate rFPG [22], excluding the DPM model with no observation features. Moreover, we constructed new HM-POMDPs by extending models used for RQ1 and by scaling up the original models in [22]. Appendix A in [34] provides more details on the benchmarks. The model statistics in Tables 1 and 2 are: |S| is the maximum number of reachable states across all POMDPs, and |I|, |Z|, |A| are the number of POMDPs, observations, and actions, respectively. The goal is to maximize the reachability probability or the reachability reward. Our benchmarks are challenging for DRL due to sparse rewards and dynamic discrete action spaces as discussed in Appendix B.6 in [34].

Experimental Setting. We run experiments on a Ryzen 7 7840U processor with 32 GB of RAM. LEXPOP uses the PPO [63] implementation in the TensorFlow Agents framework [25], and a vectorized simulator implemented in the JAX framework [9]. Verification of FSCs and worst-case POMDP synthesis are performed using the PAYNT [5] toolkit based on the Storm model checker [28]. Neural policies are evaluated using 512 independent simulations in a given POMDP, performed in parallel using the vectorized simulator. We truncate simulations after a fixed number of steps (typically 600) to have at least 512 finished episodes. For the single-POMDP setting, such an evaluation provides us with statistically reliable

estimates of the value. In the case of HM-POMDPs, the independent simulations are performed on a multiset of POMDPs in the current POMDP buffer (see Algorithm 1). Appendix D.1 in [34] provides more details. For automata learning, we use ALERGIA [49] implemented in AALpy [55]. The LEXPOP implementation is publicly available in [33].

RQ1: POMDP Planning

Setting. We train each neural policy for 4000 iterations on every model. Following DRL training, we extract FSCs with both extraction methods. The training takes about 23 minutes, while FSC extraction takes between 2 and 8 minutes. For a fair comparison, we configured SAYNT with a 30-minute timeout on the FSC synthesis. SAYNT is deterministic and therefore runs only once. For LEXPOP, we repeat each run 10 times with fixed random seeds and report the interquartile mean (IQM) and the interquartile range (IQR).

Main results for RQ1. Table 1 shows the FSC values achieved by SAYNT, the average empirical value of the learned neural policy, and the average value of the FSC extracted by LEXPOP from the neural policies using the two variants of the extraction method. The table also reports the average sizes of the FSCs produced by LEXPOP, see Appendix C.2 in [34] for more details.

RQ1a. Table 1 shows that the learned neural policies provide better empirical values than the verified value of the FSCs produced by SAYNT when the model gets larger, and comparable ones for

