

# Controller Synthesis from Deep Reinforcement Learning Policies

Florent Delgrange<sup>1,2</sup>, Guy Avni<sup>3</sup>, Anna Lukina<sup>4</sup>, Christian Schilling<sup>5</sup>, Ann Nowé<sup>1</sup>, Guillermo A. Pérez<sup>2,6</sup>

<sup>1</sup>Vrije Universiteit Brussel, Belgium <sup>2</sup>University of Antwerp, Belgium <sup>3</sup>University of Haifa, Israel <sup>4</sup>Delft University of Technology, The Netherlands <sup>5</sup>Aalborg University, Denmark <sup>6</sup>Flanders Make, Belgium

## Abstract

We propose a novel framework to controller design in environments with a two-level structure: a known high-level graph (“map”) in which each vertex is populated by a Markov decision process, called a “room”. The framework “separates concerns” by using different design techniques for low- and high-level tasks. We apply reactive synthesis for high-level tasks: given a specification as a logical formula over the high-level graph and a collection of low-level policies given on “concise” *latent* structures, we construct a “planner” that selects which low-level policy to apply in each room. We develop a reinforcement learning procedure to train low-level policies on latent structures, which unlike previous approaches, circumvents a model distillation step. It pairs the policy with probably approximately correct guarantees on its performance and abstraction quality, which are lifted to guarantees on the high-level task. These formal guarantees are the main advantage of the framework. Other advantages include scalability (rooms are large and their dynamics is unknown) and reusability of low-level policies. We demonstrate feasibility in challenging case studies involving agent navigation in environments with moving obstacles and visual inputs.

## 1 Introduction

We consider the fundamental problem of constructing control *policies* for environments modeled as *Markov decision processes* (MDPs) with formal guarantees. We deal with long-horizon tasks in environments with some prior knowledge of the environment structure: the input to our method is a (high-level) *map* given as a *graph*, where each vertex is populated by an (possibly large) MDP with unknown dynamics called a *room*, and the long-horizon task is given on the map. We argue that such settings arise naturally. Our running example is a robot for delivering a package in a warehouse with moving obstacles (e.g., forklifts, workers, or other robots); while it is infeasible to model the low-level interactions of the agent with its immediate surroundings, modeling the high-level map of the rooms in the warehouse requires minimal engineering effort. We list other examples of two-level domains with prior knowledge of the high-level architecture and in which our method is relevant: (i) *routing* (Junges and Spaan 2022): the network topology, e.g., connection between routers, is often known but

modeling low-level routing decisions is intricate; (ii) *skill graphs* (Bagaria et al. 2021) of agents, e.g., “grab a key” and “open a door”, and their dependencies are naturally modeled as a graph; (iii) *software systems* (Ryzhyk et al. 2009), in particular *probabilistic programs* (Junges and Spaan 2022): each vertex represents a software component (an MDP in a probabilistic program) and edges capture dependencies or interactions.

Our framework “separates concerns” by using different design techniques for low- and high-level tasks with complementary benefits and drawbacks. For high-level tasks, we apply *reactive synthesis* (Pnueli and Rosner 1989), which constructs an optimal policy *based on a model of the environment and objectives specified as a logical formula*. Its main advantage is a *guarantee that the policy satisfies the specification*. Also, synthesis *allows users an intuitive and natural specification language*. The reliance on an explicit model of the environment hinders scalability as well as reasoning about domains with partially-known dynamics. Hence, we solve low-level tasks via *reinforcement learning* (RL (Sutton and Barto 1998)). In particular, we may use *deep RL* (DRL (Mnih et al. 2013)), which is successful in domains of *high-dimensional feature spaces with unknown dynamics*. However, RL generally lacks formal guarantees and struggles with long-term objectives, where one needs to deal with the notorious problem of sparse rewards (Ladosz et al. 2022) by guiding the agent (Liu et al. 2022), which in turn poses an engineering effort.

**Framework (Fig. 1)** We output a *two-level controller* for an agent, consisting of a collection of low-level policies  $\Pi$  and a high-level *planner*  $\tau$ . When the agent enters a room corresponding to a vertex  $v$  of the map, the planner chooses an outgoing edge  $e$  and deploys the associated policy  $\pi_{v,e} \in \Pi$ . The agent follows  $\pi_{v,e}$  until it exits the room. For example,  $e$  can model a door between two rooms. Note that the agent may exit from direction  $e' \neq e$ . It is thus key to have an estimate on the success probability of  $\pi_{v,e}$  when designing  $\tau$ .

We obtain low-level policies by developing a novel RL procedure that is run locally in each room  $v$  and outputs *latent policies*  $\pi_{v,e}$ , for each direction  $e$ . These policies are represented on a concise model of the room (Fig. 1(b)). We stress that we only assume simulation access to the rooms; *the latent policies are learned*. Importantly, each latent policy is paired with probably approximately correct (PAC)

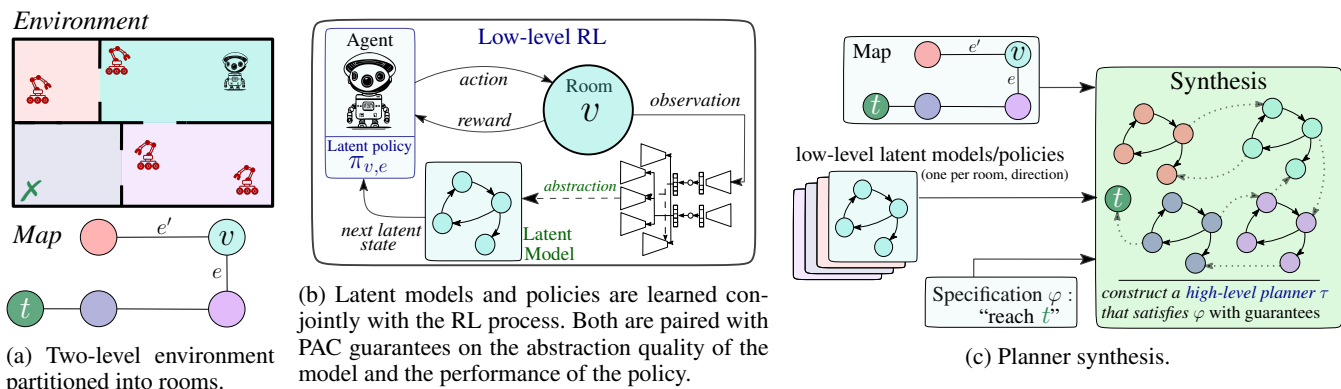


Figure 1: **(a)** Environment in which the agent (top-right) needs to reach the target (green, bottom-left) while avoiding moving adversaries (in red). The target appears in the map as a dedicated vertex  $t$ . **(b)** The agent is trained to exit *each* room, in *every possible* direction. Training is performed in *parallel simulations*. An abstraction of the environment is learned via neural networks, yielding a latent model for each room. Simultaneously, a policy is learned via RL on the learned latent representation, which guarantees the agent’s low-level behavior conformity through PAC bounds. *More details in Sect. 4.* **(c)** Given a high-level description of the environment, a collection of latent models and policies for each room, and the specifications, synthesis outputs a high-level planner guaranteed to satisfy the specifications. The challenge resides in the way the low-level components are merged to apply synthesis while maintaining their guarantees. *More details in Sect. 5.*

guarantees on its performance.

Finally, given a map, a collection of policies  $\Pi$ , and a high-level specification  $\varphi$  given as a logical formula over the map, we design an algorithm to find a planner  $\tau$  that optimizes for  $\varphi$  while lifting the guarantees on the policies in  $\Pi$  to  $\tau$  (Fig. 1(c)).

**Advantages.** We point to advantages of the framework. First and foremost, it *provides guarantees on the operation of the controller*. A key design objective is to ease the engineering burden: reward engineering is only done locally (for each room), and the high-level map and tasks are given in an intuitive specification language. Second, our framework enables *reusability*: a policy  $\pi_{v,e}$ , including its guarantees, is reusable across similar rooms  $v'$  and when the high-level task or structure changes. Finally, our framework offers a remedy for the notorious challenge of sparse rewards in RL.

**Case study.** We complement our theoretical results with two case studies which illustrate the feasibility of the approach. We consider a domain in which an agent needs to reach a distant location while avoiding mobile adversarial obstacles with stochastic dynamics. The first case study is a grid world, while the second case study is a vision-based Doom environment (Kempka et al. 2016). We show that DQN struggles to find a policy in our domain, even with reward shaping. In the rooms, we demonstrate our novel procedure for training concise latent policies directly. We synthesize a planner based on the latent policies and show the following results. First, our two-level controller achieves high success probability, demonstrating that our approach overcomes the challenge of sparse rewards. Second, the values predicted in the latent model are close to those observed, demonstrating the quality of our automatically constructed model.

**Contributions.** We outline our key theoretical contributions.

- (i) *Learning guarantees for low-level policies.* We tie between the *values* (the probability that the low-level ob-

jective is satisfied) of the latent model and that of the environment via a loss function (Thm. 1) and demonstrate that PAC bounds can be computed for these value differences (Thm. 2).

- (ii) *Guarantees on the synthesized controller.* We prove memory bounds on the size of an optimal high-level planner (Thm. 3). Moreover, we show that an optimal planner can be obtained by solving an MDP whose size is proportional to the size of the map, i.e., disregarding the size of the rooms (Thm. 4).
- (iii) *Unified learning and synthesis guarantees.* We show that the learning guarantees for the low-level policies can be lifted to the two-level controller. Specifically, minimizing the loss function to learn an abstraction of each room independently (and in parallel) guarantees that the values obtained under the two-level controller in the abstraction closely match those obtained in the true two-level environment (Thm. 5).

**Related work.** Hierarchical RL (HRL) (Pateria et al. 2022) (see also, the *option* framework (Sutton et al. 1999)) is an approach that outputs two-level controllers. Our approach is very different despite similarity in terms of outputs and motivations (e.g., both enable reusability and modularity). The most significant difference is that *our framework provides guarantees*, which *HRL generally lacks*. In our framework, high-level planners are synthesized based on prior knowledge of the environment (the map) and only after the low-level policies are learned. In HRL, both low-level policies and two-level controllers can be learned concurrently and with no prior knowledge. Another difference is that in HRL, the low-level objectives generally need to be learned, whereas in our approach they are known. We argue that the “separation of concerns” in our framework eases the engineering burden while HRL notoriously requires significant engineering efforts. Finally, unlike option-inspired ap-

proaches, where the integration of high- and low-level components results in a “semi-”Markovian process, our framework ensures that a small amount of memory for the high-level planner is sufficient to enable the agent to operate within a *fully* Markovian process. This facilitates the design of high-level planning and synthesis solutions.

*Distillation* (Hinton et al. 2015) is an established approach: a neural network (NN) is trained then *distilled* into a concise *latent* model. Verification of NN controllers is challenging, e.g., (Amir et al. 2021). Verification based distillation is a popular approach in which verification is applied to a latent policy, e.g., Ernst et al. (2005); Delgrange et al. (2022); Bastani et al. (2018); Bacci et al. (2021); Carr et al. (2021). In contrast, *we study controller-synthesis based on latent policies*. To our knowledge, only Alamdari et al. (2020) develops a synthesis based on distillation approach, but with no guarantees. In addition, *we develop a novel training procedure that trains a latent policy directly and circumvents the need for model distillation*. We stress that the abstraction is learned unlike Roderick et al. (2018); Jothimurugan et al. (2021).

*CLAPS* (Žikelić et al. 2023) is an approach that outputs a two-level controller with correctness guarantees, which is very different from ours. Low-level policies are paired with a super-martingale on the environment states that gives rise to reach-avoid guarantees. In contrast, our policies are trained on a learned latent model, which we accompany with PAC guarantees on the quality of the abstraction. They further assume prior knowledge of the transitions whereas we only assume simulation access, their policy is limited to be stationary and deterministic whereas our policies are general, and their high-level structure arises from the logical specification whereas ours arises from the environment’s structure.

Formal reasoning in RL is a timely research agenda. Safety objectives in RL are intractable (Alur et al. 2022). *Shielding* (Alshiekh et al. 2018; Könighofer et al. 2022) circumvents the difficulty of ensuring safety during training by monitoring a policy at runtime and blocking unsafe actions. Shielding has been applied to low-level policies in a hierarchical controller (Xiong et al. 2022). The limitation of this approach is that interference with the trained policy might break its guarantees. LTL objectives add intractability (Yang et al. 2021) to the already complex hierarchical scenarios in RL (Kulkarni et al. 2016) and only allow for PAC guarantees if the MDP structure is known (Fu and Topcu 2014). Reactive synthesis is applied by Nayak et al. (2023) to obtain low-level controllers, but scalability is a shortcoming of synthesis. Approaches encouraging but not ensuring safety use constrained policy optimization (Achiam et al. 2017), safe padding in small steps (Hasanbeig et al. 2020), time-bounded safety (Giacobbe et al. 2021), safety-augmented MDPs (Sootla et al. 2022), differentiable probabilistic logic (Yang et al. 2023), or distribution sampling (Badings et al. 2023).

## 2 Preliminaries

**Markov Decision Processes (MDPs).** Let  $\Delta(\mathcal{X})$  denote the set of distributions on  $\mathcal{X}$ . An *MDP* is a tuple  $\mathcal{M} =$

$\langle \mathcal{S}, \mathcal{A}, \mathbf{P}, \mathbf{I} \rangle$  with states  $\mathcal{S}$ , actions  $\mathcal{A}$ , transition function  $\mathbf{P}: \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ , and initial distribution  $\mathbf{I} \in \Delta(\mathcal{S})$ . An agent interacts with  $\mathcal{M}$  as follows. At each step, the agent is in some state  $s \in \mathcal{S}$ . It performs an action  $a \in \mathcal{A}$  and subsequently goes to the next state according to the transition function:  $s' \sim \mathbf{P}(\cdot | s, a)$ . A *policy*  $\pi: \mathcal{S} \rightarrow \Delta(\mathcal{A})$  prescribes which action to choose at each step and gives rise to a distribution over *paths* of  $\mathcal{M}$ , denoted by  $\Pr_{\pi}^{\mathcal{M}}$ . The probability of finite paths is defined inductively. Trivial paths  $s \in \mathcal{S}$  have probability  $\Pr_{\pi}^{\mathcal{M}}(s) = \mathbf{I}(s)$ . Paths  $\rho = s_0, s_1, \dots, s_n$  have probability  $\Pr_{\pi}^{\mathcal{M}}(s_0, s_1, \dots, s_{n-1}) \cdot \mathbb{E}_{a \sim \pi(\cdot | s_{n-1})} \mathbf{P}(s_n | s_{n-1}, a)$ .


**Limiting behaviors in MDPs.** The *transient measure*  $\mu_{\pi}^n(s' | s) = \mathbb{P}_{\rho \sim \Pr_{\pi}^{\mathcal{M}}}[\rho \in \{s_0, \dots, s_n | s_n = s'\} | s_0 = s]$  gives the probability of visiting each state  $s'$  after exactly  $n$  steps starting from  $s \in \mathcal{S}$ . Under policy  $\pi$ ,  $C \subseteq \mathcal{S}$  is a *bottom strongly connected component* (BSCC) of  $\mathcal{M}$  if (i)  $C$  is a maximal subset satisfying  $\mu_{\pi}^n(s' | s) > 0$  for any  $s, s' \in C$  and some  $n \geq 0$ , and (ii)  $\mathbb{E}_{a \sim \pi(\cdot | s)} \mathbf{P}(C | s, a) = 1$  for all  $s \in \mathcal{S}$ . MDP  $\mathcal{M}$  is *ergodic* if, under any stationary policy  $\pi$ , the set of reachable states  $Reach(\mathcal{M}, \pi) = \{s \in \mathcal{S} | \exists n \geq 0, \mathbb{E}_{s_0 \sim \mathbf{I}} \mu_{\pi}^n(s | s_0) > 0\}$  consist of a unique aperiodic BSCC. Then, for  $s \in \mathcal{S}$ , the *stationary distribution* of  $\mathcal{M}$  under  $\pi$  is given by  $\xi_{\pi} = \lim_{n \rightarrow \infty} \mu_{\pi}^n(\cdot | s)$ .

**Objectives and values.** A qualitative *objective* is a set of infinite paths  $\mathbb{O} \subseteq \mathcal{S}^{\omega}$ . For  $B, T \subseteq \mathcal{S}$ , we consider *reach-avoid objectives*  $\mathbb{O}(T, B) = \{s_0, s_1, \dots | \exists i. s_i \in T \text{ and } \forall j \leq i, s_j \notin B\}$  (or just  $\mathbb{O}$  if clear from context) where the goal is to reach a “target” in  $T$  while avoiding the “bad” states  $B$ . Henceforth, fix a *discount factor*  $\gamma \in (0, 1)$ . In this work, we consider *discounted value functions* (see, e.g., (de Alfaro et al. 2003)). The *value* of any state  $s \in \mathcal{S}$  for policy  $\pi$  w.r.t. objective  $\mathbb{O}$  is denoted by  $V^{\pi}(s, \mathbb{O})$  and corresponds to the probability of satisfying  $\mathbb{O}$  from state  $s$  as  $\gamma$  goes to one, i.e.,  $\lim_{\gamma \rightarrow 1} V^{\pi}(s, \mathbb{O}) = \mathbb{P}_{\rho \sim \Pr_{\pi}^{\mathcal{M}}}[\rho \in \mathbb{O} | s_0 = s]$ . In particular, for the reach-avoid objective  $\mathbb{O}(T, B)$ ,  $V^{\pi}(s, \mathbb{O})$  corresponds to the discounted probability of visiting  $T$  for the first time while avoiding  $B$ , i.e.,  $V^{\pi}(s, \mathbb{O}) = \mathbb{E}_{\rho \sim \Pr_{\pi}^{\mathcal{M}}}[\sup_{i \geq 0} \gamma^i \cdot \mathbb{1}\{s_i \in T \wedge \forall j \leq i, s_j \notin B\} | s_0 = s]$ , where  $s_i, s_j$  are respectively the  $i^{\text{th}}, j^{\text{th}}$  state of  $\rho$ . We are particularly interested in the values obtained from the beginning of the execution, written  $V_{\mathbf{I}}^{\pi}(\mathbb{O}) = \mathbb{E}_{s_0 \sim \mathbf{I}}[V^{\pi}(s_0, \mathbb{O})]$ . We may sometimes omit  $\mathbb{O}$  and simply write  $V^{\pi}$  and  $V_{\mathbf{I}}^{\pi}$ .


**Reinforcement learning** obtains a policy in a model-free way. Executing action  $a_i$  in state  $s_i$  and transitioning to  $s_{i+1}$  incurs a reward  $r_i = rew(s_i, a_i, s_{i+1})$ , computed via a *reward function*  $rew: \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ . An RL agent’s goal is to learn a policy  $\pi^*$  maximizing the return  $\mathbb{E}_{\rho \sim \Pr_{\pi^*}^{\mathcal{M}}}[\sum_{i \geq 0} \gamma^i r_i]$ . The agent is trained by interacting with the environment in episodic simulations, each ending in one of three ways: success, failure, or an eventual reset.

## 3 Problem Formulation

In this section, we formally model a two-level environment and state the problem of two-level controller synthesis. The



(a) A two-level model of a simple grid world environment.




(b) A two-level model for which an optimal planner requires memory, here flattened in 2D.

Figure 2: (a) Top: The high-level graph  $\mathcal{G}$  with two rooms  $R_0 = \ell(v_0)$  and  $R_1 = \ell(u)$ . Middle: Part of the explicit MDP for the bottom layer; e.g., the MDP  $R_0$  contains 16 states. Traversing the edge  $\langle \langle s_2, v_0 \rangle, \langle s_0, u \rangle \rangle$  corresponds to exiting  $R_0$  and entering  $R_1$  from direction  $d_1 = \langle v_0, u \rangle$ . The goal of  $\odot$  is to reach  $u'$  by exiting the room  $R_1$  from direction  $d_2 = \langle u, u' \rangle$  while avoiding the moving adversaries  $\ominus$ . For  $i \in \{0, 1\}$ , the entrance function  $\mathcal{I}_{R_i}$  models the distribution from which the initial location of  $\ominus$  in  $R_i$  is drawn. (b) A room with four policies for a planner to choose from; e.g.,  $\pi_{\rightarrow}(\cdot | s_1)$  leads to  $B_{\ell(u)}$  and  $\pi_{\uparrow}(\cdot | s_1)$  leads to  $s_3$ . Note that while these are deterministic policies, in general the policies in rooms are probabilistic.

environment MDP is given by a high-level *map*: an undirected graph whose vertices are associated with “low-level” MDPs called *rooms* (Fig. 2(a)). A two-level controller consists of two components and operates as follows. In each room, we assume access to a set of *low-level policies*, each optimizing a local (room) reach-avoid objective (Fig. 2(b)). When transitioning to a new room, a high-level *planner* selects the next low-level policy.

**Two-level model.** A room  $R = \langle \mathcal{S}_R, \mathcal{A}_R, \mathbf{P}_R, D_R, \mathcal{I}_R, \mathcal{O}_R \rangle$  consists of  $\mathcal{S}_R, \mathcal{A}_R, \mathbf{P}_R$  as in an MDP, a set of *directions*  $D_R$ , an *entrance function*  $\mathcal{I}_R: D_R \rightarrow \Delta(\mathcal{S}_R)$  taking a direction from which the room is entered and producing an initial distribution over states, and an *exit function*  $\mathcal{O}_R: D_R \rightarrow 2^{\mathcal{S}_R}$  returning a set of *exit states* from the room in a given direction  $d \in D_R$ . States are assigned to at most one exit, i.e., if  $s \in \mathcal{O}_R(d)$  and  $s \in \mathcal{O}_R(d')$ , then  $d' = d$ .