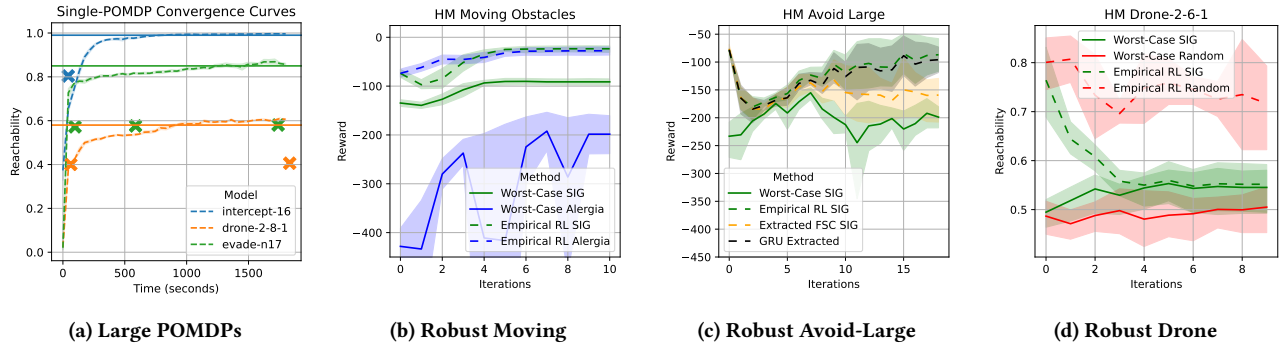


Figure 3: Selected convergence curves showing average values and IQR (if applicable): (a) Large single-POMDP benchmarks: Empirical probabilities of neural policies (dashed lines), and values of the FSC produced by SAYNT (crosses) and extracted from the final neural policy (solid lines). (b,c) Performance loss of the FSC extraction: Empirical value of the neural policy on the set of POMDPs used for training (dashed lines) and robust values of the extracted FSCs in the HM-POMDPs (solid lines). Additionally, in (c), the orange line shows the empirical value of the extracted FSCs on the set of POMDPs used for training, and the black line shows the empirical value achieved by the same self-interpretable network without the quantized layers on the same set of POMDPs. (d) Impact of the worst-case POMDP selection in HM-POMDPs: Lines as in (b). Green lines show the results using the worst-case POMDP selection, red lines show the results for random selection. Additional convergence curves are in Appendix E in [34].

small models, e.g., `rocks-16`, where SAYNT benefits from the full knowledge of the model. Figure 3 (a) provides a more in-depth performance assessment on large models. The plots compare the empirical values of the reachability probability achieved by neural policies (dashed lines) with the FSC values synthesized by SAYNT (crosses). We observe that for large models, SAYNT relies primarily on an initial memory-less policy and cannot improve over time, whereas PPO continuously improves the policy. The main bottleneck of SAYNT is the size of the considered FSCs, combined with the large number of beliefs that must be explored. **Summary: For a large majority of models, PPO learns policies with better empirical values than verified controllers produced by SAYNT.**

RQ1b. Next, we examine whether LEXPOP can extract an FSC from the given neural policy with sufficient fidelity. The results in Table 1 generally show that the extracted FSC policies perform similarly to the original neural policies and thus significantly outperform FSCs produced by SAYNT except for the `rocks-16` model. This holds for both extraction approaches, as confirmed by a Sign test, which yields a p-value of approximately 0.001 for all models except for the mentioned `rocks-16` and `maze-10`, as shown in Appendix D in [34]. Our results also confirm that we can extract small FSCs that achieve similar performance. Another benefit compared to SAYNT controllers is the simpler structure of our extracted FSCs. In contrast, SAYNT typically uses a more complex structure to represent policies that combine FSCs with belief-based randomized policies [4]. **Summary: Although abstraction of FSCs incurs a slight degradation of the achieved values, LEXPOP extracts verified FSCs that outperform the FSCs produced by SAYNT.**

RQ1c. Table 1 shows that the SIG extraction provides an advantage over ALERGIA [49] for the `rocks` and `network` models. In these cases, it yields FSCs with higher values that are close to the empirical values of the neural policies, whereas for the rest of the models, the results are mostly similar. However, the main advantage of the

SIG extraction is its flexibility. The size of the final FSC learned by ALERGIA and the extraction time cannot be easily controlled – they mainly depend on the input data. In contrast, self-interpretation can provide the best FSC with a user-defined number of memory nodes. This significantly affects verification time, which increases with the controller’s size. In the worst case, we observed a six-fold increase in verification time when using ALERGIA. **Summary: The proposed SIG extraction produces smaller FSCs than automata learning while preserving the performance.**

RQ2: Robust POMDP Planning

Setting. All experiments run with a 1-hour timeout and 10 different seeds. The number of outer iterations (Line 12 of Algorithm 1) varies between each model and extraction method. Outer iterations of the algorithm might use different sets of POMDPs when using a different extraction algorithm, since the extracted FSCs affect the selection of worst-case POMDPs.

Main results. Table 2 includes the main results for RQ2. It shows the average and the best robust value achieved by FSCs produced by RFPG and by the two variants of LEXPOP. It also reports the average size of the extracted FSCs (see Appendix C.2 in [34]). In the case of RFPG, the size of the FSCs is predetermined by a heuristic and varies across observations, with an upper bound of 4 [22].