*Example 1 (Room).* Consider the grid world below as a room  $R$  populated by an adversary  $\ominus$ . One can encode



the position of  $\ominus$  in  $\mathcal{S}_R$  and its behaviors through  $\mathbf{P}_R$ . This can be achieved by, e.g., considering states of the form  $s = \langle (x_1, y_1), (x_2, y_2) \rangle \in \mathcal{S}_R$  where  $(x_1, y_1)$  is the position

of  $\odot$  and  $(x_2, y_2)$  the one of  $\ominus$  in the grid. Note that the position of  $\ominus$  depends on the direction from which  $\odot$  enters  $R$ . Here,  $\odot$  enters from the left in direction  $\rightarrow$  to the states of  $R$  distributed according to the entrance function  $\mathcal{I}_R(\cdot | d = \rightarrow)$  (the tiling patterns highlight its support). While  $\odot$  enters (deterministically) in the leftmost cell (yellow tiling),  $\mathcal{I}_R$  allows to (probabilistically) model the possible positions of  $\ominus$  when entering the room (red tiling) from direction  $d = \rightarrow$ . When reaching the green area, depicting states from  $\mathcal{O}_R(\rightarrow)$ ,  $\odot$  exits  $R$  by the right direction  $\rightarrow$ .

A *map* is a graph  $\mathcal{G} = \langle \mathcal{V}, E \rangle$  with vertices  $\mathcal{V}$  and undirected edges  $E \subseteq \mathcal{V} \times \mathcal{V}$ . As  $\mathcal{G}$  is undirected, for any  $u, v \in \mathcal{V}$ ,  $\langle u, v \rangle \in E$  if and only if  $\langle v, u \rangle \in E$ . The *neighbors* of

$v \in \mathcal{V}$  are  $N(v) = \{u \in \mathcal{V} | \langle u, v \rangle \in E\}$  and the outgoing edges from  $v$  are  $out(v) = \{e = \langle v, u \rangle \in E\}$ . A *two-level model*  $\mathcal{H} = \langle \mathcal{G}, \mathcal{R}, \ell, v_0, \langle d_0, d_1 \rangle \rangle$  consists of a map  $\mathcal{G} = \langle \mathcal{V}, E \rangle$ , a set of rooms  $\mathcal{R}$ , a labeling  $\ell: \mathcal{V} \rightarrow \mathcal{R}$  of each vertex  $v \in \mathcal{V}$  with a room  $\ell(v)$  and directions  $D_{\ell(v)} = out(v)$ , an initial room  $v_0 \in \mathcal{V}$ , and directions  $d_0, d_1 \in out(v_0)$  in which  $v_0$  is respectively entered and must be exited.

Fix a two-level model  $\mathcal{H} = \langle \mathcal{G}, \mathcal{R}, \ell, v_0, \langle d_0, d_1 \rangle \rangle$ . Intuitively, the *explicit MDP*  $\mathcal{M}$  corresponding to  $\mathcal{H}$  is obtained by “stitching” MDPs  $R \in \mathcal{R}$  corresponding to neighboring rooms (Fig. 2(a)). Formally,  $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathbf{P}, \mathbf{I} \rangle$ , where  $\mathcal{S} = \{ \langle s, v \rangle : s \in \mathcal{S}_{\ell(v)}, v \in \mathcal{V} \}$ ,  $\mathcal{A} = \bigcup_{R \in \mathcal{R}} \mathcal{A}_R \cup \{a_{exit}\}$ . The initial distribution  $\mathbf{I}$  simulates starting in room  $\ell(v_0)$  from direction  $d_0$ ; thus, for each  $s \in \mathcal{S}_{\ell(v_0)}$ ,  $\mathbf{I}(\langle s, v_0 \rangle) = \mathcal{I}_{\ell(v_0)}(s | d_0)$ . The transitions  $\mathbf{P}$  coincide with  $\mathbf{P}_R$  for non-exit states. Let  $d = \langle v, u \rangle \in E$  with  $v \in N(u)$ ;  $\mathcal{O}_R(d)$  are the exit states in room  $R$  associated with  $v$  in direction  $d$ , and  $\mathcal{I}_{\ell(u)}(\cdot | d)$  is the entrance distribution in  $R$  associated with  $u$  in direction  $d$ . The successor state of  $s \in \mathcal{O}_R(d)$  follows  $\mathcal{I}_{\ell(u)}(\cdot | d)$  when  $a_{exit}$  is chosen. Each path  $\rho$  in  $\mathcal{M}$  corresponds to a unique path  $(\rho)$  in  $\mathcal{G}$ .

**High-level reach and low-level reach-avoid objectives.** The high-level reachability objective we consider is “ $\diamond T$ ,” where  $T \subseteq \mathcal{V}$  is a subset of vertices in the graph of  $\mathcal{H}$ . Here,  $\diamond T$  is a temporal logic notation meaning “eventually visit the set  $T$ .” Formally, a path  $\rho$  in  $\mathcal{M}$  satisfies  $\diamond T$  iff  $path(\rho)$  visits a vertex  $v$  in  $T$ . The low-level safety objective is defined over states of the rooms in  $\mathcal{R}$ . For each room  $R$ , let  $B_R \subseteq \mathcal{S}_R$  be a set of “bad” states. For room  $R$  and direction  $d \in D_R$ , the reach-avoid objective  $\mathbb{O}_R^d \in \mathcal{S}_R^*$  is  $\{s_0, \dots, s_n | s_n \in \mathcal{O}_R(d) \text{ and } s_i \notin B_R \text{ for all } i \leq n\}$ , i.e., exit  $R$  via  $d$  avoiding  $B_R$ .

**High-level control.** We define a *high-level planner*  $\tau: \mathcal{V}^* \rightarrow E$  and a set of low-level policies  $\Pi$  such that, for each room  $R \in \mathcal{R}$  and a direction  $d \in D_R$ ,  $\Pi$  contains a policy  $\pi_{R,d}$  for the objective  $\mathbb{O}_R^d$ . The pair  $\pi = \langle \tau, \Pi \rangle$  is a *two-level controller* for  $\mathcal{H}$ , defined

inductively as follows. Consider the initial vertex  $v_0 \in \mathcal{V}$  (Fig. 2(a)). Let  $d_0 = \tau(\epsilon) \in \text{out}(v_0)$  ( $\epsilon$  being the empty sequence). Control in  $\ell(v_0)$  follows  $\pi_{\ell(v_0), d_1}$ . Let  $\rho$  be a path in  $\mathcal{H}$  ending in  $s \in \mathcal{S}_R$ , for some room  $R = \ell(v)$ . If  $s$  is not an exit state of  $R$ , then control follows a policy  $\pi_{R, d}$  with  $d = \langle v, u \rangle$  and  $u \in N(v)$ . If  $s$  is an exit state in direction  $d$  and  $\text{path}(\rho)$  ends in  $v$ , i.e.,  $s \in \mathcal{O}_R(d)$ , then  $a_{\text{exit}}$  is taken in  $s$  and the next state is an initial state in  $R' = \ell(u)$  drawn from  $\mathcal{I}_{R'}(d)$ . The planner chooses a direction  $d' = \tau(\text{path}(\rho) \cdot u) \in \text{out}(u)$  to exit  $R'$ . Control of  $R'$  proceeds with the low-level policy  $\pi_{R', d'}$ . Note that  $\pi$  is a policy in the explicit MDP  $\mathcal{M}$ .

**Problem 1.** *Given a two-level model  $\mathcal{H} = \langle \mathcal{G}, \mathcal{R}, \ell, v_0, \langle d_0, d_1 \rangle \rangle$ , discount factor  $\gamma \in (0, 1)$ , high-level objective  $\diamond T$ , and low-level objectives  $\{\odot_R^d \mid R \in \mathcal{R}, d \in D_R\}$ , construct a two-level controller  $\pi = \langle \tau, \Pi \rangle$  maximizing the probability of satisfying the objectives.*

## 4 Obtaining Low-Level RL policies

There are fundamental challenges in reasoning about policies obtained via RL—especially those obtained via DRL, which are typically represented by large NNs. We develop a novel, unified approach which outputs a latent model together with a concise policy. The idea is to learn a *tractable* latent model for each room, where the values of the low-level objectives can be explicitly computed. Each latent model is accompanied by *probably approximately correct* (PAC) guarantees on their abstraction quality. We first focus on those guarantees. In the next section, we will then focus on how to synthesize a planner (with guarantees) based on these learned models and policies. Proofs of our claims are given in Appendix C.

### 4.1 Quantifying the quality of the abstraction

Throughout this section, we fix an MDP environment  $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathbf{P}, \mathbf{I} \rangle$ . A *latent model* abstracts a concrete MDP and is itself an MDP  $\overline{\mathcal{M}} = \langle \overline{\mathcal{S}}, \mathcal{A}, \overline{\mathbf{P}}, \overline{\mathbf{I}} \rangle$  whose state space is linked to  $\mathcal{M}$  via a *state-embedding function*  $\phi: \mathcal{S} \rightarrow \overline{\mathcal{S}}$ . We focus on latent MDPs with a finite state space, where values can be exactly computed.

Let  $\overline{\pi}$  be a policy in  $\overline{\mathcal{M}}$ , called a *latent policy*. The key feature is that  $\phi$  allows to control  $\mathcal{M}$  using  $\overline{\pi}$ : for each state  $s \in \mathcal{S}$ , let  $\overline{\pi}(\cdot \mid s)$  in  $\mathcal{M}$  follow the distribution  $\overline{\pi}(\cdot \mid \phi(s))$  in  $\overline{\mathcal{M}}$ . Abusing notation, we refer to  $\overline{\pi}$  as a policy in  $\mathcal{M}$ . We write  $\overline{V}^{\overline{\pi}}$  for the value function of  $\overline{\mathcal{M}}$  operating under  $\overline{\pi}$ .

Given  $\overline{\mathcal{M}}$  and  $\overline{\pi}$ , we bound the difference between  $V^{\overline{\pi}}$  and  $\overline{V}^{\overline{\pi}}$ ; the smaller the difference, the more accurately  $\overline{\mathcal{M}}$  abstracts  $\mathcal{M}$ . Computing  $V^{\overline{\pi}}$  is intractable. To overcome this, in the same spirit as Gelada et al. (2019); Delgrange et al. (2022), we define a local measure on the transitions of  $\mathcal{M}$  and  $\overline{\mathcal{M}}$  to bound the difference between the values obtained under  $\overline{\pi}$  (cf. Fig. 3). We define the *transition loss*  $L_{\overline{\mathbf{P}}}^{\overline{\pi}}$  w.r.t. a distance metric  $\mathcal{D}$  on distributions over  $\overline{\mathcal{S}}$ . We focus on the *total variation distance* (TV)  $\mathcal{D}(P, P') = 1/2 \|P - P'\|_1$  for  $P, P' \in \Delta(\overline{\mathcal{S}})$ . We compute  $L_{\overline{\mathbf{P}}}^{\overline{\pi}}$  by taking the expectation according to the stationary distribution  $\xi_{\overline{\pi}}$ :

$$L_{\overline{\mathbf{P}}}^{\overline{\pi}} = \mathbb{E}_{s \sim \xi_{\overline{\pi}}, a \sim \overline{\pi}(\cdot \mid s)} \mathcal{D}(\phi \mathbf{P}(\cdot \mid s, a), \overline{\mathbf{P}}(\cdot \mid \phi(s), a)). \quad (1)$$




Figure 3: To run  $\overline{\pi}$  in the original environment  $\mathcal{M}$ , (i) map  $s$  to  $\phi(s) = \overline{s}$ , (ii) draw  $a \sim \overline{\pi}(\cdot \mid \overline{s})$ .  $L_{\mathbf{P}}$  measures the gap (in red) between latent states produced via  $\overline{s}_1 = \phi(s')$  with  $s' \sim \mathbf{P}(\cdot \mid s, a)$  (shortened as  $\overline{s}_1 \sim \phi \mathbf{P}(\cdot \mid s, a)$ ) and those produced directly in the latent space:  $\overline{s}_2 \sim \overline{\mathbf{P}}(\cdot \mid \overline{s}, a)$ .

The superscript is omitted when clear from the context. Efficiently sampling from the stationary distribution can be done via randomized algorithms, even for unknown probabilities (Lovász and Winkler 1995; Propp and Wilson 1998).

Recall that RL is episodic, terminating when the objective is satisfied/violated or via a reset. We thus restrict  $\mathcal{M}$  to an *episodic process*, which implies ergodicity of both  $\mathcal{M}$  and  $\overline{\mathcal{M}}$  under mild conditions (Huang 2020).

**Assumption 1** (Episodic process). *The environment  $\mathcal{M}$  has a reset state  $s_{\text{reset}}$  such that (i)  $s_{\text{reset}}$  is almost surely visited under any policy, and (ii)  $\mathcal{M}$  follows the initial distribution once reset:  $\mathbf{P}(\cdot \mid s_{\text{reset}}, a) = \mathbf{I}$  for any  $a \in \mathcal{A}$ . The latent model  $\overline{\mathcal{M}}$  is also episodic with reset state  $\phi(s_{\text{reset}})$ .*

**Assumption 2.** *The abstraction preserves information regarding the objectives. Formally, let  $\langle T, \overline{T} \rangle, \langle B, \overline{B} \rangle \subseteq \mathcal{S} \times \overline{\mathcal{S}}$  be sets of target and bad states, respectively. Then, for  $\mathcal{X} \in \{T, B\}$ ,  $s \in \mathcal{X}$  iff  $\phi(s) \in \overline{\mathcal{X}}$ .<sup>1</sup> We consider the objective  $\odot(T, B)$  in  $\mathcal{M}$  and  $\odot(\overline{T}, \overline{B})$  in  $\overline{\mathcal{M}}$ .*

The following lemma establishes a bound on the difference in values based on  $L_{\mathbf{P}}$ . Notably, as  $L_{\mathbf{P}}$  goes to zero, the two models *almost surely* have the same values from every state.

**Lemma 1** (Delgrange et al. 2022). Let  $\overline{\pi}$  be a latent policy and  $\xi_{\overline{\pi}}$  be the unique stationary measure of  $\overline{\mathcal{M}}$ , then the *average value difference* is bounded by  $L_{\mathbf{P}}$ :  $\mathbb{E}_{s \sim \xi_{\overline{\pi}}} |V^{\overline{\pi}}(s) - \overline{V}^{\overline{\pi}}(\phi(s))| \leq \frac{\gamma L_{\mathbf{P}}}{1 - \gamma}$ .

The next theorem provides a more transparent bound applicable to the initial distribution, removing the need of the expectation in Lem. 1. The proof follows from plugging the stationary distribution in  $s_{\text{reset}}$  into Lem. 1 and observing that  $1/\xi_{\overline{\pi}}(s_{\text{reset}})$  is the *average episode length* (Serfozo 2009).

**Theorem 1.** The value difference from the initial states is bounded by  $L_{\mathbf{P}}$ :  $|V_{\mathbf{I}}^{\overline{\pi}} - \overline{V}_{\mathbf{I}}^{\overline{\pi}}| \leq \frac{L_{\mathbf{P}}}{\xi_{\overline{\pi}}(s_{\text{reset}})(1 - \gamma)}$ .

### 4.2 PAC estimates of the abstraction quality

Thm. 1 establishes a bound on the quality of the abstraction based on  $L_{\mathbf{P}}$  and  $\xi_{\overline{\pi}}(s_{\text{reset}})$ . Computing these quantities, however, is not possible in practice since the transition probabilities of  $\mathcal{M}$  are unknown, and even if they were known, the expectation over  $\mathcal{S}$  deems the computation infeasible.

Instead, we obtain PAC bounds on  $\xi_{\overline{\pi}}(s_{\text{reset}})$  and  $L_{\mathbf{P}}$  by simulating  $\mathcal{M}$ . The estimate of  $\xi_{\overline{\pi}}(s_{\text{reset}})$  is obtained by taking the portion of visits to  $s_{\text{reset}}$  in a simulation and Hoeffding's inequality. The estimate of  $L_{\mathbf{P}}$  is obtained as follows.

<sup>1</sup>By labeling states with atomic propositions, a standard in model checking (Delgrange et al. 2022).



When the simulation goes from  $s$  to  $s'$  following action  $a$ , we add a “reward” of  $\widehat{\mathbf{P}}(\phi(s') \mid \phi(s), a)$ . Since  $L_{\mathbf{P}}$  is a loss, we subtract the average reward from 1.

**Lemma 2.** Let  $\{\langle s_t, a_t, s'_t \rangle : 1 \leq t \leq \mathcal{T}\}$  be a set of  $\mathcal{T}$  transitions drawn from  $\xi_{\bar{\pi}}$  by simulating  $\mathcal{M}_{\bar{\pi}}$ . Let  $\widehat{L}_{\mathbf{P}} = 1 - 1/\mathcal{T} \sum_{t=1}^{\mathcal{T}} \widehat{\mathbf{P}}(\phi(s'_t) \mid \phi(s_t), a_t)$  and  $\widehat{\xi}_{\text{reset}} = 1/\mathcal{T} \sum_{t=0}^{\mathcal{T}} \mathbb{1}\{s_t = s_{\text{reset}}\}$ . Then, for all  $\varepsilon, \delta > 0$  and  $\mathcal{T} \geq \lceil -\log(\zeta)/2\varepsilon^2 \rceil$ , with at least probability  $1 - \delta$  we have that

- (i) if  $\zeta \leq \delta$ ,  $\widehat{L}_{\mathbf{P}} + \varepsilon > L_{\mathbf{P}}$ ,
- (ii) if  $\zeta \leq \delta/2$ ,  $\widehat{L}_{\mathbf{P}} + \varepsilon > L_{\mathbf{P}}$  and  $\xi_{\bar{\pi}}(s_{\text{reset}}) > \widehat{\xi}_{\text{reset}} - \varepsilon$ .

The following theorem has two key implications: (i) it establishes a lower bound on the minimum number of samples necessary to calculate the PAC upper bound for the average value difference; (ii) it suggests an online algorithm with a termination criterion for the value difference bound obtained from the initial states.

**Theorem 2** (The value bounds are PAC learnable). Let  $\{\langle s_t, a_t, s'_t \rangle : 1 \leq t \leq \mathcal{T}\}$  be  $\mathcal{T}$  transitions drawn from  $\xi_{\bar{\pi}}$  by simulating  $\mathcal{M}$  under  $\bar{\pi}$ . Then, for any  $\varepsilon, \delta > 0$ ,  $\mathcal{T} \geq \lceil -\gamma' \log(\delta') / (2\varepsilon^2(1-\gamma)\zeta) \rceil$ , with at least probability  $1 - \delta$ ,

- (i)  $\mathbb{E}_{s \sim \xi_{\bar{\pi}}} |V^{\bar{\pi}}(s) - \bar{V}^{\bar{\pi}}(\phi(s))| \leq \frac{\gamma \widehat{L}_{\mathbf{P}}}{1-\gamma} + \varepsilon$  with  $\delta' = \delta$ ,  $\gamma' = \gamma^2$ , and  $\zeta = 1$ , and
- (ii)  $|V_{\mathbf{I}}^{\bar{\pi}} - \bar{V}_{\mathbf{I}}^{\bar{\pi}}| \leq \frac{\widehat{L}_{\mathbf{P}}}{\widehat{\xi}_{\text{reset}}(1-\gamma)} + \varepsilon$  with  $\delta' = \delta/2$ ,  $\gamma' = (\widehat{L}_{\mathbf{P}} + \widehat{\xi}_{\text{reset}}(1 + \varepsilon(1-\gamma)))^2$ , and  $\zeta = \widehat{\xi}_{\text{reset}}^4$ .

Unlike (i), which enables precomputing the required number of samples to estimate the bound, (ii) allows estimating it with a probabilistic algorithm, almost surely terminating but without predetermined endpoint since  $\mathcal{T}$  relies in that case on the current approximations of  $L_{\mathbf{P}}$  and  $\xi_{\bar{\pi}}$ .

### 4.3 Obtaining latent policies during training

As highlighted in the last section, our guarantees rely on learning a policy on the representation induced by a suitable, latent abstraction. Accordingly, we propose a DRL procedure that trains the policy and the latent model *simultaneously*. Previous approaches used a two-step process: train a policy  $\pi$  in  $\mathcal{M}$  and then *distill* it. In contrast, our one-step approach alternates between optimizing a latent policy  $\bar{\pi}$  via DQN (Mnih et al. 2015) and representation learning through *Wasserstein auto-encoded MDPs* (WAE-MDPs, (Delgrange et al. 2023)). This process *avoids the distillation step* by directly learning  $\bar{\pi}$  and minimizing  $L_{\mathbf{P}}$ . That way, the DQN policy is directly optimized on the learned latent space (cf. Fig. 1(b)). We call this procedure *WAE-DQN*.

The combination of these techniques is nontrivial and requires addressing stability issues (details in Appendix D). To summarize, WAE-DQN ensures the following properties: (i)  $\phi$  groups states with close values, supporting the learning of  $\bar{\pi}$ ; (ii)  $\bar{\pi}$  prescribes the same actions for states with close behaviors, improving robustness and enabling reuse of the latent space for rooms with similar structure.