RQ2a. The extraction of FSCs from robustified neural policies is generally more challenging, since the hidden states have to handle the partial observability as well as the uncertainty in the environment dynamics. As a result, the hidden state space can be significantly more complex. Figure 3 (b) shows a nontrivial performance gap between the neural policy (dashed lines) evaluated on the set of POMDPs used for training and the robust value of the extracted FSC, evaluated on the whole HM-POMDP, in the moving-obstacle model. Moreover, we observe instability in the extraction using

ALERGIA, which results in a less robust FSC compared to the FSC extracted using the SIG method.

Figure 3 (c) shows the results for avoid-large model where the performance gap increases as the neural policy improves. To investigate this phenomenon, we consider a GRU [16] modification of self-interpretable networks, see Appendix C.2 in [34], where we do not use the quantized layers and rely on continuous values of the GRU feedback. The performance of the modified network (black line) aligns with that of the original neural policy. This shows that the performance gap of the extracted FSCs is caused by the imposed memory limitation, not by the training process. **Summary: For complex HM-POMDP benchmarks, the FSC extraction is challenging due to larger memory requirements. The SIG method provides better and more stable extraction.**

RQ2b. We now compare the performance of LEXPOP with RFPG. Table 2 shows similar trends as in RQ1: LEXPOP achieves worse results than RFPG on small models, but significantly outperforms RFPG on larger models. The statistical significance is confirmed with a Brunner Munzel test, as shown in Appendix D in [34]. The poor performance of RFPG on large models is caused by the expensive parameter updates, which rely on gradient computation through the complete state space. This is exacerbated in the case of the drone model and the larger version of the moving_obstacles model. For these models, we needed to reduce the size of FSCs synthesized by RFPG to 1, since the gradient computation was extremely expensive and could not finish the first iteration. **Summary: LEXPOP scales considerably better than RFPG and thus produces significantly better FSCs on large models.**

RQ2c. Our results confirm the observations from RFPG [22]: For HM-POMDPs describing a large set of POMDPs, a random selection of POMDPs for training robust policies is not sufficient. Figure 3 (d) illustrates this phenomenon on the drone-2-6-1 model. Using randomly selected POMDPs makes the robust training unstable: the red dashed line shows the empirical value of the neural policy evaluated on the set of POMDPs used for training. More importantly, there is a very large performance gap between the neural policy and the robust value of the extracted FSC (the red solid line), regardless of the extraction method. On the other hand, using the worst-case POMDPs (the green lines) leads to more stable learning, a smaller performance gap, and consequently to an FSC having a considerably better robust value. **Summary: The worst-case POMDP selection is crucial for solving HM-POMDPs.**

7 RELATED WORK

POMDP solvers. Existing work includes (monolithic) mixed integer linear program formulation for FSC optimization [3, 41, 75], game-based abstraction techniques [74], or the approximate analysis of belief space [65] implemented e.g. in SARSOP [43] that has been recently extended to unbounded reachability objectives [29]. SAYNT [4] combines an approximate exploration of belief states and symbolic search techniques over FSCs. Carr et al. [11] proposed for the first time to extract FSCs from RNNs to compute verified FSC policies for POMDPs. Yet, this work employs a simplistic, non-vectorized approach that optimizes MDP policies to train the RNN and an extraction method that lacks sufficient flexibility in FSC size. Note that, despite using the RNN policy, this approach does not

employ DRL to optimize the RNNs. Overall, our approach significantly improves scalability by enabling the computation of FSC policies for larger POMDP models.

Robust FSCs for POMDPs. Galeslout et al. [22] introduce HM-POMDPs that describe finitely many variations of a POMDP model, and propose the rFPG algorithm that uses model-based subgradient ascent to iteratively optimize FSCs over worst-case POMDPs. Robust POMDPs [8, 40, 58] represent infinite sets of POMDPs, encapsulating HM-POMDPs, where FSCs have been optimized through convex programming [18, 67]. Galeslout et al. [23] introduce pessimistic iterative planning using RNNs for robust POMDPs. While conceptually similar, they rely on model-based approximations to train the RNNs instead of DRL, hindering scalability and optimality.


Finite-state model extraction from RNNs. The extraction of FSCs originates from extracting finite-state models from RNNs, which has been explored already in the 1990s [57, 78] and has recently gained attention. There are two major approaches: automata learning [51, 53, 71] from the trajectories sampled from the RNN, and quantization of hidden states of RNNs [23, 39]. In line with Koul et al. [39], Carr et al. [10, 11, 12], and Galeslout et al. [23], we also experiment with ways to extract an FSC directly from the RNN. Compared to these, our proposed architecture enhances control over memory size and stabilizes learning by eliminating the straight-through gradient estimate by reparametrization using Gumbel softmax distribution [35]. Though similar in principle to the work of Aichernig et al. [1], our architecture improves upon the prior work by enabling stochastic action and memory state-updates, and achieving faster convergence.

Robust DRL. Various DRL methods address model uncertainty. The most similar approaches to ours are task-robust model-agnostic meta-learning [17], which combines meta-learning with worst-case optimization, and robust adversarial learning [24, 32, 60, 79], which frames learning as a stochastic game between a protagonist and adversary. Since we aim to extract a universal FSC, meta-learning is of limited applicability. In relation to [60], our approach can be seen as using PAYNT as an adversary, which is possible thanks to explicit knowledge of the model family. Other notable methods include model-free worst-case estimation [21], training-time perturbations of states, actions, or observations [46], and reward shaping for worst-case optimization, e.g., via state entropy [50].

8 CONCLUSION

We proposed LEXPOP, a learning-based approach that constructs FSCs for a given POMDP or HM-POMDP. First, LEXPOP trains an RNN-based neural policy using DRL, which allows us to scale to large model sizes. Second, LEXPOP extracts FSCs from the neural policies using automata learning as well as a novel extraction technique. The performance of such FSCs can be formally verified on large POMDPs and for HM-POMDPs, worst-case POMDPs can be effectively computed to iteratively construct robust FSCs. Our experiments demonstrate that LEXPOP scales up to large (HM-)POMDPs with state spaces that were previously infeasible. We plan to improve the robust planning loop by extracting programmatic policies [70]. Furthermore, we plan to scale the loop even further by extending SIG towards model learning.

ACKNOWLEDGMENTS

 This work has been executed under the project VASSAL: “Verification and Analysis for Safety and Security of Applications in Life” funded by the European Union under Horizon Europe WIDERA Coordination and Support Action/Grant Agreement No.101160022. Additionally, this work has been partially funded by the Czech Science Foundation grant GA23-06963S (VESCAA), the Vienna Science and Technology Fund (WWTF) project 10.47379/ICT22-023, and the ERC Starting Grant DEUCE (101077178). Martin Kurečka was supported by the Brno Ph.D. Talent program funded by the Brno City Municipality.