## 5 Obtaining a Planner

Fix  $\Pi$  as a collection of low-level, latent policies. In this section, we show that synthesizing a planner reduces to constructing a policy in a succinct model, where the action space coincides with the edges of the map  $\mathcal{G}$  (i.e., the choices of the planner). In the following, we describe the chain of reductions leading to this result. An overview is given in Fig. 4. We further discuss the memory requirements of the planner. Precisely, we study the following problem:

**Problem 2.** Given a two-level model  $\mathcal{H}$ , a collection of latent policies  $\Pi$ , and an objective  $\mathbb{O}$ , construct a planner  $\tau$  such that the controller  $\langle \tau, \Pi \rangle$  is optimal for  $\mathbb{O}$  in  $\mathcal{H}$ .

**Memory bounds.** Observe that planners require memory:

*Example 2.* Consider again Fig. 2(b). To reach  $v_3$  and avoid  $B_{\ell(u)}$  from  $u$ ,  $\tau$  must remember from where the room  $\ell(u)$  is entered:  $\tau$  must choose  $\uparrow$  from  $v_1$ , and  $\rightarrow$  from  $v_2$ .

Next, we establish a memory bound for an optimal planner. Upon entering a room  $R \in \mathcal{R}$ , the planner selects a direction  $d \in E$ , so the policy operating in  $R$  is  $\bar{\pi}_{R,d} \in \Pi$ , optimizing the objective  $\mathbb{O}_R^d$  to exit  $R$  via  $d$ . We construct an MDP plan  $\mathcal{M}_{\Pi} = \langle \mathcal{S}_{\Pi}, \mathcal{A}_{\Pi}, \mathbf{P}_{\Pi}, \mathbf{I}_{\Pi} \rangle$  to simulate this interaction. A state  $s^* = \langle s, v, u \rangle \in \mathcal{S}_{\Pi}$  represents  $\mathcal{H}$  being at vertex  $v$ , the room  $R = \ell(v)$  at state  $s$ , and the operating policy  $\bar{\pi}_{R,d=\langle v,u \rangle}$ . For non-exit states  $s$ , the transition function  $\mathbf{P}_{\Pi}(\cdot \mid s^*)$  follows  $\mathbf{P}_R(\cdot \mid s, a)$  with  $a \sim \bar{\pi}_{R,d}(\cdot \mid s)$ ; for exit states, the planner chooses direction  $d' \in D_{R'}$  for the next room  $R' = \ell(u)$ , where  $\mathbf{P}_{\Pi}(\cdot \mid s, d')$  follows the entrance function  $\mathcal{I}_{R'}(\cdot \mid d)$  from  $d = \langle v, u \rangle$ . Construction details are in Appendix E.

An optimal stationary policy exists for  $\mathcal{M}_{\Pi}$  (Puterman 1994) and can be implemented by a planner that memorizes the room’s entry direction. This requires *memory of size*  $|\mathcal{V}|$ , as decisions depend on any of the  $|\mathcal{V}|$  preceding vertices.

**Theorem 3.** Given low-level policies  $\Pi$ , there is a  $|\mathcal{V}|$ -memory planner  $\tau$  maximizing  $\mathbb{O}$  in  $\mathcal{H}$  iff there is a deterministic stationary policy  $\pi^*$  maximizing  $\mathbb{O}$  in  $\mathcal{M}_{\Pi}$ .

**Planner synthesis.** As a first step, we construct a *succinct MDP*  $\mathcal{M}_{\Pi}^{\mathcal{G}}$  that preserves the value of  $\mathcal{M}_{\Pi}$ . States of  $\mathcal{M}_{\Pi}^{\mathcal{G}}$  are pairs  $\langle v, u \rangle$  indicating room  $R = \ell(u)$  is entered via direction  $d = \langle v, u \rangle$ . As in  $\mathcal{M}_{\Pi}$ , a planner selects an exit direction  $d' = \langle u, v' \rangle$  for  $R$ . We use the following trick. Recall that we consider discounted properties; when  $R$  is exited via direction  $d'$  after  $j$  steps, the utility is  $\gamma^j$ . In  $\mathcal{M}_{\Pi}^{\mathcal{G}}$ , we set the probability of transitioning to  $v'$  upon choosing  $d'$  to the expected value achieved by policy  $\bar{\pi}_{R,d'}$  in  $R$ . Precisely, let  $\mathcal{M}_{\Pi}^{\mathcal{G}} = \langle \mathcal{S}, \mathcal{A}, \mathbf{P}, \mathbf{I} \rangle$  with  $\mathcal{S} = E \cup \{\perp\}$ ,  $\mathcal{A} = E$ ,  $\mathbf{I}(d_0) = 1$ ,

$$\mathbf{P}(\langle u, t \rangle \mid \langle v, u \rangle, d) = \mathbb{E}_{s \sim \mathcal{I}_{\ell(u)}(\cdot \mid \langle v, u \rangle)} \left[ V^{\bar{\pi}_{\ell(u),d}} \left( s, \mathbb{O}_{\ell(u)}^d \right) \right], \quad (2)$$

and  $\mathbf{P}(\perp \mid \langle v, u \rangle, d) = 1 - \mathbf{P}(\langle u, t \rangle \mid \langle v, u \rangle, d)$  for any  $\langle v, u \rangle \in E$  with target direction  $d = \langle u, t \rangle \in D_{\ell(u)}$ , while  $\mathbf{P}(\perp \mid \perp, d) = 1$ . The sink state  $\perp$  captures when low-level policies do not satisfy the objective.

**Theorem 4.** Let  $\langle \tau, \Pi \rangle$  be a  $|\mathcal{V}|$ -memory controller for  $\mathcal{H}$  and  $\pi$  be an equivalent policy in  $\mathcal{M}_{\Pi}$ , the values obtained





Figure 4: Chain of reductions for synthesizing a planner  $\tau$  in a two-level model  $\mathcal{H}$ .  $\mathcal{H}$  can be formulated as an explicit MDP  $\mathcal{M}$ . Once the low-level policies  $\Pi$  are learned (Fig. 1(b)), the synthesis problem reduces to constructing a stationary policy in an MDP plan  $\mathcal{M}_{\Pi}$  where  $\Pi$  is fixed and the state space of  $\mathcal{M}_{\Pi}$  encodes the directions chosen in each room. From this policy, one can derive a  $|\mathcal{V}|$ -memory planner  $\tau$  for  $\mathcal{H}$  (Thm. 3). Finally, finding a policy in  $\mathcal{M}_{\Pi}$  is equivalent to finding a policy in a succinct model  $\mathcal{M}_{\Pi}^G$  where (i) the state space corresponds to the directions from which rooms are entered, (ii) the actions to the choices of the planner, and (iii) the transition probabilities to the values achieved by the latent policy chosen (Thm. 4).

Figure 5: Uniform distribution  $\mathbf{I}_R$  (blue) and entrance function  $\mathcal{I}_R$  (red:  $\downarrow$ , green:  $\uparrow$ ). Assume  $\tau$  chooses  $\rightarrow$  in  $R$ . At training time, as  $\mathbf{I}_R$  is uniform, each state is included in the support of distribution of visited states  $\xi_{\bar{\pi}_{R,\rightarrow}}$ . Yet under a high-level controller,  $R$  is entered w.r.t.  $\mathcal{I}_R(\cdot | d \in \{\downarrow, \uparrow\})$ . To exit on the right, all states need not be visited under  $\bar{\pi}_{R,\rightarrow}$  so the distribution over visited states may differ.



under  $\pi$  for  $\odot$  in  $\mathcal{M}_{\Pi}$  are equal to those under  $\tau$  obtained in  $\mathcal{M}_{\Pi}^G$  for the reachability objective to states  $\mathcal{V} \times T$ .

We are ready to describe the algorithm to synthesize a planner. Note that the values  $V^{\bar{\pi}_{R,d}}$  in Eq. (2) are either unknown or computationally intractable. Instead, we leverage the latent model to evaluate the *latent value* of each low-level objective using standard techniques for discounted reachability objectives (de Alfaro et al. 2003). We construct  $\bar{\mathcal{M}}_{\Pi}^G$  similar to  $\mathcal{M}_{\Pi}^G$  and obtain the controller  $\langle \tau, \Pi \rangle$  by computing a planner  $\tau$  optimizing the values of  $\bar{\mathcal{M}}_{\Pi}^G$  (Puterman 1994). As  $\bar{\mathcal{M}}_{\Pi}^G$  and  $\mathcal{M}_{\Pi}^G$  have identical state spaces, planners for  $\bar{\mathcal{M}}_{\Pi}^G$  are compatible with  $\mathcal{M}_{\Pi}^G$ .

**Lifting the guarantees.** We now lift the guarantees for low-level policies to a planner operating on the two-level model, overcoming the following challenge. To learn one latent model per room  $R$  and the set of low-level policies  $\Pi$ , we run WAE-DQN independently (and possibly in parallel) in each room  $R$  (Fig. 1(b)). Viewing  $R$  as an MDP, we obtain a transition loss  $L_{\mathbf{P}}^{R,d}$  for every direction  $d$ , associated with latent policy  $\bar{\pi}_{R,d} \in \Pi$ .

Independent training introduces complications. Each room  $R$  has its own initial distribution  $\mathbf{I}_R$ , while at synthesis time, the initial distribution depends on the controller  $\pi = \langle \tau, \Pi \rangle$  and marginalizes  $\mathcal{I}_R(\cdot | d)$  over directions  $d$  chosen by  $\tau$ . Recall that  $L_{\mathbf{P}}^{R,d}$  is the TV between original and latent transition functions, averaged over  $\xi_{\bar{\pi}_{R,d}}$ , i.e., states likely to be visited under  $\bar{\pi}$  when using  $\mathbf{I}_R$  as the entrance function. The latter differs from  $\mathcal{I}_R$ , used at synthesis time. As  $\xi_{\bar{\pi}_{R,d}}$  may not align with the state distribution visited under the two-level controller  $\pi$ ,  $L_{\mathbf{P}}^{R,d}$  (and thus the guarantees from

the latent model) may become obsolete or non-reusable.

Fig. 5 illustrates the distribution shift. A detailed analysis is given in Appendix G.

Fortunately, as we will show in the following theorem, it turns out that if the initial distribution  $\mathbf{I}_R$  of each room  $R$  is well designed and provides sufficient coverage of the state space of  $R$ , it is possible to learn a *latent entrance function*  $\bar{\mathcal{I}}_R$  so that the guarantees associated with each room can be lifted to the two-level controller.

**Theorem 5.** Let  $\langle \tau, \Pi \rangle$  be a  $|\mathcal{V}|$ -memory controller for  $\mathcal{H}$  and  $\pi$  be an equivalent stationary policy in  $\mathcal{M}_{\Pi}$ .

- (Entrance loss) Define  $\bar{\mathcal{I}}_R: D_R \rightarrow \Delta(\bar{\mathcal{S}})$  and

$$L_{\mathcal{I}} = \mathbb{E}_{R,d \sim \xi_{\pi}} \mathcal{D}(\phi \mathcal{I}_R(\cdot | d), \bar{\mathcal{I}}_R(\cdot | d)),$$

where  $\xi_{\pi}$  is the stationary measure of  $\mathcal{M}_{\Pi}$  under  $\pi$ ;

- (State coverage) Assume the projection of the BSCC of  $\mathcal{M}_{\Pi}$  under  $\pi$  onto  $S_R$  is included in the BSCC of  $R$  under  $\bar{\pi}_{R,d}$  for all training rooms  $R \in \mathcal{R}$  and  $d \in D_R$ ;

Then,  $\exists K \geq 0: |V_{\mathbf{I}}^{\mathcal{M}_{\Pi}, \pi} - \bar{V}_{\bar{\mathcal{I}}}^{\bar{\mathcal{M}}_{\Pi}^G, \tau}| \leq \frac{L_{\mathcal{I}} + K \cdot \mathbb{E}_{R,d \sim \xi_{\pi}} L_{\mathbf{P}}^{R,d}}{\xi_{\pi}(s_{\text{reset}}) \cdot (1-\gamma)}$ .

Essentially, under mild conditions, the guarantees obtained for *individually trained* rooms can be *reused* for the entire two-level environment. By minimizing losses within each room *independently*, the true environment’s values increasingly align with those computed in the latent space for the high-level objective.

This theorem is the building block that enables our technique, as low-level latent policies are trained individually and in parallel *before* performing synthesis.

## 6 Case Studies

While the focus of this work is primarily of a theoretical nature, we show in the following that our theory is grounded through a navigation domain involving an agent required to reach a distant location while avoiding moving adversaries. We consider two challenging case studies. The first one consists of a large grid world of scalable size with a nontrivial observation space. The second one is a large `ViZDOM` environment (Kempka et al. 2016) with visual inputs.

To the best of our knowledge, our framework is currently the only one allowing formally verifying the values of the specification in a learned model, providing PAC bounds on the abstraction quality of this model, and synthesizing a controller in such large environments with guarantees. Thus,




Figure 6: Evaluation of WAE-DQN (low-level) and DQN (high-level) policies respectively in each room/direction and in a 9-room,  $20 \times 20$  grid world (avg. over 30 rollouts).

this section aims to show the following: (1) our method successfully trains latent policies in non-trivial settings; (2) the theoretical bounds are a good prediction for the observed behavior; (3) our low-level policies are reusable and can be composed into a strong global policy. Extra details about our setting and results are in Appendix H.

**Grid world.** The grid world environments consist of  $N$  rooms of  $m \times n$  cells, each containing at most  $l$  possible items: walls, entries/exits, power-ups, and  $A$  adversaries. The latter patrol, moving between rooms, with varying stochastic behaviors (along walls, chase the agent, or fully random). The rooms need not be identical. Each state features (i) a bitmap of rank 4 and shape  $[N, l, m, n]$  and (ii) step, power-up, and life-point (LP) counters. Note that the resulting state space is very large and policies may require, e.g., convolutional NNs to process the observations. Fig. 6 shows that DRL (here, DQN with SOTA extensions and reward shaping, (Hessel et al. 2018; Ng et al. 1999)) struggles to learn for 9 rooms/11 adversaries, while applying WAE-DQN independently in each room successfully allows learning to satisfy low-level reach-avoid objectives. A video of a synthesized controller in such an environment is available at <https://youtu.be/crowN8-GaRg>.

**ViZDoom.** We designed a map for the first-person shooter game *Doom* consisting of  $N = 8$  distinct rooms. The map includes  $A$  adversaries that actively pursue and attack the agent, reducing the agent’s health upon successful hits. Additional adversaries spawn randomly on the map (every  $\sim 60$  steps). Similar to the grid world environment, adversaries can move freely between rooms. The agent has the ability to shoot; however, missed shots incur a negative reward during the RL phase, penalizing wasted ammunition. The agent’s observations consist of (i) a single frame of the game (*visual input*), (ii) the velocity of the agent along the  $x, y$  axes, (iii) the agent’s angle w.r.t. the map, and (iv) its current health. Notice that the resulting state space is inherently colossal due to the inclusion of these variables. We are not aware of synthesis techniques that can provide guarantees in such settings. A demonstration of a synthesized controller can be viewed at <https://youtu.be/BAVLmsWEaQY>.

**Results.** We use WAE-DQN to train low-level latent models and policies in a 9-room,  $20 \times 20$  grid world as well as in the ViZDoom environment. At the start of each episode, the agent is placed in a random room, and the episode concludes successfully when the agent reaches a sub-goal. Leveraging the representation learning capabilities

	$N$	LP	$A$	avg. return ( $\gamma = 1$ )	latent value	avg. value (original)
Grid World	9	1	11	$0.5467 \pm 0.1017$	0.1378	$0.07506 \pm 0.01664$
	9	3	11	$0.7 \pm 0.09428$	0.4343	$0.01 \pm 0.00163$
	25	3	23	$0.4933 \pm 0.09832$	0.1763	$0.007833 \pm 0.002131$
	25	5	23	$0.5667 \pm 0.07817$	0.346	$0.00832 \pm 0.00288$
	49	7	47	$0.02667 \pm 0.01491$	0.004229	$5.565e-6 \pm 7e-6$
Doom	8	/	8	$0.89333 \pm 0.059628$	0.24171	$0.23405 \pm 0.014781$
	8	/	14	$0.78 \pm 0.064979$	0.16459	$0.16733 \pm 0.023117$
	8	/	20	$0.39333 \pm 0.11643$	0.086714	$0.06898 \pm 0.017788$

Table 1: Synthesis for  $\gamma = 0.99$ . Avg. return is the observed, empirical probability of reaching the high-level goal when running the synthesized two-level controller in the environment. This metric serves as a reference for the controller’s performance. Latent value is the predicted value of the high-level objective computed in the latent model. Avg. value is the empirical value of this objective approximated by simulating the environment under the controller.

of WAE-MDPs, the latent space generalizes over all rooms: we only train 4 policies (one for each direction). PAC bounds for each direction are reported in Tab. 2 ( $\epsilon = 0.01$ ,  $\delta = 0.05$ ). The

$d$	Grid World	ViZDoom
$\rightarrow$	0.50412	0.32011
$\leftarrow$	0.77787	0.44883
$\uparrow$	0.49631	0.37931
$\downarrow$	0.48058	0.48108

lower the bounds, the more accurately the latent model is guaranteed to represent the true underlying dynamics (Thm. 2). From those policies, we apply our synthesis procedure to construct a two-level controller. The results are in Tab. 1. To emphasize the reusability of the low-level components, we modify the environments by significantly increasing both the number of rooms and adversaries in the grid world (up to 50 each) and the initial number of adversaries in ViZDoom (from 8 to 20), while keeping the same latent models and policies unchanged.

In the grid world, the predicted latent values are consistent with the observed ones and comprised between the approximate return and values in the environment (averaged over 30 rollouts). In ViZDoom, the PAC bounds (Tab. 2) are lower, theoretically indicating that the latent model is of higher quality and greater accuracy. This theoretical insight is supported by the results, as the latent values are closer to the empirical, observed ones.

## 7 Conclusion

Our approach enables synthesis in environments where traditional formal synthesis does not scale. Given a high-level map, we integrate RL in the low-level rooms by training latent policies, which ensure PAC bounds on their value function. Composing with the latent policies allows to construct a high-level planner in a two-level model, where the guarantees can be lifted. Experiments show the feasibility in scenarios that are even challenging for pure DRL.

While we believe the map is a mild requirement, future work involves its relaxation to “emulate” synthesis with only the specification as input (end-to-end). In that sense, integrating skill discovery (Bagaria et al. 2021) or goal-oriented (Liu et al. 2022) RL are promising directions.

Another aspect is to refine the PAC bounds, being currently quite conservative, and obtain an estimate efficiently.



## Acknowledgments

We thank Sterre Lutz and Willem Röpke for providing valuable feedback during the preparation of this manuscript.

This research received support from the Belgian Flemish Government’s AI Research Program and DESCARTES iBOF project, the Dutch Research Council (NWO) Talent Programme (VI.Veni.222.119), Independent Research Fund Denmark (10.46540/3120-00041B), DIREC - Digital Research Centre Denmark (9142-0001B), Villum Investigator Grant S4OS (37819), and ISF grant (1679/21).

## References

- Abels, A.; Roijers, D. M.; Lenaerts, T.; Nowé, A.; and Steckelmacher, D. 2019. Dynamic Weights in Multi-Objective Deep Reinforcement Learning. In *ICML*, volume 97 of *PMRL*, 11–20. PMLR.
- Achiam, J.; Held, D.; Tamar, A.; and Abbeel, P. 2017. Constrained Policy Optimization. In *ICML*, volume 70, 22–31. PMLR.
- Alamdari, P. A.; Avni, G.; Henzinger, T. A.; and Lukina, A. 2020. Formal Methods with a Touch of Magic. In *FMCAD*, 138–147. IEEE.
- Alegre, L. N.; Bazzan, A. L. C.; Roijers, D. M.; Nowé, A.; and da Silva, B. C. 2023. Sample-Efficient Multi-Objective Learning via Generalized Policy Improvement Prioritization. In *AAMAS*, 2003–2012. ACM.
- Alshiekh, M.; Bloem, R.; Ehlers, R.; Könighofer, B.; Niekum, S.; and Topcu, U. 2018. Safe Reinforcement Learning via Shielding. In *AAAI*, 2669–2678. AAAI Press.
- Alur, R.; Bansal, S.; Bastani, O.; and Jothimurugan, K. 2022. A Framework for Transforming Specifications in Reinforcement Learning. In *Principles of Systems Design - Essays Dedicated to Thomas A. Henzinger on the Occasion of His 60th Birthday*, volume 13660 of *LNCS*, 604–624. Springer.
- Amir, G.; Schapira, M.; and Katz, G. 2021. Towards Scalable Verification of Deep Reinforcement Learning. In *FM-CAD*, 193–203. IEEE.
- Bacci, E.; Giacobbe, M.; and Parker, D. 2021. Verifying Reinforcement Learning up to Infinity. In *IJCAI*, 2154–2160. ijcai.org.
- Badings, T. S.; Romao, L.; Abate, A.; Parker, D.; Poonawala, H. A.; Stoelinga, M.; and Jansen, N. 2023. Robust Control for Dynamical Systems with Non-Gaussian Noise via Formal Abstractions. *J. Artif. Intell. Res.*, 76: 341–391.
- Bagaria, A.; Senthil, J. K.; and Konidaris, G. 2021. Skill Discovery for Exploration and Planning using Deep Skill Graphs. In *ICML*, volume 139 of *PMLR*, 521–531. PMLR.
- Baier, C.; de Alfaro, L.; Forejt, V.; and Kwiatkowska, M. 2018. Model Checking Probabilistic Systems. In *Handbook of Model Checking*, 963–999. Springer.
- Baier, C.; and Katoen, J. 2008. *Principles of model checking*. MIT Press. ISBN 978-0-262-02649-9.
- Bastani, O.; Pu, Y.; and Solar-Lezama, A. 2018. Verifiable Reinforcement Learning via Policy Extraction. In *NeurIPS*, 2499–2509.
- Bellemare, M. G.; Dabney, W.; and Munos, R. 2017. A Distributional Perspective on Reinforcement Learning. In *ICML*, volume 70 of *PMLR*, 449–458. PMLR.
- Carr, S.; Jansen, N.; and Topcu, U. 2021. Task-Aware Verifiable RNN-Based Policies for Partially Observable Markov Decision Processes. *J. Artif. Intell. Res.*, 72: 819–847.
- Chatterjee, K.; Majumdar, R.; and Henzinger, T. A. 2006. Markov Decision Processes with Multiple Objectives. In *STACS*, volume 3884 of *LNCS*, 325–336. Springer.
- Dalrymple, D.; Skalse, J.; Bengio, Y.; Russell, S.; Tegmark, M.; Seshia, S.; Omohundro, S.; Szegedy, C.; Goldhaber, B.; Ammann, N.; Abate, A.; Halpern, J.; Barrett, C.; Zhao, D.; Zhi-Xuan, T.; Wing, J.; and Tenenbaum, J. 2024. Towards Guaranteed Safe AI: A Framework for Ensuring Robust and Reliable AI Systems. arXiv:2405.06624.
- de Alfaro, L.; Henzinger, T. A.; and Majumdar, R. 2003. Discounting the Future in Systems Theory. In *ICALP*, volume 2719 of *LNCS*, 1022–1037. Springer.
- Delgrange, F.; Katoen, J.; Quatmann, T.; and Randour, M. 2020. Simple Strategies in Multi-Objective MDPs. In *TACAS*, volume 12078 of *LNCS*, 346–364. Springer.
- Delgrange, F.; Nowé, A.; and Pérez, G. A. 2022. Distillation of RL Policies with Formal Guarantees via Variational Abstraction of Markov Decision Processes. In *AAAI*, 6497–6505. AAAI Press.
- Delgrange, F.; Nowé, A.; and Pérez, G. A. 2023. Wasserstein Auto-encoded MDPs: Formal Verification of Efficiently Distilled RL Policies with Many-sided Guarantees. In *ICLR*. OpenReview.net.
- Ernst, D.; Geurts, P.; and Wehenkel, L. 2005. Tree-based batch mode reinforcement learning. *JMLR*, 6(Apr): 503–556.
- Eteessami, K.; Kwiatkowska, M. Z.; Vardi, M. Y.; and Yannakakis, M. 2008. Multi-Objective Model Checking of Markov Decision Processes. *Log. Methods Comput. Sci.*, 4(4).
- Forejt, V.; Kwiatkowska, M. Z.; and Parker, D. 2012. Pareto Curves for Probabilistic Model Checking. In *ATVA*, volume 7561 of *LNCS*, 317–332. Springer.
- Fu, J.; and Topcu, U. 2014. Probably Approximately Correct MDP Learning and Control With Temporal Logic Constraints. In *Robotics: Science and Systems X*.
- Gelada, C.; Kumar, S.; Buckman, J.; Nachum, O.; and Bellemare, M. G. 2019. DeepMDP: Learning Continuous Latent Space Models for Representation Learning. In *ICML*, volume 97, 2170–2179. PMLR.
- Germain, M.; Gregor, K.; Murray, I.; and Larochelle, H. 2015. MADE: Masked Autoencoder for Distribution Estimation. In *ICML*, volume 37 of *JMLR*, 881–889. JMLR.org.
- Giacobbe, M.; Hasanbeig, M.; Kroening, D.; and Wijk, H. 2021. Shielding Atari Games with Bounded Prescience. In *AAMAS*, 1507–1509. ACM.
- Givan, R.; Dean, T. L.; and Greig, M. 2003. Equivalence notions and model minimization in Markov decision processes. *Artif. Intell.*, 147(1-2): 163–223.

- Hansson, H.; and Jonsson, B. 1994. A Logic for Reasoning about Time and Reliability. *Formal Aspects Comput.*, 6(5): 512–535.
- Hartmanns, A.; Junges, S.; Katoen, J.; and Quatmann, T. 2018. Multi-cost Bounded Reachability in MDP. In *TACAS*, volume 10806 of *LNCS*, 320–339. Springer.
- Hasanbeig, M.; Abate, A.; and Kroening, D. 2020. Cautious Reinforcement Learning with Logical Constraints. In *AA-MAS*, 483–491.
- Hayes, C. F.; Radulescu, R.; Bargiacchi, E.; Källström, J.; Macfarlane, M.; Reymond, M.; Verstraeten, T.; Zintgraf, L. M.; Dazeley, R.; Heintz, F.; Howley, E.; Irissappane, A. A.; Mannion, P.; Nowé, A.; de Oliveira Ramos, G.; Restelli, M.; Vamplew, P.; and Roijers, D. M. 2022. A practical guide to multi-objective reinforcement learning and planning. *AAMAS*, 36(1): 26.
- Hessel, M.; Modayil, J.; van Hasselt, H.; Schaul, T.; Ostrovski, G.; Dabney, W.; Horgan, D.; Piot, B.; Azar, M. G.; and Silver, D. 2018. Rainbow: Combining Improvements in Deep Reinforcement Learning. In *AAAI*, 3215–3222. AAAI Press.
- Hinton, G. E.; Vinyals, O.; and Dean, J. 2015. Distilling the Knowledge in a Neural Network. *CoRR*, abs/1503.02531.
- Huang, B. 2020. Steady State Analysis of Episodic Reinforcement Learning. In *NeurIPS*.
- Jothimurugan, K.; Bastani, O.; and Alur, R. 2021. Abstract Value Iteration for Hierarchical Reinforcement Learning. In *AISTATS*, volume 130, 1162–1170. PMLR.
- Junges, S.; and Spaan, M. T. J. 2022. Abstraction-Refinement for Hierarchical Probabilistic Models. In *CAV*, volume 13371 of *LNCS*, 102–123. Springer.
- Kempka, M.; Wydmuch, M.; Runc, G.; Toczek, J.; and Jaskowski, W. 2016. ViZDoom: A Doom-based AI research platform for visual reinforcement learning. In *CIG*, 1–8. IEEE.
- Könighofer, B.; Bloem, R.; Ehlers, R.; and Pek, C. 2022. Correct-by-Construction Runtime Enforcement in AI - A Survey. In *Principles of Systems Design - Essays Dedicated to Thomas A. Henzinger on the Occasion of His 60th Birthday*, volume 13660 of *LNCS*, 650–663. Springer.
- Kulkarni, T. D.; Narasimhan, K.; Saeedi, A.; and Tenenbaum, J. 2016. Hierarchical Deep Reinforcement Learning: Integrating Temporal Abstraction and Intrinsic Motivation. In *NeurIPS*, 3675–3683.
- Ladosz, P.; Weng, L.; Kim, M.; and Oh, H. 2022. Exploration in deep reinforcement learning: A survey. *Inf. Fusion*, 85: 1–22.
- Larsen, K. G.; and Skou, A. 1989. Bisimulation Through Probabilistic Testing. In *POPL*, 344–352. ACM Press.
- LeCun, Y.; Boser, B. E.; Denker, J. S.; Henderson, D.; Howard, R. E.; Hubbard, W. E.; and Jackel, L. D. 1989. Backpropagation Applied to Handwritten Zip Code Recognition. *Neural Comput.*, 1(4): 541–551.
- Liu, M.; Zhu, M.; and Zhang, W. 2022. Goal-Conditioned Reinforcement Learning: Problems and Solutions. In *IJCAI*, 5502–5511. ijcai.org.
- Lovász, L.; and Winkler, P. 1995. Exact Mixing in an Unknown Markov Chain. *Electron. J. Comb.*, 2.
- Maddison, C. J.; Mnih, A.; and Teh, Y. W. 2017. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. In *ICLR*. OpenReview.net.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D.; and Riedmiller, M. A. 2013. Playing Atari with Deep Reinforcement Learning. *CoRR*, abs/1312.5602.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M. A.; Fidjeland, A.; Ostrovski, G.; Petersen, S.; Beattie, C.; Sadik, A.; Antonoglou, I.; King, H.; Kumaran, D.; Wierstra, D.; Legg, S.; and Hassabis, D. 2015. Human-level control through deep reinforcement learning. *Nat.*, 518(7540): 529–533.
- Nayak, S. P.; Egidio, L. N.; Rossa, M. D.; Schmuck, A.; and Jungers, R. M. 2023. Context-triggered Abstraction-based Control Design. *IEEE Open Journal of Control Systems*, 2: 277–296.
- Ng, A. Y.; Harada, D.; and Russell, S. 1999. Policy Invariance Under Reward Transformations: Theory and Application to Reward Shaping. In *ICML*, 278–287. Morgan Kaufmann.
- O’Cinneide, C. A. 1993. Entrywise perturbation theory and error analysis for Markov chains. *Numerische Mathematik*, 65(1): 109–120.
- Pateria, S.; Subagdja, B.; Tan, A.; and Quek, C. 2022. Hierarchical Reinforcement Learning: A Comprehensive Survey. *ACM Comput. Surv.*, 54(5): 109:1–109:35.
- Pnueli, A. 1977. The Temporal Logic of Programs. In *FOCS*, 46–57. IEEE Computer Society.
- Pnueli, A.; and Rosner, R. 1989. On the Synthesis of a Reactive Module. In *POPL*, 179–190. ACM Press.
- Propp, J. G.; and Wilson, D. B. 1998. How to Get a Perfectly Random Sample from a Generic Markov Chain and Generate a Random Spanning Tree of a Directed Graph. *J. Algorithms*, 27(2): 170–217.
- Puterman, M. L. 1994. *Markov decision processes: Discrete stochastic dynamic programming*. Wiley. ISBN 978-0-47161977-2.
- Reymond, M.; Bargiacchi, E.; and Nowé, A. 2022. Pareto Conditioned Networks. In *AAMAS*, 1110–1118. (IFAA-MAS).
- Reymond, M.; and Nowé, A. 2019. Pareto-DQN: Approximating the Pareto front in complex multi-objective decision problems. In *ALA*.
- Roderick, M.; Grimm, C.; and Tellex, S. 2018. Deep Abstract Q-Networks. In *AAMAS*, 131–138.
- Ryzhyk, L.; Chubb, P.; Kuz, I.; Sueur, E. L.; and Heiser, G. 2009. Automatic device driver synthesis with termite. In *SOSP*, 73–86. ACM.
- Serfozo, R. 2009. *Basics of Applied Stochastic Processes. Probability and Its Applications*. Springer Berlin Heidelberg. ISBN 9783540893325.

- Sootla, A.; Cowen-Rivers, A. I.; Jafferjee, T.; Wang, Z.; Mguni, D. H.; Wang, J.; and Ammar, H. 2022. Sauté RL: Almost Surely Safe Reinforcement Learning Using State Augmentation. In *ICML*, volume 162, 20423–20443. PMLR.
- Sutton, R. S.; and Barto, A. G. 1998. *Reinforcement learning - an introduction*. MIT Press. ISBN 978-0-262-19398-6.
- Sutton, R. S.; Precup, D.; and Singh, S. 1999. Between MDPs and Semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning. *Artif. Intell.*, 112(1-2): 181–211.
- Tsitsiklis, J. N. 1994. Asynchronous Stochastic Approximation and Q-Learning. *Mach. Learn.*, 16(3): 185–202.
- Tsitsiklis, J. N.; and Roy, B. V. 1997. An analysis of temporal-difference learning with function approximation. *IEEE Trans. Autom. Control.*, 42(5): 674–690.
- van Hasselt, H.; Guez, A.; and Silver, D. 2016. Deep Reinforcement Learning with Double Q-Learning. In *AAAI*, 2094–2100. AAAI Press.
- Watanabe, K.; van der Vegt, M.; Hasuo, I.; Rot, J.; and Junges, S. 2024. Pareto Curves for Compositionally Model Checking String Diagrams of MDPs. In *TACAS*, volume 14571 of *LNCS*, 279–298. Springer.
- Wiewiora, E. 2003. Potential-Based Shaping and Q-Value Initialization are Equivalent. *J. Artif. Intell. Res.*, 19: 205–208.
- Xiong, Z.; Agarwal, I.; and Jagannathan, S. 2022. HiSaRL: A Hierarchical Framework for Safe Reinforcement Learning. In *SafeAI*, volume 3087 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Yang, C.; Littman, M. L.; and Carbin, M. 2021. Reinforcement Learning for General LTL Objectives Is Intractable. *CoRR*, abs/2111.12679.
- Yang, W.; Marra, G.; Rens, G.; and Raedt, L. D. 2023. Safe Reinforcement Learning via Probabilistic Logic Shields. In *IJCAI*, 5739–5749. ijcai.org.
- Zhang, A.; McAllister, R. T.; Calandra, R.; Gal, Y.; and Levine, S. 2021. Learning Invariant Representations for Reinforcement Learning without Reconstruction. In *ICLR*. OpenReview.net.
- Žikelić, D.; Lechner, M.; Verma, A.; Chatterjee, K.; and Henzinger, T. A. 2023. Compositional Policy Learning in Stochastic Control Systems with Formal Guarantees. In *NeurIPS*.

## Appendix

### A Further Related Work

**On the importance of reachability specifications.** RL — and in particular, DRL — algorithms lack both theoretical and practical guarantees. Our approach aims to advance towards formal guarantees in partially known environments. In our work, we consider *reach-avoid specifications*, which have attracted considerable attention in recent years from the AI community — particularly within goal-oriented RL (Liu et al. 2022) and reliable safe AI (Dalrymple et al. 2024), both prominent research areas.

Importantly, *reachability properties are building blocks to verify specifications in stochastic systems*, e.g., for LTL (Pnueli 1977) or PCTL (Hansson and Jonsson 1994). Specifically, verifying specifications in an MDP typically boils down to checking the reachability to recurrent regions within a product of the MDP and an (omega-regular) automaton, or a tree decomposition of the formula, involving repeated reachability to satisfiability regions within the MDP (Baier and Katoen 2008; Baier et al. 2018). Although safety can be reduced to reachability — minimize the value of reaching bad states — we directly included safety in the specifications for convenience, given its widespread demand in RL. The analysis of such properties is necessary and the first step to enable reactive synthesis in settings like ours when limited information about the environment’s dynamics is available.

**Multi-objective reasoning.** The framework introduced in this paper provides latent models and policies that allow to formally reason about the behaviors of the agent. Real-world systems are complex and often involve multiple trade-offs between (possibly conflicting) constraints, costs, rewards, and specifications. In fact, the willingness to achieve sub-goals at the lower level of the environment while ensuring that a set of safety requirements are met is a typical example of a multi-objective problem. In essence, then, our problem involves multiple objectives, not just at the same decision level, but in a multi-level classification of decisions.

Our framework tackles *one* aspect of multi-objective decision making, which we note is not standard: traditional methods (Reymond and Nowé 2019; Abels et al. 2019; Reymond et al. 2022; Hayes et al. 2022; Alegre et al. 2023; Chatterjee et al. 2006; Etesami et al. 2008; Forejt et al. 2012; Hartmanns et al. 2018; Delgrange et al. 2020) involve the ability to reason about the multiple trade-offs by conducting multi-objective analyses (e.g., generating the *Pareto curve/set/frontier*, embedding all the compromises). In contrast, we focus on dealing and *composing with the different objectives* in order to satisfy the high-level specification.

We note that Watanabe et al. (2024) consider multi-level environments while approximating Pareto curves to deal with the compromises incurred by the low-level tasks. However, that approach relies on a model and thus exhibits tractability issues while being inapplicable when the dynamics are not fully known. Furthermore, the formalization of our multi-level environment is more permissive and allows to encode information from neighboring rooms (e.g., obstacles or adversaries moving between rooms), which also requires memory for the planner (see Sect. 5 for more information on memory requirements).

### B Remark about Episodic Processes and Ergodicity

Assumption 1 implies ergodicity of both  $\mathcal{M}$  and  $\overline{\mathcal{M}}$  under mild conditions (Huang 2020). In ergodic MDPs, each state is almost surely visited infinitely often (Baier and Katoen 2008). Thus, for unconstrained reachability goals ( $B = \emptyset$ ), while a discount factor still provides insights into how quickly the objective is achieved, optimizing the values associated with reaching the target  $T$  before the episode concludes ( $B = \{s_{\text{reset}}\}$ ) is often more appealing. This involves finding a policy  $\pi$  maximizing  $V_{\Gamma}^{\pi}(\mathbb{O}(T, B = \{s_{\text{reset}}\}))$ . In essence, this is how an RL agent is trained: learning to fulfill the low-level objective before the episode concludes.

### C Proofs from Sect. 4

**Notation.** For convenience, in the remaining of this Appendix, we may write  $s, a \sim \xi_{\pi}$  as shorthand for the distribution over  $\mathcal{S} \times \mathcal{A}$  obtained by sampling  $s$  from  $\xi_{\pi}$  and then sampling  $a$  from  $\pi(\cdot | s)$ .

*Proof of Thm. 1.* Note that

$$|V^{\bar{\pi}}(s, \mathbb{O}) - \bar{V}^{\bar{\pi}}(\phi(s), \mathbb{O})| \leq \frac{1}{\xi_{\bar{\pi}}(s)} \mathbb{E}_{s' \sim \xi_{\bar{\pi}}} |V^{\bar{\pi}}(s', \mathbb{O}) - \bar{V}^{\bar{\pi}}(\phi(s'), \mathbb{O})|$$

for any  $s \in \mathcal{S}$ . Since  $s_{\text{reset}}$  is almost surely visited episodically, *restarting* the MDP (i.e., visiting  $s_{\text{reset}}$ ) is a measurable event,



meaning that  $s_{\text{reset}}$  has a non-zero probability  $\xi_{\bar{\pi}}(s_{\text{reset}}) \in (0, 1)$ . This gives us:

$$\begin{aligned}
& |V_{\mathbf{I}}^{\bar{\pi}}(\mathbb{O}) - \bar{V}_{\mathbf{I}}^{\bar{\pi}}(\mathbb{O})| \\
&= |\mathbb{E}_{s \sim \mathbf{I}} V^{\bar{\pi}}(s, \mathbb{O}) - \mathbb{E}_{\bar{s} \sim \bar{\mathbf{I}}} \bar{V}^{\bar{\pi}}(\bar{s}, \mathbb{O})| \\
&= \frac{1}{\gamma} |\mathbb{E}_{s \sim \mathbf{I}} [\gamma \cdot V^{\bar{\pi}}(s, \mathbb{O})] - \mathbb{E}_{\bar{s} \sim \bar{\mathbf{I}}} [\gamma \cdot \bar{V}^{\bar{\pi}}(\bar{s}, \mathbb{O})]| \\
&= \frac{1}{\gamma} |V^{\bar{\pi}}(s_{\text{reset}}, \mathbb{O}) - \bar{V}^{\bar{\pi}}(\phi(s_{\text{reset}}), \mathbb{O})| \quad (\text{by Assumption 1}) \\
&\leq \frac{1}{\gamma \xi_{\bar{\pi}}(s_{\text{reset}})} \mathbb{E}_{s \sim \xi_{\bar{\pi}}} |V^{\bar{\pi}}(s, \mathbb{O}) - \bar{V}^{\bar{\pi}}(\phi(s), \mathbb{O})| \\
&\leq \frac{L_{\mathbf{P}}}{\xi_{\bar{\pi}}(s_{\text{reset}})(1-\gamma)}. \quad (\text{by Lem. 1})
\end{aligned}$$

□

*Proof of Lem. 2.* By definition of the total variation distance, we have

$$\begin{aligned}
L_{\mathbf{P}} &= \mathbb{E}_{s, a \sim \xi_{\bar{\pi}}} D(\phi(\mathbf{P}(\cdot | s, a)), \bar{\mathbf{P}}(\cdot | \phi(s), a)) \\
&= \mathbb{E}_{s, a \sim \xi_{\bar{\pi}}} \left[ \frac{1}{2} \sum_{\bar{s}' \in \bar{\mathcal{S}}} |\mathbb{P}_{s' \sim \mathbf{P}(\cdot | s, a)}[\phi(s') = \bar{s}'] - \bar{\mathbf{P}}(\bar{s}' | s, a)| \right] \\
&= \mathbb{E}_{s, a \sim \xi_{\bar{\pi}}} \left[ \frac{1}{2} \sum_{\bar{s}' \in \bar{\mathcal{S}}} |\mathbb{E}_{s' \sim \mathbf{P}(\cdot | s, a)} \mathbf{1}\{\phi(s') = \bar{s}'\} - \bar{\mathbf{P}}(\bar{s}' | s, a)| \right].
\end{aligned}$$

Notice that this quantity cannot be approximated from samples distributed according to  $\xi_{\bar{\pi}}$  alone: intuitively, we need to have access to the original transition function  $\mathbf{P}$  to be able to estimate the expectation  $\mathbb{E}_{s' \sim \mathbf{P}(\cdot | s, a)} \mathbf{1}\{\phi(s') = \bar{s}'\}$  for each single point drawn from  $\xi_{\bar{\pi}}$ .