REFERENCES

- [1] Bernhard K. Aichernig, Sandra König, Cristinel Mateis, Andrea Pferscher, and Martin Tappler. 2024. Learning minimal automata with recurrent neural networks. *Softw. Syst. Model.* (2024).
- [2] Christopher Amato, Daniel S. Bernstein, and Shlomo Zilberstein. 2010. Optimizing fixed-size stochastic controllers for POMDPs and decentralized POMDPs. *Auton. Agents Multi Agent Syst.* (2010).
- [3] Christopher Amato, Blai Bonet, and Shlomo Zilberstein. 2010. Finite-State Controllers Based on Mealy Machines for Centralized and Decentralized POMDPs. In *AAAI*.
- [4] Roman Andriushchenko, Alexander Bork, Milan Ceska, Sebastian Junges, Joost-Pieter Katoen, and Filip Macák. 2023. Search and Explore: Symbiotic Policy Synthesis in POMDPs. In *CAV*.
- [5] Roman Andriushchenko, Milan Ceska, Sebastian Junges, Joost-Pieter Katoen, and Simon Stupinsky. 2021. PAYNT: A Tool for Inductive Synthesis of Probabilistic Programs. In *CAV*.
- [6] Blai Bonet and Hector Geffner. 2009. Solving POMDPs: RTDP-Bel vs. Point-based Algorithms. In *IJCAI*.
- [7] Eline M. Bovy, Caleb Probine, Marnix Suilen, Ufuk Topcu, and Nils Jansen. 2025. Multi-Environment POMDPs: Discrete Model Uncertainty Under Partial Observability. In *NeurIPS*.
- [8] Eline M. Bovy, Marnix Suilen, Sebastian Junges, and Nils Jansen. 2024. Imprecise Probabilities Meet Partial Observability: Game Semantics for Robust POMDPs. In *IJCAI*.
- [9] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2018. *JAX: composable transformations of Python+NumPy programs*. <http://github.com/google/jax>
- [10] Steven Carr, Nils Jansen, and Ufuk Topcu. 2020. Verifiable RNN-Based Policies for POMDPs Under Temporal Logic Constraints. In *IJCAI*.
- [11] Steven Carr, Nils Jansen, and Ufuk Topcu. 2021. Task-Aware Verifiable RNN-Based Policies for Partially Observable Markov Decision Processes. *J. Artif. Intell. Res.* (2021).
- [12] Steven Carr, Nils Jansen, Ralf Wimmer, Alexandru Constantin Serban, Bernd Becker, and Ufuk Topcu. 2019. Counterexample-Guided Strategy Improvement for POMDPs Using Recurrent Neural Networks. In *IJCAI*.
- [13] Rafael C. Carrasco and José Oncina. 1994. Learning Stochastic Regular Grammars by Means of a State Merging Method. In *ICGI*, Rafael C. Carrasco and José Oncina (Eds.).
- [14] Milan Ceska, Nils Jansen, Sebastian Junges, and Joost-Pieter Katoen. 2019. Shepherding Hordes of Markov Chains. In *TACAS*.
- [15] Krishnendu Chatterjee, Martin Chmelik, Raghav Gupta, and Ayush Kanodia. 2016. Optimal cost almost-sure reachability in POMDPs. *Artificial Intelligence* (2016).
- [16] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the Properties of Neural Machine Translation: Encoder–Decoder Approaches. In *SSST*.
- [17] Liam Collins, Aryan Mokhtari, and Sanjay Shakkottai. 2020. Task-robust model-agnostic meta-learning. In *NeurIPS*.
- [18] Murat Cubuktepe, Nils Jansen, Sebastian Junges, Ahmadreza Marandi, Marnix Suilen, and Ufuk Topcu. 2021. Robust Finite-State Controllers for Uncertain POMDPs. In *AAAI*.
- [19] Gabriel Dulac-Arnold, Nir Levine, Daniel J. Mankowitz, Jerry Li, Cosmin Paduraru, Sven Gowal, and Todd Hester. 2021. Challenges of real-world reinforcement learning: definitions, benchmarks and analysis. *Mach. Learn.* (2021).
- [20] Xin Fei, Azzedine Boukerche, and Richard Yu. 2010. A POMDP Based K-Coverage Dynamic Scheduling Protocol for Wireless Sensor Networks. In *GLOBECOM*.
- [21] Uri Gadot, Kaixin Wang, Navdeep Kumar, Kfir Yehuda Levy, and Shie Mannor. 2024. Bring Your Own (Non-Robust) Algorithm to Solve Robust MDPs by Estimating The Worst Kernel. In *ICML*.
- [22] Maris F. L. Galesloot, Roman Andriushchenko, Milan Česka, Sebastian Junges, and Nils Jansen. 2025. Robust Finite-Memory Policy Gradients for Hidden-Model POMDPs. In *IJCAI*.
- [23] Maris F. L. Galesloot, Marnix Suilen, Thiago D. Simão, Steven Carr, Matthijs T. J. Spaan, Ufuk Topcu, and Nils Jansen. 2025. Pessimistic Iterative Planning with RNNs for Robust POMDPs. In *ECAI*.
- [24] Adam Gleave, Michael Dennis, Cody Wild, Neel Kant, Sergey Levine, and Stuart Russell. 2021. Adversarial Policies: Attacking Deep Reinforcement Learning. In *ICLR*.
- [25] Sergio Guadarrama, Anoop Korattikara, Oscar Ramirez, Pablo Castro, Ethan Holly, Sam Fishman, Ke Wang, Ekaterina Gonina, Neal Wu, Efi Kokioyopoulou, Luciano Sbaiz, Jamie Smith, Gábor Bartók, Jesse Berent, Chris Harris, Vincent Vanhoucke, and Eugene Brevdo. 2018. TF-Agents: A library for Reinforcement Learning in TensorFlow. <https://github.com/tensorflow/agents>. <https://github.com/tensorflow/agents> [Online; accessed 25-June-2025].
- [26] Eric A. Hansen. 1997. An Improved Policy Iteration Algorithm for Partially Observable MDPs. In *NIPS*.
- [27] Matthew J. Hausknecht and Peter Stone. 2015. Deep Recurrent Q-Learning for Partially Observable MDPs. In *AAAI*.
- [28] Christian Hensel, Sebastian Junges, Joost-Pieter Katoen, Tim Quatmann, and Matthias Volk. 2022. The probabilistic model checker Storm. *Int. J. Softw. Tools Technol. Transf.* (2022).
- [29] Qi Heng Ho, Martin Feather, Federico Rossi, Zachary Sunberg, and Morteza Lahijanian. 2024. Sound Heuristic Search Value Iteration for Undiscounted POMDPs with Reachability Objectives. In *UAI*. PMLR.
- [30] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural Computation* 9, 8 (Nov. 1997), 1735–1780.
- [31] Karel Horák, Branislav Bosanský, and Krishnendu Chatterjee. 2018. Goal-HSVI: Heuristic Search Value Iteration for Goal POMDPs. In *IJCAI*.
- [32] Sandy Huang, Nicolas Papernot, Ian Goodfellow, Yan Duan, and Pieter Abbeel. 2017. Adversarial Attacks on Neural Network Policies. *CoRR* (2017).
- [33] David Hudák, Maris F. L. Galesloot, Martin Tappler, Martin Kurečka, Nils Jansen, and Milan Česka. 2026. Lexpop Synthesis Repository (AAMAS Camera-Ready Branch). https://github.com/DaveHudiny/r1_synthesis/tree/AAMAS-Camera-Ready