Instead, consider now the following upper bound on  $L_{\mathbf{P}}$ :

$$L_{\mathbf{P}} \leq \mathbb{E}_{s, a \sim \xi_{\bar{\pi}}} \mathbb{E}_{s' \sim \mathbf{P}(\cdot | s, a)} D(\phi(\cdot | s'), \bar{\mathbf{P}}(\cdot | \bar{s}, a)) = L_{\mathbf{P}}^{\uparrow},$$

where  $\phi(\bar{s}' | s')$  is defined as  $\mathbf{1}\{\phi(s') = \bar{s}'\}$  for any  $s' \in \bar{\mathcal{S}}$ . This bound directly follows from Jensen's inequality. We know from (Delgrange et al. 2022) that  $\hat{L}_{\mathbf{P}} + \varepsilon \leq L_{\mathbf{P}}^{\uparrow}$  with probability at most  $\exp(-2\mathcal{T}\varepsilon^2)$ . We recall the proof for the sake of presentation:

$$\begin{aligned}
& L_{\mathbf{P}}^{\uparrow} \\
&= \mathbb{E}_{s, a, s' \sim \xi_{\bar{\pi}}} D(\phi(\cdot | s'), \bar{\mathbf{P}}(\cdot | \phi(s), a)) \\
&= \mathbb{E}_{s, a, s' \sim \xi_{\bar{\pi}}} \left[ \frac{1}{2} \sum_{\bar{s}' \in \bar{\mathcal{S}}} |\phi(\bar{s}' | s') - \bar{\mathbf{P}}(\bar{s}' | \phi(s), a)| \right] \\
&= \mathbb{E}_{s, a, s' \sim \xi_{\bar{\pi}}} \left[ \frac{1}{2} \cdot \left( (1 - \bar{\mathbf{P}}(\phi(s') | \phi(s), a)) + \sum_{\bar{s}' \in \bar{\mathcal{S}} \setminus \{\phi(s')\}} |0 - \bar{\mathbf{P}}(\bar{s}' | \phi(s), a)| \right) \right] \\
&\quad \text{(because } \phi(\bar{s}' | s') = 1 \text{ if } \phi(s') = \bar{s}' \text{ and 0 otherwise)} \\
&= \mathbb{E}_{s, a, s' \sim \xi_{\bar{\pi}}} \left[ \frac{1}{2} \cdot \left( (1 - \bar{\mathbf{P}}(\phi(s') | \phi(s), a)) + \sum_{\bar{s}' \in \bar{\mathcal{S}} \setminus \{\phi(s')\}} \bar{\mathbf{P}}(\bar{s}' | \phi(s), a) \right) \right] \\
&= \mathbb{E}_{s, a, s' \sim \xi_{\bar{\pi}}} \left[ \frac{1}{2} \cdot 2 \cdot (1 - \bar{\mathbf{P}}(\phi(s') | \phi(s), a)) \right] \\
&= \mathbb{E}_{s, a, s' \sim \xi_{\bar{\pi}}} [1 - \bar{\mathbf{P}}(\phi(s') | \phi(s), a)].
\end{aligned}$$

By Hoeffding's inequality, we obtain that  $\hat{L}_{\mathbf{P}} + \varepsilon \leq L_{\mathbf{P}}^{\uparrow}$  with probability at most  $\exp(-2\mathcal{T}\varepsilon^2)$ . Equivalently, this means that  $\hat{L}_{\mathbf{P}} + \varepsilon > L_{\mathbf{P}}^{\uparrow}$  with at least probability  $1 - \exp(-2\mathcal{T}\varepsilon^2)$ . The fact that  $\hat{L}_{\mathbf{P}} + \varepsilon > L_{\mathbf{P}}^{\uparrow} \geq L_{\mathbf{P}}$  finally yields the bound.

By applying Hoeffding's inequality again, we obtain that with at most probability  $\exp(-2\mathcal{T}\varepsilon^2)$ , we have  $\widehat{\xi}_{\text{reset}} - \varepsilon \geq \xi_{\pi}(s_{\text{reset}})$ . By the union bound, we have

$$\mathbb{P}\left(\widehat{L}_{\mathbf{P}} + \varepsilon \leq L_{\mathbf{P}}^{\uparrow} \text{ or } \widehat{\xi}_{\text{reset}} - \varepsilon \geq \xi_{\pi}(s_{\text{reset}})\right) \leq \exp(-2\mathcal{T}\varepsilon^2) + \exp(-2\mathcal{T}\varepsilon^2).$$

Finding a  $\mathcal{T} \geq 0$  which yields  $\delta \geq 2\exp(-2\mathcal{T}\varepsilon^2)$  is sufficient to ensure the bound. In that case, we have

$$\delta \geq 2\exp(-2\mathcal{T}\varepsilon^2) \Leftrightarrow \delta/2 \geq \exp(-2\mathcal{T}\varepsilon^2) \Leftrightarrow \log(\delta/2) \geq -2\mathcal{T}\varepsilon^2 \Leftrightarrow \mathcal{T} \geq \frac{-\log(\delta/2)}{2\varepsilon^2}. \quad (3)$$

Then, we have that with at least probability  $1 - \delta$ ,  $\widehat{L}_{\mathbf{P}} + \varepsilon > L_{\mathbf{P}}$  and  $\widehat{\xi}_{\text{reset}} - \varepsilon < \xi_{\pi}(s_{\text{reset}})$  if  $\mathcal{T} \geq \lceil -\log(\delta)/2\varepsilon^2 \rceil$ .  $\square$

*Proof of Thm. 2.* Let  $\zeta, \delta > 0$ , then we know by Lem. 1, Thm. 1, and Lem. 2 that

- (i)  $\mathbb{E}_{s \sim \xi_{\pi}} |V^{\pi}(s) - \bar{V}^{\pi}(\phi(s))| \leq \frac{\gamma L_{\mathbf{P}}}{1-\gamma} \leq \frac{\gamma(\widehat{L}_{\mathbf{P}} + \zeta)}{1-\gamma}$ , with probability  $1 - \delta$ . Then, to ensure an error of at most  $\varepsilon > 0$ , we need to set  $\zeta$  such that:

$$\frac{\gamma(\widehat{L}_{\mathbf{P}} + \zeta)}{1-\gamma} \leq \frac{\gamma \widehat{L}_{\mathbf{P}}}{1-\gamma} + \varepsilon \quad \Leftrightarrow \quad \frac{\gamma \zeta}{1-\gamma} \leq \varepsilon \quad \Leftrightarrow \quad \zeta \leq \frac{\varepsilon(1-\gamma)}{\gamma}.$$

Then, by Lem. 2, we need  $\mathcal{T} \geq \left\lceil \frac{-\log \delta}{2\zeta^2} \right\rceil = \left\lceil \frac{-\gamma^2 \log \delta}{2\varepsilon^2(1-\gamma)^2} \right\rceil$  samples to provide an error of at most  $\varepsilon$  with probability  $1 - \delta$ .

- (ii)  $|V_{\mathbf{I}}^{\pi} - \bar{V}_{\mathbf{I}}^{\pi}| \leq \frac{L_{\mathbf{P}}}{\xi_{\pi}(s_{\text{reset}})(1-\gamma)} \leq \frac{\widehat{L}_{\mathbf{P}} + \zeta}{(\widehat{\xi}_{\text{reset}} - \zeta)(1-\gamma)}$  with probability at least  $1 - \delta$ . Then, to ensure an error of at most  $\varepsilon > 0$ , we need to set  $\zeta$  such that:

$$\begin{aligned} \frac{\widehat{L}_{\mathbf{P}}}{\widehat{\xi}_{\text{reset}} \cdot (1-\gamma)} + \varepsilon &\geq \frac{\widehat{L}_{\mathbf{P}} + \zeta}{(\widehat{\xi}_{\text{reset}} - \zeta)(1-\gamma)} \\ \Leftrightarrow (\widehat{\xi}_{\text{reset}} - \zeta) \left( \frac{\widehat{L}_{\mathbf{P}}}{\widehat{\xi}_{\text{reset}}} + \varepsilon(1-\gamma) \right) &\geq \widehat{L}_{\mathbf{P}} + \zeta \\ \Leftrightarrow \widehat{L}_{\mathbf{P}} + \widehat{\xi}_{\text{reset}} \cdot \varepsilon(1-\gamma) - \frac{\widehat{L}_{\mathbf{P}} \cdot \zeta}{\widehat{\xi}_{\text{reset}}} - \varepsilon \cdot \zeta(1-\gamma) &\geq \widehat{L}_{\mathbf{P}} + \zeta \\ \Leftrightarrow \widehat{\xi}_{\text{reset}} \cdot \varepsilon(1-\gamma) \geq \zeta + \frac{\widehat{L}_{\mathbf{P}} \cdot \zeta}{\widehat{\xi}_{\text{reset}}} + \varepsilon \cdot \zeta(1-\gamma) &= \zeta \left( 1 + \frac{\widehat{L}_{\mathbf{P}}}{\widehat{\xi}_{\text{reset}}} + \varepsilon(1-\gamma) \right) \\ \Leftrightarrow \frac{\widehat{\xi}_{\text{reset}} \cdot \varepsilon(1-\gamma)}{1 + \frac{\widehat{L}_{\mathbf{P}}}{\widehat{\xi}_{\text{reset}}} + \varepsilon(1-\gamma)} \geq \zeta &\Leftrightarrow \frac{\widehat{\xi}_{\text{reset}}^2 \cdot \varepsilon(1-\gamma)}{\widehat{L}_{\mathbf{P}} + \widehat{\xi}_{\text{reset}} \cdot (1 + \varepsilon(1-\gamma))} \geq \zeta. \end{aligned}$$

Notice that this upper bound on  $\zeta > 0$  is well defined since

- (a)  $\widehat{\xi}_{\text{reset}}^2 \cdot \varepsilon(1-\gamma) > 0$ , and (b)  $\widehat{L}_{\mathbf{P}} + \widehat{\xi}_{\text{reset}}(1 + \varepsilon(1-\gamma)) > 0$ .

Then, setting  $\zeta \leq \frac{\widehat{\xi}_{\text{reset}}^2 \cdot \varepsilon(1-\gamma)}{\widehat{L}_{\mathbf{P}} + \widehat{\xi}_{\text{reset}}(1 + \varepsilon(1-\gamma))}$  means by Lem. 2 that we need

$$\mathcal{T} \geq \left\lceil \frac{-\log(\delta/2)}{2\zeta^2} \right\rceil \geq \left\lceil \frac{-\log(\delta/2) \left( \widehat{L}_{\mathbf{P}} + \widehat{\xi}_{\text{reset}}(1 + \varepsilon(1-\gamma)) \right)^2}{2\widehat{\xi}_{\text{reset}}^4 \cdot \varepsilon^2(1-\gamma)^2} \right\rceil$$

samples to provide an error of at most  $\varepsilon$  with probability at least  $1 - \delta$ .  $\square$

## D WAE-DQN

In this section, we give additional details on WAE-DQN, which combines representation (WAE-MDP) and policy (DQN) learning. Before presenting the algorithm, we briefly recall basic RL concepts. **Q-Learning.** Q-learning is an RL algorithm whose goal is to learn the optimal solution of the Bellman equation (Puterman 1994):  $Q^*(s, a) = \mathbb{E}_{s' \sim \mathcal{P}(\cdot|s,a)} [rew(s, a, s') + \gamma \cdot \max_{a' \in \mathcal{A}} Q^*(s', a')]$  for any  $(s, a) \in \mathcal{S} \times \mathcal{A}$ , with

$$\mathbb{E}_{s_0 \sim \mathcal{I}} \left[ \max_{a \in \mathcal{A}} Q^*(s_0, a) \right] = \max_{\pi} \mathbb{E}_{\rho \sim \text{Pr}_{\pi}^{\mathcal{M}}} \left[ \sum_{i \geq 0} \gamma^i \cdot r_i \right].$$

To do so, Q-learning relies on learning *Q-values* iteratively: at each step  $i \geq 0$ , a transition  $\langle s, a, r, s' \rangle$  is drawn in  $\mathcal{M}$ , and  $Q_{i+1}(s, a) = Q_i(s, a) + \alpha(r + \gamma \max_{a' \in \mathcal{A}} Q_i(s', a') - Q_i(s, a))$  for a given learning rate  $\alpha \in (0, 1)$ . Under some assumptions,  $Q_i$  is guaranteed to converge to  $Q^*$  (Tsitsiklis 1994). Q-learning is implemented by maintaining a table of size  $|\mathcal{S} \times \mathcal{A}|$  of the Q-values. This is intractable for environments with large or continuous state spaces.

**Deep Q-networks (DQN,** (Mnih et al. 2015)) is an established technique to scale Q-learning (even for continuous state spaces), at the cost of convergence guarantees, by approximating the Q-values in parameterized NNs. By fixing a network  $Q(\cdot, \theta)$  and, for stability (Tsitsiklis and Roy 1997), periodically fixing a parameter assignment  $\hat{\theta}$ , DQN obtains *the target network*  $Q(\cdot, \hat{\theta})$ . Q-values are then optimized by applying gradient descent on the following loss function:

$$L_{\text{DQN}}(\theta) = \mathbb{E}_{s,a,r,s' \sim \mathcal{B}} \left( r + \gamma \max_{a' \in \mathcal{A}} Q(s', a'; \hat{\theta}) - Q(s, a; \theta) \right)^2, \quad (4)$$

where  $\pi^\epsilon$  is an  $\epsilon$ -greedy exploration strategy, i.e.,  $\pi^\epsilon(a | s) = (1 - \epsilon) \mathbf{1}\{a = \arg \max_{a'} Q(s, a')\} + \epsilon/|\mathcal{A}|$  for some  $\epsilon \in (0, 1)$ . In practice,  $\xi_\pi$  is emulated by a *replay buffer*  $\mathcal{B}$  where encountered transitions are stored and then sampled later on to minimize  $L_{\text{DQN}}(\theta)$ .

**Wasserstein auto-encoded MDP (WAE-MDP,** (Delgrange et al. 2023)) is a distillation technique providing PAC guarantees. Given an MDP  $\mathcal{M}$ , a policy  $\pi$  and the number of states in  $\bar{\mathcal{M}}$ , the transition probabilities and embedding function  $\phi$  (both modeled by NNs) are learned by minimizing  $L_{\mathcal{P}}$  via gradient descent. Also, a policy  $\bar{\pi}$  in  $\bar{\mathcal{M}}$  is distilled such that  $\bar{\mathcal{M}}$  exhibits *bisimilarly close* (Larsen and Skou 1989; Givan et al. 2003; Delgrange et al. 2022) behaviors to  $\mathcal{M}$  when executing  $\bar{\pi}$ , providing PAC guarantees on the difference of the two values from Lem. 1. WAE-MDPs enjoy *representation* guarantees that any states clustered to the same latent representation yield close values when  $L_{\mathcal{P}}$  is minimized (Delgrange et al. 2022): for any latent policy  $\bar{\pi}$  and  $s_1, s_2 \in \mathcal{S}$ ,  $\phi(s_1) = \phi(s_2)$  implies  $|V^{\bar{\pi}}(s_1) - V^{\bar{\pi}}(s_2)| \leq \frac{\gamma L_{\mathcal{P}}}{1-\gamma} (1/\xi_{\bar{\pi}}(s_1) + 1/\xi_{\bar{\pi}}(s_2))$ .

**WAE-DQN.** Our procedure (Fig. 7) unifies the training and distillation steps (Alg. 1). Intuitively, a WAE-MDP and a (latent) DQN policy are learned in round-robin fashion: the WAE-MDP produces the input representation (induced by  $\phi$ ) that the DQN agent uses to optimize its policy  $\bar{\pi}$ . At each step  $t = 1, \dots, \mathcal{T}$ , the environment is explored via a strategy to collect transitions in a replay buffer. Each training step consists of two optimization rounds. First, we optimize the parameters of  $\bar{\mathcal{P}}$  and  $\phi$ . Second, we optimize DQN’s parameters to learn the policy as in DQN. DQN may further backpropagate gradients through  $\phi$ . We use a *target embedding function*  $\hat{\phi}$  for stability purposes, similar to (Zhang et al. 2021). This is consistent with DQN’s target-networks approach: the weights of  $\hat{\phi}$  are periodically synchronized with those of  $\phi$ . Then,  $\hat{\phi}$  is paired with the DQN’s target network, which allows avoiding oscillations and shifts in the representation (a.k.a. moving target issues).

WAE-DQN learns a tractable model of the environment in parallel to the agent’s policy (Algorithm 1). Precisely, the algorithm alternates between optimizing the quality of the abstraction as well as the representation of the original state space via a WAE-

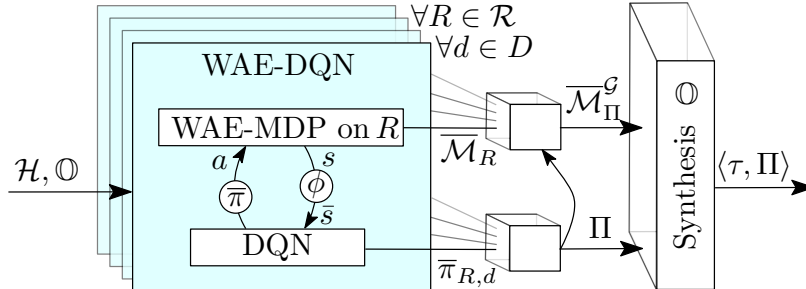


Figure 7: Given  $\mathcal{H}$  and  $\mathcal{O}$ , we run WAE-DQN in each room  $R \in \mathcal{R}$  and direction  $d \in D$  in parallel, yielding embedding  $\phi$ , latent MDPs, and policies  $\Pi$  with PAC guarantees. We then synthesize planner  $\tau$  to maximize  $\mathcal{O}$  in succinct model  $\bar{\mathcal{M}}_{\Pi}^{\mathcal{G}}$ , aggregated as per the map of  $\mathcal{H}$ , given as graph  $\mathcal{G}$ .

---

**Algorithm 1: WAE-DQN**

---

**Input:** steps  $\mathcal{T}$ , model updates  $\mathcal{N}$ , batch sizes  $B_{\text{WAE}}, B_{\text{DQN}}$ , and  $\alpha, \epsilon \in (0, 1)$ ;

Initialize the target parameters:  $\langle \hat{\iota}, \hat{\theta}_{\text{DQN}} \rangle \leftarrow$  **copy** the parameters  $\langle \iota, \theta_{\text{DQN}} \rangle$

Initialize replay buffer  $\mathcal{B}$  with transitions from random exploration of  $\mathcal{M}$

**for**  $t \in \{1, \dots, \mathcal{T}\}$  with  $s_0 \sim \mathbf{I}$  **do**

Embed  $s_t$  into the latent space:  $\bar{s} \leftarrow \phi(s_t)$

Choose action  $a_t$ :  $\begin{cases} \text{w.p. } (1-\epsilon), \text{ define } a_t = \arg \max_a Q(\bar{s}, a), \text{ and} \\ \text{w.p. } \epsilon, \text{ draw } a_t \text{ uniformly from } \mathcal{A} \end{cases}$

Execute  $a_t$  in the environment  $\mathcal{M}$ , receive reward  $r_t$ , and observe  $s_{t+1}$

Store the transition in the replay buffer:  $\mathcal{B} \leftarrow \mathcal{B} \cup \{\langle s_t, a_t, r_t, s_{t+1} \rangle\}$

**repeat**  $\mathcal{N}$  **times**

Sample a batch of size  $B_{\text{WAE}}$  from  $\mathcal{B}$ :  $X \leftarrow \{\langle s, a, r, s' \rangle_i\}_{i=1}^{B_{\text{WAE}}} \sim \mathcal{B}$

Update  $\iota$  and  $\theta_{\text{WAE}}$  on the batch  $X$  by minimizing the WAE-MDP loss (including  $L_{\mathbf{P}}$ ) for the latent policy  $\bar{\pi}^\epsilon$   
 $\triangleright$  *details in (Delgrange et al. 2023)*

**for**  $i \in \{1, \dots, B_{\text{DQN}}\}$  **do**

Sample a transition from  $\mathcal{B}$ :  $s, a, r, s' \sim \mathcal{B}$

Compute the target:  $\hat{y} \leftarrow r + \gamma \max_{a' \in \mathcal{A}} Q(\phi(s'); \hat{\iota}, a'; \hat{\theta}_{\text{DQN}})$

Compute the DQN loss (Eq. 4):  $L_i \leftarrow (Q(\phi(s); \iota, a; \theta_{\text{DQN}}) - \hat{y})^2$

Update  $\iota$  and  $\theta_{\text{DQN}}$  by minimizing  $1/B_{\text{DQN}} \sum_{i=1}^{B_{\text{DQN}}} L_i$

Update the target params.:  $\hat{\iota} \leftarrow \alpha \cdot \iota + (1 - \alpha) \cdot \hat{\iota}$ ;  $\hat{\theta} \leftarrow \alpha \cdot \theta_{\text{DQN}} + (1 - \alpha) \cdot \hat{\theta}$

**return**  $\phi, \bar{\mathcal{M}}, \text{ and } \bar{\pi}$

---

MDP, and optimizing a latent policy via DQN. We respectively denote the parameters of the state embedding function  $\phi$ , those of the latent transition function  $\bar{\mathbf{P}}$ , and those of the Deep Q-networks by  $\iota, \theta_{\text{WAE}}$ , and  $\theta_{\text{DQN}}$ .

## E Explicit Construction of the MDP Plan

Along this section, fix a two-level model  $\mathcal{H} = \langle \mathcal{G}, \ell, \mathcal{R}, v_0, \langle d_0, d_1 \rangle \rangle$  with its explicit MDP representation  $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathbf{P}, \mathbf{I} \rangle$ .

To enable high-level reasoning when the rooms are aggregated into a unified model, we add the following assumption.

**Assumption 3.** *All rooms  $R \in \mathcal{R}$  share the same reset state  $s_{\text{reset}}$  in  $\mathcal{H}$ .*

Note that Assumption 3 is a technicality that can be trivially met in every two-level model  $\mathcal{H}$ : it just requires that when a reset is triggered in a room  $R$  of  $\mathcal{H}$ , the whole model is globally reset, and not only  $R$ , locally.

We define an MDP  $\mathcal{M}_{\Pi}$ , called an *MDP plan*, such that policies in  $\mathcal{M}_{\Pi}$  correspond to planners. Recall that the actions that a planner performs consist of choosing a policy once entering a room. Accordingly, we define  $\mathcal{M}_{\Pi} = \langle \mathcal{S}_{\Pi}, \mathcal{A}_{\Pi}, \mathbf{P}_{\Pi}, \mathbf{I}_{\Pi} \rangle$ . States in  $\mathcal{S}_{\Pi}$  keep track of the location in a room as well as the target of the low-level policy that is being executed. Formally,

$$\mathcal{S}_{\Pi} = (\cup_{R \in \mathcal{R}} (\mathcal{S}_R \setminus \{s_{\text{reset}}\}) \times E) \cup \{s_{\text{reset}}, \perp\},$$

where a pair  $\langle s, v, u \rangle \in \mathcal{S}_{\Pi}$  means that the current room is  $v$ , the target of the low-level policy is to exit the room in direction  $d = \langle v, u \rangle$ , and the current state is  $s \in \mathcal{S}_{\ell(v)}$ . Following Assumption 3, the rooms share the reset state  $s_{\text{reset}}$ , and  $\perp$  is a special sink state that we add for technical reasons to disable actions in states. The initial distribution  $\mathbf{I}_{\Pi}$  has for support  $\{\langle s, v, u \rangle \in \mathcal{S}_{\Pi} \mid v = v_0 \text{ and } \langle v, u \rangle = d_1\}$  where states  $s \in \mathcal{S}_{\ell(v_0)}$  are distributed according to  $\mathcal{I}_{\ell(v_0)}(\cdot \mid d_0)$ . Actions chosen correspond to those of the planner — only required when entering a room — so the action space is  $\mathcal{A}_{\Pi} = E \cup \{*\}$ , where  $d \in E$  means that the low-level policy that is executed exits via direction  $d$ , and  $*$  is a special action that is used inside a room, indicating no change to the low-level policy. Note that once  $d$  is chosen, we only allow exiting the room through direction  $d$ . We define the transition function. Let  $\mathbf{P}$  be the transition function of the explicit MDP  $\mathcal{M}$ . For a state  $\langle s, v, u \rangle \in \mathcal{S}_{\Pi}$  with  $d = \langle v, u \rangle$ ,

(i) if  $s$  is not an exit state, i.e.,  $s \notin \mathcal{O}_{\ell(v)}(d)$ , then the action is chosen by the low-level policy  $\pi_{\ell(v), d}$ , and the next state is chosen according to the transitions of  $\ell(v)$ : for every  $s' \in \mathcal{S}_{\ell(v)} \setminus \{s_{\text{reset}}\}$ ,

$$\mathbf{P}_{\Pi}(\langle s', v, u \rangle \mid \langle s, v, u \rangle, *) = \mathbb{E}_{a \sim \pi_{\ell(v), d}(\cdot \mid s)} \mathbf{P}(\langle s', v \rangle \mid \langle s, v \rangle, a); \quad (5)$$

(ii) if  $s$  is an exit state in direction  $d$ , i.e.,  $s \in \mathcal{O}_{\ell(v)}(d)$ , the next room is entered according to the entrance function from direction  $d$  and the planner needs to choose a new target direction  $d'$ : for every  $s' \in \mathcal{S}_{\ell(u)} \setminus \{s_{\text{reset}}\}$  and edge  $d' = \langle u, t \rangle \in \text{out}(u)$ :

$$\mathbf{P}_{\Pi}(\langle s', u, t \rangle \mid \langle s, v, u \rangle, d') = \mathbf{P}(\langle s', u \rangle \mid \langle s, v \rangle, a_{\text{exit}}) = \mathcal{I}_{\ell(u)}(s' \mid d) \quad (6)$$



(iii) the reset state is handled exactly as in the explicit model  $\mathcal{M}$ :

$$\mathbf{P}_{\Pi}(s_{\text{reset}} \mid \langle s, v, u \rangle, *) = \mathbb{E}_{a \sim \pi_{\ell(v),d}} \mathbf{P}(s_{\text{reset}} \mid \langle s, v \rangle, a),$$

and  $\mathbf{P}_{\Pi}(\cdot \mid s_{\text{reset}}, a) = \mathbf{I}_{\Pi}$  for any  $a \in \mathcal{A}_{\Pi}$ ;

(iv) any other undefined distribution transitions deterministically to the sink state  $\perp$  so that  $\mathbf{P}_{\Pi}(\perp \mid \perp, a) = 1$  for any  $a \in \mathcal{A}_{\Pi}$ .

**Proper policies.** We say that a policy  $\pi$  for  $\mathcal{M}_{\Pi}$  is *proper* if the decisions of  $\pi$  ensure to almost surely avoid  $\perp$ , i.e.,  $V^{\pi}(s, \mathbb{O}(T = \{\perp\}, B = \emptyset)) = 0$  for all states  $s \in \mathcal{S}_{\Pi} \setminus \{\perp\}$ . Note that *improper* policies strictly consist of those which prescribe to not follow the low-level policy corresponding to the current objective and do not select a new target direction when exiting.

*In the following proofs, we restrict our attention to proper policies.*

**Property 1** (High-level objective in the MDP plan). In  $\mathcal{M}_{\Pi}$ , the high-level objective  $\mathbb{O}$  translates to the reach-avoid objective  $\mathbb{O}(\mathbf{T}, \mathbf{B})$  where  $\mathbf{T} = \{\langle s, v, u \rangle \in \mathcal{S}_{\Pi} \mid v \in T\}$  and  $\mathbf{B} = \{\langle s, v, u \rangle \in \mathcal{S}_{\Pi} \mid s \notin B_{\ell(v)}\}$  for the high-level objective  $\diamond T$  so that  $B_R$  is the set to avoid in room  $R$ .

## F Proofs from Sect. 5

**Lemma 3** (Equivalence of policies in the two-level model and plan). There exists an equivalence between planners with memory of size  $|\mathcal{V}|$  in the two-level model  $\mathcal{H}$  and proper deterministic stationary policies in the MDP plan  $\mathcal{M}_{\Pi}$  that preserves the values of their respective objective under equivalent planners and policies.

*Proof.* Let  $\tau$  be a planner for  $\mathcal{H}$  with memory of size  $|\mathcal{V}|$ . Let us encode  $\tau$  as a finite Mealy machine whose inputs are graph vertices  $\mathcal{V}$  and outputs are directions, i.e.,  $\tau = \langle \mathcal{Q}, \tau_a, \tau_u, q_0 \rangle$  where  $\mathcal{Q}$  is a set of memory states with  $|\mathcal{Q}| = |\mathcal{V}|$ ,  $\tau_a: \mathcal{V} \times \mathcal{Q} \rightarrow E$  is the next action function,  $\tau_u: \mathcal{V} \times \mathcal{Q} \times E \rightarrow \mathcal{Q}$  is the memory update function, and  $q_0$  is the initial memory state.

Let us consider the two-level controller  $\langle \tau, \Pi \rangle$  as a policy in the explicit MDP  $\mathcal{M}$ . Since  $\tau$  is a planner, we require that

1.  $\tau_a(v_0, q_0) = d_1$ , and
2. if  $\tau_a(v, q) = d$ , then  $d \in \text{out}(v)$  for any  $v \in \mathcal{V}, q \in \mathcal{Q}$ .

Intuitively,  $\tau_a$  chooses the direction to follow in the current room based on the current memory state  $q$ , and  $\tau_u$  describes how to update the memory, based on the current room, the current memory state, and the direction chosen. By definition of the two-level controller  $\langle \tau, \Pi \rangle$  (see Sect. 3),  $\tau_a$  is used at each time step in the current room, to know which low-level policy to execute, and  $\tau_u$  is triggered once an exit state is reached, to switch to the next memory state that will determine the direction to follow in the next room.

Then,  $\Pr_{\langle \tau, \Pi \rangle}^{\mathcal{M}}$  is a distribution over the product of the paths of  $\mathcal{M}$  and the sequence of memory states of  $\tau$ . Following the definition of the controller  $\langle \tau, \Pi \rangle$  (cf. Sect. 3), the measure  $\Pr_{\langle \tau, \Pi \rangle}^{\mathcal{M}}$  can be obtained inductively as follows. For a state  $\langle s, v \rangle \in \mathcal{S}$ ,  $\Pr_{\langle \tau, \Pi \rangle}^{\mathcal{M}}(s, v, q) = \mathcal{I}_{\ell(v_0)}(s \mid d_0)$  if  $v_0 = v$  and  $q = q_0$ , and assigns a zero probability otherwise. The probability of a path  $\rho = s_0, v_0, q_0, \dots, s_{t-1}, v_{t-1}, q_{t-1}, s_t, v_t, q_t$  is given as follows

- (a) if  $s_{t-1}$  is not an exit state, the low-level policy is executed in direction  $d = \tau_a(v_{t-1}, q_{t-1})$  and both the current vertex and memory state must remain unchanged:

$$\Pr_{\langle \tau, \Pi \rangle}^{\mathcal{M}}(s_0, v_0, q_0, \dots, s_{t-1}, v_{t-1}, q_{t-1}) \cdot \mathbb{E}_{a \sim \pi_{\ell(v_t),d}(\cdot \mid s)} \mathbf{P}(\langle s_t, v_t \rangle \mid \langle s_{t-1}, v_{t-1} \rangle, a)$$

if  $s_{t-1} \notin \mathcal{O}_{\ell(v_{t-1})}(d)$  with  $d = \tau_a(v_{t-1}, q_{t-1})$ ,  $v_t = v_{t-1}$ , and  $q_t = q_{t-1}$ ;

- (b) if  $s_{t-1}$  is an exit state in the direction prescribed in  $q_{t-1}$ , then this direction should point to  $v_t$  and the memory state must be updated to  $q_t$ :

$$\Pr_{\langle \tau, \Pi \rangle}^{\mathcal{M}}(s_0, v_0, q_0, \dots, s_{t-1}, v_{t-1}, q_{t-1}) \cdot \mathcal{I}_{\ell(v_t)}(s_t \mid d)$$

if  $s_{t-1} \in \mathcal{O}_{\ell(v_{t-1})}(d)$  with  $d = \tau_a(v_{t-1}, q_{t-1}) = \langle v_{t-1}, v_t \rangle$ , and  $q_t = \tau_u(v_{t-1}, q_{t-1}, d)$ ;

- (c) if  $s_{t-1}$  is the reset state, by Assumptions 1 and 3, the planner must be reset as well:

$$\Pr_{\langle \tau, \Pi \rangle}^{\mathcal{M}}(s_0, v_0, q_0, \dots, s_{t-1}, v_{t-1}, q_{t-1}) \cdot \mathcal{I}_{\ell(v_0)}(s_t \mid d_0)$$

if  $s_{t-1} = s_{\text{reset}}$ ,  $v_t = v_0$ , and  $q_t = q_0$ ;

- (d) zero otherwise.

Notice that renaming  $\mathcal{Q}$  to  $\mathcal{V}$  so that for all  $q \in \mathcal{Q}$ ,  $q$  is changed to  $u \in \mathcal{V}$  (i.e.,  $q \mapsto u$ ) whenever  $\tau_a(v, q) = \langle v, u \rangle$  is harmless, since the probability measure remains unchanged. From now on, we consider that  $\mathcal{Q}$  has been renamed to  $\mathcal{V}$  in this manner.

Now, define the relation  $\equiv$  between planners in  $\mathcal{M}$  and policies in  $\mathcal{M}_\Pi$  as<sup>2</sup>

$$\tau \equiv \pi \quad \text{if and only if}$$

$$\pi(\langle s, v, u \rangle) = \begin{cases} \tau_a(u, \cdot) \circ \tau_u(v, u, d = \cdot) \circ \tau_a(v, u) & \text{if } s \in \mathcal{O}_{\ell(v)}(\langle v, u \rangle) \\ * & \text{otherwise.} \end{cases}$$

By construction of  $\mathcal{M}_\Pi$ , modulo the renaming of  $\mathcal{Q}$  to  $\mathcal{V}$ ,  $\Pr_{\langle \tau, \Pi \rangle}^{\mathcal{M}} = \Pr_{\pi}^{\mathcal{M}_\Pi}$  for any  $\tau, \pi$  in relation  $\tau \equiv \pi$ : condition (i) is equivalent to (a), condition (ii) is equivalent to (b), and condition (iii) is equivalent to (c). Note that the only policies  $\pi$  which cannot be in relation with some planner  $\tau$  are *improper policies*, i.e., those choosing actions leading to the sink state  $\perp$  (see condition (iv)). Such policies are discarded by assumption.

The result follows from the fact that, modulo the renaming of  $\mathcal{Q}$  to  $\mathcal{V}$ , planners and policies in relation  $\equiv$  lead to the same probability space.  $\square$

**Theorem 6.** For a fixed collection of low-level policies  $\Pi$ , a memory of size  $|\mathcal{V}|$  is necessary and sufficient for the planner to maximize the values of  $\mathbb{O}$  in the two-level model  $\mathcal{H}$ .

*Proof.* The necessity of a memory of size  $|\mathcal{V}|$  is shown in Example 2. The sufficiency follows from Thm. 3 and the fact that a deterministic stationary policy is sufficient to maximize constrained, discounted reachability objectives in MDPs (Puterman 1994; Baier and Katoen 2008) (in particular in  $\mathcal{M}_\Pi$ ).

To see how, let  $\pi^*$  be a proper optimal deterministic stationary policy in  $\mathcal{M}_\Pi$ . Note that one can always find a proper optimal policy from an improper one: if  $\pi^*$  is improper, it is necessarily because a prohibited action has been chosen *after* having reached the target, which can be replaced by any other action without changing the value of the objective. Consider a planner  $\tau$  in the two-level model  $\mathcal{H}$  which is equivalent to  $\pi^*$  (Lem. 3). Then,  $\tau$  is optimal for the high-level objective in  $\mathcal{H}$  (since the probability space of the two models is the same), and  $\tau$  uses a memory of size  $|\mathcal{V}|$ .  $\square$

**Succinct MDP** In the following, we take a closer look at the construction of the succinct MDP  $\mathcal{M}_\Pi^G$ . We then prove Thm. 4.

The next example illustrates how setting transition probabilities to be expected values maintains the values between the models.

*Example 3.* Consider the explicit model of Fig. 2(a), projected on two dimensions in Fig. 8. Each directed arrow corresponds to a transition with a non-zero probability. A state of the form  $\langle s, v \rangle$  indicates that the agent is in state  $s$  of room  $\ell(v)$ . Consider a path  $\rho$  that enters  $\ell(v_0) = R_0$ , exits after  $i = 3$  steps ( $s_0 \rightarrow s_1 \rightarrow s_2 \xrightarrow{a_{exit}}$ ), enters  $\ell(u) = R_1$ , exits after  $j = 3$  steps ( $s_0 \rightarrow s_1 \rightarrow s_2 \xrightarrow{a_{exit}}$ ), and finally reaches the high-level goal. The prefix of  $\rho$  in  $R_0$  is discounted to  $\gamma^3$  when the agent exits. Similarly, the suffix of  $\rho$  in  $R_1$  is discounted to  $\gamma^3$ . Once in the goal, the agent gets a “reward” of one (the goal is reached). The discounted reward obtained along  $\rho$  is thus  $\gamma^{i+j} = \gamma^6$ . In expectation, this corresponds to multiplying the values in the individual rooms and, in turn, with the semantics of  $\mathcal{M}_\Pi^G$  where probabilities are multiplied along a path.




Figure 8: Projection of Fig. 2(a) on two dimensions

Explicitly, the transition function can be re-formalized as follows. Let  $v, u \in \mathcal{V}$ ,  $d \in E$ , and  $d' \in E \cup \{\perp\}$ ,  $\mathbf{P}$  defined as

$$\mathbf{P}(d' \mid \langle v, u \rangle, d) = \begin{cases} \mathbb{E}_{s \sim \mathcal{I}_{\ell(u)}(\cdot \mid \langle v, u \rangle)} V^{\bar{\pi}_{\ell(u), d}}(s) & \text{if } d = d' \in \text{out}(u), \\ 1 - \mathbf{P}(d \mid \langle v, u \rangle, d) & \text{if } d' = \perp \text{ and } d \in \text{out}(u), \\ 1 & \text{if } d' = \perp \text{ and } d \notin \text{out}(u), \text{ and} \\ 0 & \text{otherwise,} \end{cases}$$

(7)

<sup>2</sup>Notice the slight asymmetry induced by Mealy machines: while the policy must decide the next direction in exit states, the planner just need update its memory state (Eq. (b)).

while  $\mathbf{P}(\perp \mid \perp, d) = 1$ .

For convenience, in the following, we assume that  $s_{\text{reset}} \in B_R$  for each room  $R$ , which is consistent with the remark made in Appendix B. The construction of  $\mathcal{M}_{\Pi}^G$  and Theorem 4 can be generalized by additionally wisely handling the reset state in  $\mathbf{P}$ .

For the sake of clarity, we formally restate Thm. 4:

**Theorem 7** (Value equality in the succinct model). Let  $\langle \tau, \Pi \rangle$  be a hierarchical controller for  $\mathcal{H}$  with a  $|\mathcal{V}|$ -memory planner  $\tau$ . Denote by  $V_{\mathcal{M}_{\Pi}}^{\pi}(\mathbb{O})$  the initial value of  $\mathcal{M}_{\Pi}$  running under a policy  $\pi$  equivalent<sup>3</sup> to  $\tau$  in  $\mathcal{M}_{\Pi}$  for the reach-avoid objective  $\mathbb{O}$  of Property 1. Moreover, denote by  $V_{\mathcal{M}_{\Pi}^G}^{\tau}(\diamond T)$  the initial value obtained in  $\mathcal{M}_{\Pi}^G$  when the agent follows the decisions of  $\tau$  for the reachability objective to states of the set  $\mathcal{V} \times T$ , i.e., the reach-avoid objective  $\mathbb{O}(\mathcal{V} \times T, \emptyset)$ . Then, assuming  $v_0 \notin T$  (the case where  $v_0 \in T$  is trivial),

$$V_{\mathcal{M}_{\Pi}}^{\pi}(\mathbb{O}) = V_{\mathcal{M}_{\Pi}^G}^{\tau}(\diamond T).$$

*Proof.* Given any MDP  $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathbf{P}, \mathbf{I} \rangle$ , we start by recalling the definition of the value function of any reach-avoid objective of the form  $\mathbb{O}(T, B)$  with  $T, B \subseteq \mathcal{S}$  for a discount factor  $\gamma \in (0, 1)$  and a policy  $\pi$ :

$$V_{\mathbf{I}}^{\pi}(\mathbb{O}) = \mathbb{E}_{\rho \sim \text{Pr}_{\pi}^{\mathcal{M}}} \left[ \sup_{i \geq 0} \gamma^i \mathbb{1} \{s_i \in T\} \cdot \mathbb{1} \{\forall j \leq i, s_j \notin B\} \right], \quad (8)$$

where  $s_i$  denotes the  $i^{\text{th}}$  state of  $\rho$ . Intuitively, this corresponds to the expected value of the discount scaled to the time step of *the first visit of the set  $T$* , ensuring that the set of bad states  $B$  is not encountered before this first visit.

First, notice that the reach-avoid property can be merely reduced to a simple reachability property by making absorbing the states of  $B$  (Baier and Katoen 2008). Precisely, write  $\mathcal{M}^{\circ B}$  for the MDP  $\mathcal{M}$  where we make all states from  $B$  absorbing, i.e., where  $\mathbf{P}$  is modified so that  $\mathbf{P}(s \mid s, a) = 1$  for any  $s \in B$  and  $a \in \mathcal{A}$ . Then, one can get rid of the indicator  $\mathbb{1} \{\forall j \leq i, s_j \notin B\}$  in Eq. (8) by considering infinite paths of  $\mathcal{M}^{\circ B}$ :

$$\begin{aligned} V_{\mathbf{I}}^{\pi}(\mathbb{O}) &= \mathbb{E}_{\rho \sim \text{Pr}_{\pi}^{\mathcal{M}}} \left[ \sup_{i \geq 0} \gamma^i \mathbb{1} \{s_i \in T\} \cdot \mathbb{1} \{\forall j \leq i, s_j \notin B\} \right] \\ &= \mathbb{E}_{\rho \sim \text{Pr}_{\pi}^{\mathcal{M}^{\circ B}}} \left[ \sup_{i \geq 0} \gamma^i \mathbb{1} \{s_i \in T\} \right]. \end{aligned}$$

Second, define

$$\text{Paths}_{\diamond T}^{\text{fin}} = \{\rho = s_0, s_1, \dots, s_i \mid s_i \in T \text{ and } s_j \notin T \text{ for all } j < t\}$$

as the set of finite paths that end up in  $T$ , with  $T$  being visited for the first time. Then, one can get rid of the supremum of Eq. (8) follows:

$$\begin{aligned} V_{\mathbf{I}}^{\pi}(\mathbb{O}) &= \mathbb{E}_{\rho \sim \text{Pr}_{\pi}^{\mathcal{M}^{\circ B}}} \left[ \sup_{i \geq 0} \gamma^i \mathbb{1} \{s_i \in T\} \right] \\ &= \mathbb{E}_{\rho \sim \text{Pr}_{\pi}^{\mathcal{M}^{\circ B}}} \left[ \sum_{t=0}^{\infty} \gamma^t \cdot \mathbb{1} \left\{ \text{pref}(\rho, t) \in \text{Paths}_{\diamond T}^{\text{fin}} \right\} \right], \quad (9) \end{aligned}$$

where  $\text{pref}(\rho, t) = s_0, s_1, \dots, s_t$  yields the prefix of  $\rho = s_0, s_1, \dots$  which ends up in the  $t^{\text{th}}$  state  $s_t$ . The attentive reader may have noticed that the resulting expectation can be seen as the expectation of a discounted cumulative reward signal (or a *discounted return*, for short), where a reward of one is incurred when visiting  $T$  for the first time. Taking it a step further, define the reward function

$$\text{rew}(s, a, s') = \begin{cases} 1 - \gamma & \text{if } s \in T, \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

Then, the value function can be re-written as

$$\begin{aligned} V_{\mathbf{I}}^{\pi}(\mathbb{O}) &= \mathbb{E}_{\rho \sim \text{Pr}_{\pi}^{\mathcal{M}^{\circ B}}} \left[ \sum_{t=0}^{\infty} \gamma^t \cdot \mathbb{1} \left\{ \text{pref}(\rho, t) \in \text{Paths}_{\diamond T}^{\text{fin}} \right\} \right] \\ &= \mathbb{E}_{\rho \sim \text{Pr}_{\pi}^{\mathcal{M}^{\circ T \cup B}}} \left[ \sum_{t=0}^{\infty} \gamma^t \cdot r_t \right]. \end{aligned}$$

For any state  $s \in T$ , notice that since  $T$  is absorbing in  $\mathcal{M}^{\circ T \cup B}$ ,

$$V^{\pi}(s, \mathbb{O}) = 1. \quad (10)$$

<sup>3</sup>cf. Lemma 3.