- [34] David Hudák, Maris F. L. Galesloot, Martin Tappler, Martin Kurečka, Nils Jansen, and Milan Česka. 2026. Finite-State Controllers for (Hidden-Model) POMDPs using Deep Reinforcement Learning. arXiv:2602.08734 [cs.AI] <https://arxiv.org/abs/2602.08734>
- [35] Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical Reparameterization with Gumbel-Softmax. In *ICLR*.
- [36] Sebastian Junges, Nils Jansen, Ralf Wimmer, Tim Quatmann, Leonore Winterer, Joost-Pieter Katoen, and Bernd Becker. 2018. Finite-state controllers of POMDPs via parameter synthesis. In *UAI*.
- [37] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* (1998).
- [38] Mykel J. Kochenderfer, Tim A. Wheeler, and Kyle H. Wray. 2022. *Algorithms for Decision Making*. MIT Press.
- [39] Anurag Koul, Alan Fern, and Sam Greycanus. 2019. Learning Finite State Representations of Recurrent Policy Networks. In *ICLR*.
- [40] Merlijn Krales, Eline M. Bovy, Maris F. L. Galesloot, Thiago D. Simão, and Nils Jansen. 2025. On Evaluating Policies for Robust POMDPs. In *NeurIPS*.
- [41] Akshat Kumar and Shlomo Zilberstein. 2015. History-based controller design and optimization for partially observable MDPs. In *ICAPS*.
- [42] Hanna Kurniawati. 2022. Partially Observable Markov Decision Processes and Robotics. *Annual Review of Control, Robotics, and Autonomous Systems* (2022).
- [43] Hanna Kurniawati, David Hsu, and Wee Sun Lee. 2008. SARSOP: Efficient Point-Based POMDP Planning by Approximating Optimally Reachable Belief Spaces. In *Robotics: Science and Systems (RSS)*. MIT Press.
- [44] Marta Kwiatkowska, Gethin Norman, and David Parker. 2011. PRISM 4.0: Verification of probabilistic real-time systems. In *CAV*.
- [45] Gaspard Lambrechts, Adrien Bolland, and Damien Ernst. 2022. Recurrent networks, hidden states and beliefs in partially observable environments. *Trans. Mach. Learn. Res.* (2022).
- [46] Zuxin Liu and Zijian Guo. 2023. On the Robustness of Safe Reinforcement Learning under Observational Perturbations. In *ICLR*.
- [47] Omid Madani, Steve Hanks, and Anne Condon. 1999. On the Undecidability of Probabilistic Planning and Infinite-Horizon Partially Observable Markov Decision Problems. In *AAAI/IAAI*.
- [48] Omid Madani, Steve Hanks, and Anne Condon. 2000. On the Undecidability of Probabilistic Planning and Infinite-Horizon Partially Observable Markov Decision Problems. *Proceedings of the National Conference on Artificial Intelligence* (2000).
- [49] Hua Mao, Yingke Chen, Manfred Jaeger, Thomas D. Nielsen, Kim G. Larsen, and Brian Nielsen. 2016. Learning deterministic probabilistic automata from a model checking perspective. *Mach. Learn.* (2016).
- [50] Pierre-François Massiani, Alexander von Rohr, Lukas Haverbeck, and Sebastian Trimpe. 2024. Viability of Future Actions: Robust Reinforcement Learning via Entropy Regularization. In *EWRL*.