It is folklore that the discounted return is the solution of the Bellman equation  $V^\pi(s, \mathbb{O}) = \gamma \mathbb{E}_{a \sim \pi(\cdot|s)} \mathbb{E}_{s' \sim \mathbf{P}(\cdot|s,a)} [\text{rew}(s, a, s') \cdot V^\pi(s', \mathbb{O})]$  for any  $s \in \mathcal{S}$  (Puterman 1994). In particular, considering the reach-avoid objective  $\mathbb{O}$ , we have by Eq. (10)

$$V^\pi(s, \mathbb{O}) = \begin{cases} \gamma \mathbb{E}_{a \sim \pi(\cdot|s)} \mathbb{E}_{s' \sim \mathbf{P}(\cdot|s,a)} [V^\pi(s', \mathbb{O})] & \text{if } s \notin T \cup B, \\ 1 & \text{if } s \in T \setminus B, \text{ and} \\ 0 & \text{otherwise, when } s \in B. \end{cases}$$

Now, let us consider the values of the MDP plan  $\mathcal{M}_\Pi$  for the reach-avoid objective  $\mathbb{O}(\mathbf{T}, \mathbf{B})$  where  $\mathbf{T} = \{\langle s, v, u \rangle \mid v \in T\}$  and  $\mathbf{B} = \{\langle s, v, u \rangle \mid s \notin B_{\ell(v)}\}$  for the high-level objective  $\diamond T$  and set of low-level objectives  $\{\mathbb{O}_R^d : R \in \mathcal{R}, d \in D_R\}$  so that  $B_R$  is the set of states to avoid in room  $R$ . Fix a  $|\mathcal{V}|$ -memory two-level controller  $\pi = \langle \tau, \Pi \rangle$  in for two-level model  $\mathcal{H}$  (which is compliant with  $\mathcal{M}_\Pi$ , see Thm. 3 and the related proof). We take a close look to the value of each state in  $\mathcal{M}_\Pi$  by following the same structure as we used for the definition of  $\mathcal{M}_\Pi$  (cf. Sect. 5). For the sake of presentation, given any pair of vertices  $v, u \in \mathcal{V}$ , we may note  $\langle s_{\text{reset}}, v, u \rangle$  to refer to the (unified, cf. Assumption 3) reset state  $s_{\text{reset}} \in \mathcal{S}_\Pi$ . Given a state  $\langle s, v, u \rangle \in \mathcal{S}_\Pi$  with direction  $d = \langle v, u \rangle$ ,

(i) if  $s$  is not an exit state, i.e., if  $s \notin \mathcal{O}_{\ell(v)}(d)$ , then

$$\begin{aligned} & V^\pi(\langle s, v, u \rangle, \mathbb{O}) \\ &= \gamma \mathbb{E}_{\langle s', v, u \rangle \sim \mathbf{P}_\Pi(\cdot | \langle s, v, u \rangle, *)} [V^\pi(\langle s', v, u \rangle, \mathbb{O})] && \text{(by Eq. (5))} \\ &= \gamma \sum_{s' \in \mathcal{S}_{\ell(v)}} \mathbf{P}_\Pi(\langle s', v, u \rangle \mid \langle s, v, u \rangle, *) \cdot V^\pi(\langle s', v, u \rangle, \mathbb{O}) \\ &= \gamma \sum_{s' \in \mathcal{S}_{\ell(v)}} \sum_{a \in \mathcal{A}_{\ell(v)}} \bar{\pi}_{\ell(v), d}(a \mid s) \cdot \mathbf{P}_{\ell(v)}(s' \mid s, a) \cdot V^\pi(\langle s', v, u \rangle, \mathbb{O}); \end{aligned}$$

(ii) if  $s$  is an exit state in the direction  $d$ , i.e.,  $s \in \mathcal{O}_{\ell(v)}(d)$ , given the direction chosen by the planner  $d' = \tau(v, u) = \langle u, t \rangle$  for some neighbor  $t \in N(u)$ , we have

$$\begin{aligned} & V^\pi(\langle s, v, u \rangle, \mathbb{O}) \\ &= \gamma \mathbb{E}_{\langle s', u, t \rangle \sim \mathbf{P}_\Pi(\cdot | \langle s, v, u \rangle, d')} [V^\pi(\langle s', u, t \rangle, \mathbb{O})] \\ &= \gamma \mathbb{E}_{s' \sim \mathcal{I}_{\ell(u)}(\cdot | d)} [V^\pi(\langle s', u, t \rangle, \mathbb{O})] && \text{(by Eq. (6))} \\ &= \gamma \sum_{s' \in \mathcal{S}_{\ell(u)}} \mathcal{I}_{\ell(u)}(s' \mid d) \cdot V^\pi(\langle s', u, t \rangle, \mathbb{O}); \end{aligned} \tag{11}$$

(iii) if  $v$  is the target, i.e.,  $v \in T$ ,  $V^\pi(s, \mathbb{O}) = 1$ ; and

(iv) otherwise, when  $s$  is a bad state, i.e.,  $s \in B_{\ell(v)}$ ,  $V^\pi(s, \mathbb{O}) = 0$ .

Take  $R = \ell(v)$ . By (i) and (ii), when  $s$  is not an exit state, i.e.,  $s \notin \mathcal{O}_{\ell(v)}(d)$ , we have

$$V^\pi(\langle s, v, u \rangle, \mathbb{O}) = \sum_{s_0, s_1, \dots, s_i \in \text{Path}_{\mathbb{O}_R^d}^{\text{fin}}} \gamma^i \Pr_{\bar{\pi}_{R,d}}^{R_s}(s_0, s_1, \dots, s_i) \cdot V^\pi(\langle s_i, v, u \rangle, \mathbb{O}),$$

so that

$$\text{Path}_{\mathbb{O}_R^d}^{\text{fin}} = \text{Path}_{\diamond \mathbb{O}_R(d)}^{\text{fin}} \setminus \{\rho = s_0, s_1, \dots, s_n \mid \exists 1 \leq i \leq n, s_i \in B_R\},$$

where we denote by  $R_s$  the room  $R$  where we change the initial distribution by the Dirac  $\mathbf{I}_R(s_0) = \mathbb{1}\{s_0 = s\}$ , and  $\Pr_{\bar{\pi}_{R,d}}^{R_s}$  is the distribution over paths of  $R$  which start in state  $s$  which is induced by the choices of the low-level latent policy  $\bar{\pi}_{R,d}$ .

Following Eq. (11), notice that  $V^\pi(\langle s_{\text{exit}}, v, u \rangle, \mathbb{O}) = V^\pi(\langle s'_{\text{exit}}, v, u \rangle, \mathbb{O})$  for any  $s_{\text{exit}}, s'_{\text{exit}} \in \mathcal{O}_R(d = \langle v, u \rangle)$ : the probability of going to the next room  $R' = \ell(u)$  from an exit state of the current room  $R$  only depends on the entrance function  $\mathcal{I}_{R'}$  and is independent from the exact exit state which allowed to leave the current room  $R$ . Therefore, we further denote by  $V^\pi(\langle \cdot, v, u \rangle, \mathbb{O})$  the value of any exit state of  $R$  in direction  $d$ , i.e.,  $V^\pi(\langle \cdot, v, u \rangle, \mathbb{O}) = V^\pi(\langle s_{\text{exit}}, v, u \rangle, \mathbb{O})$  for all  $s_{\text{exit}} \in \mathcal{O}_R(d)$ .



Then, we have

$$\begin{aligned}
& V^\pi(\langle s, v, u \rangle, \mathbb{O}) \\
&= \sum_{s_0, s_1, \dots, s_i \in \text{Path}_{\mathbb{O}_R^d}^{\text{fin}}} \gamma^i \Pr_{\bar{\pi}, d}^{R, s} (s_0, s_1, \dots, s_i) \cdot V^\pi(\langle s_i, v, u \rangle, \mathbb{O}) \\
&= \sum_{s_0, s_1, \dots, s_i \in \text{Path}_{\mathbb{O}_R^d}^{\text{fin}}} \gamma^i \Pr_{\bar{\pi}, d}^{R, s} (s_0, s_1, \dots, s_i) \cdot V^\pi(\langle \cdot, v, u \rangle, \mathbb{O}) \\
&= V^\pi(\langle \cdot, v, u \rangle, \mathbb{O}) \cdot \sum_{s_0, s_1, \dots, s_i \in \text{Path}_{\mathbb{O}_R^d}^{\text{fin}}} \gamma^i \Pr_{\bar{\pi}, d}^{R, s} (s_0, s_1, \dots, s_i) \\
&= V^\pi(\langle \cdot, v, u \rangle, \mathbb{O}) \cdot V^{\bar{\pi}, d, R}(s, \gamma), \tag{by Eq. (9)}
\end{aligned}$$

where  $V^{\bar{\pi}, d}(s, \gamma)$  denotes the value of the reach-avoid objective  $\mathbb{O}_R^d = \mathbb{O}(\mathcal{O}_R(d), B_R)$  in the room  $R$  from state  $s \in R$ . Then, by (ii), assuming  $v \notin G$ , we have

1. if  $u \notin G$ ,

$$\begin{aligned}
& V^\pi(\langle \cdot, v, u \rangle, \mathbb{O}) \\
&= \gamma \cdot \sum_{s' \in \mathcal{S}_{\ell(u)}} \mathcal{I}_{\ell(u)}(s' \mid d = \langle v, u \rangle) \cdot V^\pi(\langle s', \tau(v, u) \rangle, \mathbb{O}) \tag{12} \\
&= \gamma \cdot \sum_{s' \in \mathcal{S}_{\ell(u)}} \mathcal{I}_{\ell(u)}(s' \mid d = \langle v, u \rangle) \cdot V^{\bar{\pi}, \ell(u), \tau(v, u)}(s, \mathbb{O}_R^d) \cdot V^\pi(\langle \cdot, \tau(v \cdot u) \rangle, \mathbb{O}) \\
&= \gamma \cdot \mathbf{P}(\tau(v, u) \mid \langle v, u \rangle, \tau(v, u)) \cdot V^\pi(\langle \cdot, \tau(v \cdot u) \rangle, \mathbb{O}) \quad (\text{where } \mathbf{P} \text{ is the transition function of } \mathcal{M}_{\Pi}^G, \text{ see Eq. (2)})
\end{aligned}$$

2. if  $u \in G$ ,

$$\begin{aligned}
& V^\pi(\langle \cdot, v, u \rangle, \mathbb{O}) \\
&= \gamma \cdot \sum_{s' \in \mathcal{S}_{\ell(u)}} \mathcal{I}_{\ell(u)}(s' \mid d = \langle v, u \rangle) \cdot V^\pi(\langle s', \tau(v, u) \rangle, \mathbb{O}) \\
&= \gamma \cdot \sum_{s' \in \mathcal{S}_{\ell(u)}} \mathcal{I}_{\ell(u)}(s' \mid d = \langle v, u \rangle) \cdot 1 \quad (\text{since } \tau(v, u) = \langle u, t \rangle \text{ for some } t \in N(u)) \\
&= \gamma.
\end{aligned}$$

Now, respectively denote by  $V_{\mathcal{M}_{\Pi}}^\pi(\cdot, \mathbb{O}) := V^\pi(\cdot, \mathbb{O})$  and  $V_{\mathcal{M}_{\Pi}^G}^\tau(\cdot, \diamond T)$  the value functions of  $\mathcal{M}_{\Pi}$  and  $\mathcal{M}_{\Pi}^G$  for the objectives  $\mathbb{O}$  and  $\diamond T$ . By 1 and 2, and by construction of  $\mathcal{M}_{\Pi}^G$ , we have for any pair of vertices  $v, u \in \mathcal{V}$  that

$$V_{\mathcal{M}_{\Pi}}^\pi(\langle \cdot, v, u \rangle, \mathbb{O}) = \gamma \cdot V_{\mathcal{M}_{\Pi}^G}^\tau(\langle v, u \rangle, \diamond T),$$

On the one hand, notice that, by construction of  $\mathcal{M}_{\Pi}^G$ , we have for any pair of vertices  $\langle v, u \rangle \in E$  that the initial values  $V_{\mathcal{M}_{\Pi}^G}^\tau(\diamond T)$  are  $\mathbb{E}_{d \sim \mathbf{I}} V_{\mathcal{M}_{\Pi}^G}^\tau(d, \diamond T) = V_{\mathcal{M}_{\Pi}^G}^\tau(d_0, \diamond T)$ . On the other hand, we have

$$V_{\mathcal{M}_{\Pi}}^\pi(\mathbb{O}) = \mathbb{E}_{s' \sim \mathcal{I}_{\ell(v_0)}(\cdot \mid d_0)} V_{\mathcal{M}_{\Pi}}^\pi(\langle s', \tau(d_0) \rangle, \mathbb{O}) = 1/\gamma \cdot V_{\mathcal{M}_{\Pi}}^\pi(\cdot, d_0, \mathbb{O}). \tag{by Eq. (12)}$$

Then, we finally have:

$$V_{\mathcal{M}_{\Pi}}^\pi(\mathbb{O}) = 1/\gamma \cdot V_{\mathcal{M}_{\Pi}}^\pi(\cdot, d_0, \mathbb{O}) = \gamma/\gamma \cdot V_{\mathcal{M}_{\Pi}^G}^\tau(d_0, \diamond T) = V_{\mathcal{M}_{\Pi}^G}^\tau(\diamond T),$$

which concludes the proof.  $\square$

## G Initial Distribution Shifts: Training vs. Synthesis

Our two-level controller construction occurs in two phases. First, we create a set of low-level policies  $\Pi$  by running Algorithm 1 in each room (Sect. 4.3). Notably, training in each room is independent and can be executed in parallel. However, independent training introduces a challenge: an *initial distribution shift* emerges when combining low-level policies using a planner. Our value bounds for a room  $R$  in direction  $d$  depend on a loss  $L_{\mathbf{P}}^{R, d}$ , computed based on the stationary distribution. This distribution may significantly change depending on a planner's choices. In this section, we address this challenge by showing, under mild

assumptions on the initial distribution of each room  $R$ , that their transition losses  $L_{\mathbf{P}}^{R,d}$  obtained under any latent policy  $\bar{\pi}_{R,d}$  for direction  $d$  still guarantee to bound the gap between the values of the original and latent two-level models.

In the following, we first give details on this ‘‘distribution shift,’’ and then we prove Theorem 5 through Theorems 8 and 9.

**Training rooms.** To construct  $\Pi$ , we train low-level policies via Algorithm 1 by simulating each room individually. Precisely, for room  $R \in \mathcal{R}$  and direction  $d \in D_R$ , we train a WAE-DQN agent by considering  $R$  as episodic MDP with *some* initial distribution  $\mathbf{I}_R$ , yielding (i) low-level latent policy  $\bar{\pi}_{R,d}$ , (ii) latent MDP  $\mathcal{M}_R$ , and (iii) state-embedding function  $\phi_R$ . Since  $\bar{\pi}_{R,d}$  must learn to maximize the values of the objective  $\mathbb{O}_R^d$ , which asks to reach the exit state in direction  $d$ , we restart the simulation when the latter is visited. Formally, the related training room is an episodic MDP  $R_d = \langle \mathcal{S}_R, \mathcal{A}_R, \mathbf{P}_R^d, \mathbf{I}_R \rangle$ , where  $s_{\text{reset}} \in \mathcal{S}_R$ ,  $\mathbf{P}_R^d(\cdot | s, a) = \mathbf{P}_R(\cdot | s, a)$  when  $s \notin \mathcal{O}_R(d)$ , and  $\mathbf{P}_R^d(s_{\text{reset}} | s, a) = 1$  otherwise. We define  $\bar{\mathbf{P}}_R^d$  similarly for  $\bar{\mathcal{M}}_R$  when the direction  $d$  is considered.

**Distribution shift.** Crucially, by considering rooms individually, a noticeable *initial distribution shift* occurs when switching between training and synthesis phases. During training, there is no two-level controller, so the initial distribution of room  $R$  is just  $\mathbf{I}_R$ . During synthesis, room entries and exits are determined by the distributions influenced by the choices made by the controller in the hierarchical MDP  $\mathcal{H}$ . This implies that the induced initial distribution of each room depends on the likelihood of visiting other rooms and is further influenced by the other low-level policies.

We contend that this shift may induce significant consequences: denote by  $L_{\mathbf{P}}^{R,d}$  the transition loss of the room  $R_d$  operating under  $\bar{\pi}_{R,d}$  and by  $L_{\mathbf{P}}^{\tau, \Pi}$  the transition loss of the two-level model  $\mathcal{H}$  operating under  $\langle \tau, \Pi \rangle$ . Then, in the worst case,  $L_{\mathbf{P}}^{\tau, \Pi}$  and  $L_{\mathbf{P}}^{R,d}$  might be completely unrelated whatever the room  $R$  and direction  $d$ . To see why, recall that transition losses are defined over stationary distributions of the respective models (Eq. 1). One can see this shift as a perturbation in the transition function of the rooms. Intuitively, by Assumption 1, each room is almost surely entered infinitely often, meaning that such perturbations are also repeated infinitely often, possibly leading to completely divergent stationary distributions (O’Cinneide 1993), meaning that we lose the abstraction quality guarantees possibly obtained for each individual training room.

**Entrance loss.** Fortunately, we claim that under some assumptions, when the initial distribution of each training room  $\mathbf{I}_R$  is wisely chosen, we can still link the transition losses  $L_{\mathbf{P}}^{R,d}$  minimized in the training rooms to  $L_{\mathbf{P}}^{\tau, \Pi}$ . To provide this guarantee, the sole remaining missing component to our framework is learning a *latent entrance function*: we define the entrance loss as

$$L_{\mathcal{I}} = \mathbb{E}_{R,d \sim \xi_{\pi}} \mathcal{D}(\phi_{\mathcal{I}R}(\cdot | d), \bar{\mathcal{I}}_R(\cdot | d)), \quad (13)$$

where  $\phi_{\mathcal{I}R}(\cdot | d) = \mathbb{E}_{s \sim \mathcal{I}R(\cdot | d)} \mathbb{1}\{s = \phi_R(s)\}$ ,  $\bar{\mathcal{I}}_R: D_R \rightarrow \Delta(\bar{\mathcal{S}})$  is the latent entrance function,  $\pi$  is the stationary policy in  $\mathcal{M}_{\Pi}$  corresponding to the two-level controller  $\langle \tau, \Pi \rangle$  where  $\tau$  has a memory of size  $|\mathcal{V}|$ ,  $\mathcal{D}$  is total variation, and  $\xi_{\pi}$  is the stationary distribution induced by  $\pi$  in  $\mathcal{M}_{\Pi}$ . The measure  $\xi_{\pi}$  can also be seen as a distribution over rooms and directions chosen under the controller:<sup>4</sup>

$$\xi_{\pi}(R, d) = \mathbb{E}_{s,v,u \sim \xi_{\pi}} [\mathbb{1}\{s = s_{\text{reset}}, R = \ell(v), d = d_0\} + \mathbb{1}\{R = \ell(v), d = \langle v, u \rangle\}].$$

**Theorem 8** (Reusable RL components). Let  $\langle \tau, \Pi \rangle$  be a two-level controller in  $\mathcal{H}$  where  $\tau$  has finite memory of size  $|\mathcal{V}|$  and let  $\pi$  be the equivalent stationary policy in the MDP plan  $\mathcal{M}_{\Pi}$ . Assume (i)  $\Pi$  only consists of latent policies and (ii) for any training room  $R \in \mathcal{R}$  and direction  $d \in D_R$ , the projection<sup>5</sup> of the BSCC of  $\mathcal{M}_{\Pi}$  under  $\pi$  to  $\mathcal{S}_R$  is included in the BSCC of  $R_d$  under low-level policy  $\bar{\pi}_{R,d}$ . Let

$$\begin{aligned} \mathcal{S}_{R,d} &= \{\langle s, v, u \rangle \in \mathcal{S}_{\Pi} \mid \ell(v) = R \text{ and } \langle v, u \rangle = d\}, \\ \xi_{\pi}(s_{\text{reset}} \mid R, d) &= \mathbb{E}_{\langle s,v,u \rangle, a \sim \xi_{\pi}} [\mathbf{P}_{\Pi}(s_{\text{reset}} \mid \langle s, v, u \rangle, a) \mid \mathcal{S}_{R,d}], \text{ and} \\ \xi_{\text{continue}}^{\min} &= 1 - \max_{R \in \mathcal{R}, d \in D} (\xi_{\pi}(s_{\text{reset}} \mid \mathcal{S}_{R,d}) + \xi_{\pi}(\mathcal{O}_R(d) \times \{d\} \mid \mathcal{S}_{R,d})). \end{aligned}$$

Then, there is a  $\kappa \geq 0$  with  $L_{\mathbf{P}}^{\tau, \Pi} \leq L_{\mathcal{I}} + \frac{\kappa}{\xi_{\text{continue}}^{\min}} \mathbb{E}_{R,d \sim \xi_{\pi}} L_{\mathbf{P}}^{R,d}$ . Define the expected entrance function in room  $R$  as

$$\mathbf{I}_R^{\pi}(s) = \mathbb{E}_{\dot{s}, \langle u, v \rangle \sim \xi_{\pi}} [\mathcal{I}_R(s \mid d = \langle u, v \rangle) \mid \dot{s} \in \mathcal{O}_{\ell(u)}(\langle u, v \rangle) \text{ and } \ell(v) = R] \text{ for any } s \in \mathcal{S}_R.$$

With  $\text{supp}(P) = \{x \in \mathcal{X} \mid P(x) > 0\}$  the support of distribution  $P$ , if  $\text{supp}(\mathbf{I}_R) = \text{supp}(\mathbf{I}_R^{\pi})$ ,  $\kappa$  can be set to the maximum probability ratio of room entry during training and synthesis:

$$\kappa = \max_{R \in \mathcal{R}} \left( \max_{s \in \text{supp}(\mathbf{I}_R)} \max \left\{ \frac{\mathbf{I}_R^{\pi}(s)}{\mathbf{I}_R(s)}, \frac{\mathbf{I}_R(s)}{\mathbf{I}_R^{\pi}(s)} \right\} \right)^{|\mathcal{S}|}.$$

<sup>4</sup>For simplicity, we consider here the special state  $\langle s_{\text{reset}}, v, v_0 \rangle$  with  $\langle v, v_0 \rangle = d_0$  as the joint reset state of the model (Assumption 3).