- [51] Franz Mayr and Sergio Yovine. 2018. Regular Inference on Artificial Neural Networks. In *CD-MAKE*.
- [52] Nicolas Meuleau, Kee-Eung Kim, Leslie Pack Kaelbling, and Anthony R. Cassandra. 1999. Solving POMDPs by searching the space of finite policies. In *UAI*.
- [53] Edi Muskardın, Bernhard K. Aichernig, Ingo Pill, and Martin Tappler. 2022. Learning Finite State Models from Recurrent Neural Networks. In *IFM*.
- [54] Edi Muskardın, Martin Tappler, Bernhard K. Aichernig, and Ingo Pill. 2024. Active model learning of stochastic reactive systems (extended version). *Softw. Syst. Model.* (2024).
- [55] Edi Muskardın, Bernhard K. Aichernig, Ingo Pill, Andrea Pferscher, and Martin Tappler. 2022. AALpy: an active automata learning library. *Innovations in Systems and Software Engineering* (2022).
- [56] Tianwei Ni, Benjamin Eysenbach, and Ruslan Salakhutdinov. 2022. Recurrent Model-Free RL Can Be a Strong Baseline for Many POMDPs. In *ICML*.
- [57] Christian W. Omlin and C. Lee Giles. 1996. Extraction of rules from discrete-time recurrent neural networks. *Neural Networks* 9, 1 (1996), 41–52.
- [58] Takayuki Osogami. 2015. Robust partially observable Markov decision process. In *ICML (JMLR Workshop and Conference Proceedings, Vol. 37)*. JMLR.org, 106–115.
- [59] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. In *NeurIPS*.
- [60] Lorrel Pinto, James Davidson, Rahul Sukthankar, and Abhinav Gupta. 2017. Robust adversarial reinforcement learning. In *ICML*.
- [61] Pascal Poupart and Craig Boutilier. 2003. Bounded Finite State Controllers. In *Advances in Neural Information Processing Systems (NIPS)*. MIT Press, 823–830.
- [62] Martin L. Puterman. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Statistics.
- [63] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal Policy Optimization Algorithms. *CoRR* abs/1707.06347 (2017).
- [64] Richard D Smallwood and Edward J Sondik. 1973. The optimal control of partially observable Markov processes over a finite horizon. *Oper. Res.* (1973).
- [65] Trey Smith and Reid Simmons. 2004. Heuristic search value iteration for POMDPs. In *UAI*.
- [66] Matthijs T J Spaan. 2012. Partially Observable Markov Decision Processes. In *Reinforcement Learning: State-of-the-Art*.
- [67] Marnix Suilen, Nils Jansen, Murat Cubuktepe, and Ufuk Topcu. 2020. Robust Policy Synthesis for Uncertain POMDPs via Convex Optimization. In *IJCAI*.
- [68] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement Learning: An Introduction*.
- [69] Petr Tomášek, Karel Horak, and Branislav Bošanský. 2024. Iterative algorithms for solving one-sided partially observable stochastic shortest path games. *International Journal of Approximate Reasoning* (2024).
- [70] Abhinav Verma, Vijayaraghavan Murali, Rishabh Singh, Pushmeet Kohli, and Swarat Chaudhuri. 2018. Programmatically Interpretable Reinforcement Learning. In *Proceedings of the 35th International Conference on Machine Learning*. PMLR.
- [71] Gail Weiss, Yoav Goldberg, and Eran Yahav. 2018. Extracting Automata from Recurrent Neural Networks Using Queries and Counterexamples. In *ICML*.
- [72] Paul J. Werbos. 1990. Backpropagation through time: what it does and how to do it. *Proc. IEEE* (1990).
- [73] Daan Wierstra, Alexander Förster, Jan Peters, and Jürgen Schmidhuber. 2007. Solving deep memory POMDPs with recurrent policy gradients. In *ICANN*.
- [74] Leonore Winterer, Sebastian Junges, Ralf Wimmer, Nils Jansen, Ufuk Topcu, Joost-Pieter Katoen, and Bernd Becker. 2021. Strategy Synthesis for POMDPs in Robot Planning via Game-Based Abstractions. *IEEE Trans. Autom. Control.* (2021).
- [75] Leonore Winterer, Ralf Wimmer, Bernd Becker, and Nils Jansen. 2024. Strong Simple Policies for POMDPs. *Int. J. Softw. Tools Technol. Transf.* (2024).
- [76] Peter R. Wurman, Samuel Barrett, Kenta Kawamoto, James MacGlashan, Kaushik Subramanian, Thomas J. Walsh, Roberto Capobianco, Alisa Devlic, Franziska Eckert, Florian Fuchs, Leilani Gilpin, Piyush Khandelwal, Varun Raj Kompella, HaoChih Lin, Patrick MacAlpine, Declan Oller, Takuma Seno, Craig Sherstan, Michael D. Thomure, Houmeh Aghabozorgi, Leon Barrett, Rory Douglas, Dion Whitehead, Peter Dürr, Peter Stone, Michael Spranger, and Hiroaki Kitano. 2022. Outracing champion Gran Turismo drivers with deep reinforcement learning. *Nature* (2022).
- [77] Yang You, Vincent Thomas, Alex Schutz, Robert Skilton, Nick Hawes, and Olivier Buffet. 2025. Partially observable Monte-Carlo graph search. In *ICAPS*.
- [78] Zheng Zeng, Rodney M. Goodman, and Padhraic Smyth. 1993. Learning Finite State Machines With Self-Clustering Recurrent Networks. *Neural Computing* (1993).
- [79] Huan Zhang, Hongge Chen, Duane S. Boning, and Cho-Jui Hsieh. 2021. Robust Reinforcement Learning on State Observations with Learned Optimal Adversary. In *ICLR*.
- [80] Wenqian Zhang and Haiyan Wang. 2022. Diagnostic Policies Optimization for Chronic Diseases Based on POMDP Model. *Healthcare* (2022).
- [81] Wei Zheng, Taeho Jung, and Hai Lin. 2022. The Stackelberg equilibrium for one-sided zero-sum partially observable stochastic games. *Automatica* (2022).