<sup>5</sup>Formally speaking, this is the projection to  $\mathcal{S}_R$  of the intersection of the BSCC of  $\mathcal{M}_{\Pi}$  operating under  $\pi$  with  $\mathcal{S}_R \times D_R$ .






Figure 9: Room  $R = \ell(v)$  in the two-level model (left) and the same room taken individually (right). Both distributions  $\xi_\pi(\cdot | d)$  and  $\xi_{\bar{\pi}_{R,d}}^{R_{d,\pi}}$  correspond to the limiting distributions over  $S_R$  when  $\bar{\pi}_{R,d}$  is executed in  $R$ . The sole difference remains in the fact that the reset is considered outside  $R$  in the two-level model (Assumption 3) while it is considered to be part of the state space when  $R$  is taken individually (Assumption 1).

We call the resulting MDP the *individual room* version of  $R$  that we denote by  $R_{d,\pi}$ . The stationary distribution of the room  $R_{d,\pi}$  for the low-level policy  $\bar{\pi}_{R,d}$  is  $\xi_{\bar{\pi}_{R,d}}^{R_{d,\pi}}$ . Observe that  $\xi_{\bar{\pi}_{R,d}}^{R_{d,\pi}}$  is over  $S_R$ , which includes the reset state  $s_{\text{reset}}$ , while  $\xi_\pi(\cdot | d)$  is over the exact same state space but without the reset state (since the reset state is a special state outside  $R$ , shared by all the rooms in the two-level model; cf. Assumption 3 and the definition of  $\mathcal{M}_\Pi$ ). Furthermore, notice that, modulo this reset state, the two distributions are the same (see Fig. 9): they both consist of the limiting distribution over  $S_R$  when  $\bar{\pi}_{R,d}$  is executed in  $R$ . All the transition distributions remain the same, except those of the exit states: in the two-level model  $\mathcal{H}$ , every state  $s \in \mathcal{O}_R(d = \langle v, u \rangle)$  transitions to  $u$  deterministically, while in the individual room  $R_{d,\pi}$ , they transition to the reset state deterministically. Still, in both cases,  $R$  is entered and exited with the same probability (respectively from and to  $(S \setminus S_R \times \{v\})$  in  $\mathcal{H}$  and  $s_{\text{reset}}$  in the individual room  $R_{d,\pi}$ ). Therefore, we have:

$$\xi_\pi(s | d) = \xi_{\bar{\pi}_{R,d}}^{R_{d,\pi}}(s | S_R \setminus \{s_{\text{reset}}\}) = \frac{\xi_{\bar{\pi}_{R,d}}^{R_{d,\pi}}(s) \cdot \mathbb{1}\{s \neq s_{\text{reset}}\}}{1 - \xi_{\bar{\pi}_{R,d}}^{R_{d,\pi}}(s_{\text{reset}})}. \quad (15)$$

Instead of sampling from  $\xi_\pi(s | d)$  in Eq. (\*), we would rather like to sample from the distribution of the individual room  $\xi_{\bar{\pi}_{R,d}}^{R_{d,\pi}}(s | S_R \setminus \{s_{\text{reset}}\})$ . We have:

$$\begin{aligned} & \mathbb{E}_{s,a \sim \xi_\pi(\cdot | d)} [\mathbb{1}\{s \neq s_{\text{reset}}\} \mathbb{1}\{s \notin \mathcal{O}_{\ell(v)}(d)\} \|\phi \mathbf{P}(\cdot | s, a) - \bar{\mathbf{P}}(\cdot | \phi(s), a)\|_1] \\ &= \sum_{s \in S} \sum_{a \in \mathcal{A}} [\xi_\pi(s | d) \bar{\pi}_{R,d}(a | \phi(s)) \mathbb{1}\{s \neq s_{\text{reset}}\} \mathbb{1}\{s \notin \mathcal{O}_{\ell(v)}(d)\} \|\phi \mathbf{P}(\cdot | s, a) - \bar{\mathbf{P}}(\cdot | \phi(s), a)\|_1] \end{aligned} \quad (16)$$

$$= \sum_{s \in S} \sum_{a \in \mathcal{A}} \left[ \frac{\xi_{\bar{\pi}_{R,d}}^{R_{d,\pi}}(s) \mathbb{1}\{s \neq s_{\text{reset}}\}}{1 - \xi_{\bar{\pi}_{R,d}}^{R_{d,\pi}}(s_{\text{reset}})} \bar{\pi}_{R,d}(a | \phi(s)) \mathbb{1}\{s \neq s_{\text{reset}}\} \mathbb{1}\{s \notin \mathcal{O}_{\ell(v)}(d)\} \|\phi \mathbf{P}_R(\cdot | s, a) - \bar{\mathbf{P}}_R(\cdot | \phi(s), a)\|_1 \right] \quad (17)$$

$$= \mathbb{E}_{s,a \sim \xi_{\bar{\pi}_{R,d}}^{R_{d,\pi}}} \left[ \frac{\mathbb{1}\{s \neq s_{\text{reset}}\}}{1 - \xi_{\bar{\pi}_{R,d}}^{R_{d,\pi}}(s_{\text{reset}})} \mathbb{1}\{s \notin \mathcal{O}_{\ell(v)}(d)\} \|\phi \mathbf{P}_R(\cdot | s, a) - \bar{\mathbf{P}}_R(\cdot | \phi(s), a)\|_1 \right]$$

Notice that we can pass from Eq. (16) to (17) because we only consider states  $s \neq s_{\text{reset}}$  and  $s \notin \mathcal{O}_{\ell(v)}(d)$ . States that do not satisfy both constraints are the only ones for which  $\mathbf{P}(\cdot | s, a)$  differs from  $\mathbf{P}_R^{d,\pi}(\cdot | s, a)$  (Eq. 14). Furthermore, in that case,

we have  $\mathbf{P}_R^{d,\pi}(\cdot | s, a) = \mathbf{P}_R(\cdot | s, a)$ . Then we have:

$$\begin{aligned}
& \mathbb{E}_{s,a \sim \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}} \left[ \frac{\mathbf{1}\{s \neq s_{\text{reset}}\}}{1 - \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s_{\text{reset}})} \mathbf{1}\{s \notin \mathcal{O}_{\ell(v)}(d)\} \|\phi \mathbf{P}_R(\cdot | s, a) - \bar{\mathbf{P}}_R(\cdot | \phi(s), a)\|_1 \right] \\
&= \mathbb{E}_{s,a \sim \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}} \left[ \frac{\mathbf{1}\{s \neq s_{\text{reset}}\}}{1 - \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s_{\text{reset}})} \|\phi \mathbf{P}_R^d(\cdot | s, a) - \bar{\mathbf{P}}_R^d(\cdot | \phi(s), a)\|_1 \right] \quad (\text{by definition of } \mathbf{P}_R^d \text{ and } \bar{\mathbf{P}}_R^d) \\
&= \frac{1}{1 - \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s_{\text{reset}})} \mathbb{E}_{s,a \sim \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}} \left[ \mathbf{1}\{s \neq s_{\text{reset}}\} \|\phi \mathbf{P}_R^d(\cdot | s, a) - \bar{\mathbf{P}}_R^d(\cdot | \phi(s), a)\|_1 \right] \\
&\leq \frac{1}{1 - \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s_{\text{reset}})} \mathbb{E}_{s,a \sim \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}} \|\phi \mathbf{P}_R^d(\cdot | s, a) - \bar{\mathbf{P}}_R^d(\cdot | \phi(s), a)\|_1
\end{aligned}$$

Assuming that the projection of the BSCC of  $\mathcal{M}_\Pi$  under  $\pi$  to  $\mathcal{S}_R$  is included in the BSCC of  $R$  when it operates under  $\bar{\pi}_{R,d}$ , we have that  $\text{supp}\left(\xi_{\bar{\pi}_{R,d}}^{R_d,\pi}\right) \subseteq \text{supp}\left(\xi_{\bar{\pi}_{R,d}}^R\right)$ , where  $\xi_{\bar{\pi}_{R,d}}^R$  denotes the stationary distribution of the training room  $R_d$  under the latent policy  $\bar{\pi}_{R,d}$ . Then:

$$\begin{aligned}
& \frac{1}{1 - \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s_{\text{reset}})} \mathbb{E}_{s,a \sim \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}} \|\phi \mathbf{P}_R^d(\cdot | s, a) - \bar{\mathbf{P}}_R^d(\cdot | \phi(s), a)\|_1 \\
&= \frac{1}{1 - \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s_{\text{reset}})} \sum_{s \in \text{supp}\left(\xi_{\bar{\pi}_{R,d}}^{R_d,\pi}\right)} \sum_{a \in \mathcal{A}_R} \left[ \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s) \bar{\pi}_{R,d}(a | \phi(s)) \|\phi \mathbf{P}_R^d(\cdot | s, a) - \bar{\mathbf{P}}_R^d(\cdot | \phi(s), a)\|_1 \right] \\
&= \frac{1}{1 - \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s_{\text{reset}})} \sum_{s \in \text{supp}\left(\xi_{\bar{\pi}_{R,d}}^R\right)} \sum_{a \in \mathcal{A}_R} \left[ \frac{\xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s)}{\xi_{\bar{\pi}_{R,d}}^R(s)} \xi_{\bar{\pi}_{R,d}}^R(s) \bar{\pi}_{R,d}(a | \phi(s)) \|\phi \mathbf{P}_R^d(\cdot | s, a) - \bar{\mathbf{P}}_R^d(\cdot | \phi(s), a)\|_1 \right] \\
&= \frac{1}{1 - \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s_{\text{reset}})} \mathbb{E}_{s,a \sim \xi_{\bar{\pi}_{R,d}}^R} \left[ \frac{\xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s)}{\xi_{\bar{\pi}_{R,d}}^R(s)} \|\phi \mathbf{P}_R^d(\cdot | s, a) - \bar{\mathbf{P}}_R^d(\cdot | \phi(s), a)\|_1 \right] \\
&\leq \frac{1}{1 - \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s_{\text{reset}})} \mathbb{E}_{s,a \sim \xi_{\bar{\pi}_{R,d}}^R} \left[ \max_{s' \in \text{supp}\left(\xi_{\bar{\pi}_{R,d}}^R\right)} \left( \frac{\xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s')}{\xi_{\bar{\pi}_{R,d}}^R(s')} \right) \|\phi \mathbf{P}_R^d(\cdot | s, a) - \bar{\mathbf{P}}_R^d(\cdot | \phi(s), a)\|_1 \right] \\
&= \frac{1}{1 - \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s_{\text{reset}})} \max_{s \in \text{supp}\left(\xi_{\bar{\pi}_{R,d}}^R\right)} \left( \frac{\xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s)}{\xi_{\bar{\pi}_{R,d}}^R(s)} \right) \mathbb{E}_{s,a \sim \xi_{\bar{\pi}_{R,d}}^R} \left[ \|\phi \mathbf{P}_R^d(\cdot | s, a) - \bar{\mathbf{P}}_R^d(\cdot | \phi(s), a)\|_1 \right] \\
&= \max_{s \in \text{supp}\left(\xi_{\bar{\pi}_{R,d}}^R\right)} \left( \frac{\xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s)}{\xi_{\bar{\pi}_{R,d}}^R(s)} \right) \frac{2L_{\mathbf{P}}^{R,d}}{1 - \xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s_{\text{reset}})}
\end{aligned}$$

If the initial distributions of the individual room  $R_{d,\pi}$  and the training room  $R_d$  have the same support, then the projection and the BSCCs coincide since the same set of states is eventually visited under  $\bar{\pi}$  from states of  $\text{supp}(\mathbf{I}_R) = \text{supp}(\mathbf{I}_R^\pi)$ . Furthermore, by (O’Cinneide 1993, Thm. 1), we have

$$\max_{s \in \text{supp}\left(\xi_{\bar{\pi}_{R,d}}^R\right)} \left( \frac{\xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s)}{\xi_{\bar{\pi}_{R,d}}^R(s)} \right) \tag{18}$$

$$\begin{aligned}
&\leq \max_{s \in \text{supp}\left(\xi_{\bar{\pi}_{R,d}}^R\right)} \max \left\{ \frac{\xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s)}{\xi_{\bar{\pi}_{R,d}}^R(s)}, \frac{\xi_{\bar{\pi}_{R,d}}^R(s)}{\xi_{\bar{\pi}_{R,d}}^{R_d,\pi}(s)} \right\} \\
&\leq \left( \max_{s \in \text{supp}(\mathbf{I}_R)} \max \left\{ \frac{\mathbf{I}_R^\pi(s)}{\mathbf{I}_R(s)}, \frac{\mathbf{I}_R(s)}{\mathbf{I}_R^\pi(s)} \right\} \right)^{|\mathcal{S}|} \quad (\text{cf. Eq. (14)})
\end{aligned}$$

$$= \kappa_{R,d}; \tag{19}$$

otherwise, we set  $\kappa_{R,d}$  to  $\max_{s \in \text{supp}(\xi_{\bar{\pi}_{R,d}}^R)} \left( \frac{\xi_{\bar{\pi}_{R,d}}^{R,d}(s)}{\xi_{\bar{\pi}_{R,d}}^R(s)} \right)$ . Moreover, let  $\mathcal{S}_{R,d} = \{\langle s, v, u \rangle \in \mathcal{S}_{\Pi} \mid \ell(v) = R \text{ and } \langle v, u \rangle = d\}$  and define

$$\xi_{\pi}(s_{\text{reset}} \mid R, d) = \mathbb{E}_{\langle s, v, u \rangle, a \sim \xi_{\pi}} [\mathbf{P}_{\Pi}(s_{\text{reset}} \mid \langle s, v, u \rangle, a) \mid \mathcal{S}_{R,d}].$$

Notice that

$$\begin{aligned} \xi_{\bar{\pi}_{R,d}}^{R,d}(s_{\text{reset}}) &= \mathbb{E}_{\langle s, v, u \rangle, a \sim \xi_{\pi}} [\mathbf{P}_{\Pi}(s_{\text{reset}} \mid \langle s, v, u \rangle, a) + \mathbb{1}\{s \in \mathcal{O}_R(d)\} \mid \mathcal{S}_{R,d}] \\ &= \xi_{\pi}(s_{\text{reset}} \mid R, d) + \xi_{\pi}(\mathcal{O}_R(d) \times \{d\} \mid \mathcal{S}_{R,d}) \end{aligned}$$

by (i) the stationary property, (ii) definition of  $\mathbf{P}_R^{d,\pi}$  (cf. Eq. (14) and Fig. 9), (iii) the fact that the probability of exiting the room is equal to the probability of visiting an exit state, and (iv) the fact that resetting the room and visiting an exit state are disjoint events (when an exit state is visited, it always transitions to the next room, never to the reset state).

By putting all together, we have

$$\begin{aligned} &L_{\mathbf{P}}^{\tau, \Pi} \\ &\leq L_{\mathcal{I}} + \frac{1}{2} \mathbb{E}_{d \sim \xi_{\pi}} \mathbb{E}_{s, a \sim \xi_{\pi}(\cdot \mid d)} [\mathbb{1}\{s \neq s_{\text{reset}}\} \mathbb{1}\{s \notin \mathcal{O}_{\ell(v)}(d)\} \|\phi_{\mathbf{P}}(\cdot \mid s, a) - \bar{\mathbf{P}}(\cdot \mid \phi(s), a)\|_1] \\ &\leq L_{\mathcal{I}} + \mathbb{E}_{R, d \sim \xi_{\pi}} \frac{\kappa_{R,d} L_{\mathbf{P}}^{R,d}}{1 - \xi_{\bar{\pi}_{R,d}}^{R,d}(s_{\text{reset}})} \\ &= L_{\mathcal{I}} + \mathbb{E}_{R, d \sim \xi_{\pi}} \frac{\kappa_{R,d} L_{\mathbf{P}}^{R,d}}{1 - \xi_{\pi}(s_{\text{reset}} \mid R, d) - \xi_{\pi}(\mathcal{O}_R(d) \times \{d\} \mid \mathcal{S}_{R,d})} \\ &\leq L_{\mathcal{I}} + \mathbb{E}_{R, d \sim \xi_{\pi}} \frac{\max\{\kappa_{R^*, d^*} : R^* \in \mathcal{R}, d^* \in D_{R^*}\} L_{\mathbf{P}}^{R,d}}{1 - \max_{R^* \in \mathcal{R}, d^* \in D_{R^*}} (\xi_{\pi}(s_{\text{reset}} \mid R^*, d^*) + \xi_{\pi}(\mathcal{O}_{R^*}(d^*) \times \{d^*\} \mid \mathcal{S}_{R^*, d^*}))} \\ &\leq L_{\mathcal{I}} + \frac{\kappa}{\xi_{\text{continue}}^{\min}} \mathbb{E}_{R, d \sim \xi_{\pi}} L_{\mathbf{P}}^{R,d} \end{aligned}$$

where  $\kappa = \max\{\kappa_{R^*, d^*} : R^* \in \mathcal{R}, d^* \in D_{R^*}\}$  and

$$\xi_{\text{continue}}^{\min} = 1 - \max_{R \in \mathcal{R}, d \in D_R} (\xi_{\pi}(s_{\text{reset}} \mid R, d) + \xi_{\pi}(\mathcal{O}_R(d) \times \{d\} \mid \mathcal{S}_{R,d})). \quad (20)$$

This concludes the proof.  $\square$

**Discussion.** Assumption (ii) boils down to design an initial distribution for the simulator of each room that provides a sufficient coverage of the state space: the latter should include the states likely to be seen when the room is entered under any planner. Then, if this initial distribution is powerful enough to provide an exact coverage of the entrance states visited under the planner  $\tau$ , the multiplier of the transition loss  $\kappa$  can be determined solely based on the ratio of the initial distributions obtained during training and synthesis. We summarize the results as follows.

**Theorem 9** (Value bound in  $\mathcal{H}$ ). Under the assumptions of Thm. 8,

$$|V_{\mathbf{I}}^{\pi} - \bar{V}_{\mathbf{I}}^{\pi}| \leq \frac{L_{\mathcal{I}} + \kappa / \xi_{\text{continue}}^{\min} \mathbb{E}_{R, d \sim \xi_{\pi}} L_{\mathbf{P}}^{R,d}}{\xi_{\pi}(s_{\text{reset}})(1-\gamma)}. \quad (21)$$

## H Experiments

In this section, we provide additional details on the experiments we performed.

**Setup.** Models have been trained on a cluster running under CentOS Linux 7 (Core) composed of a mix of nodes containing Intel processors with the following CPU microarchitectures: (i) 10-core INTEL E5-2680v2, (ii) 14-core INTEL E5-2680v4, and (iii) 20-core INTEL Xeon Gold 6148. We used 8 cores and 42 GB of memory for each run during the hyperparameter search. **Learning the low-level policies.** We run WAE-DQN to learn the set of low-level policies  $\Pi$  along with their latent-space models. Recall the representation quality guarantees of our algorithm (cf. Sect. 4.3): the same latent space can be used for rooms sharing similar features. We leverage this property to learn only four latent policies (one per direction). In other words, only one pair of latent MDP and policy is learned per direction, which encompasses and generalizes the behavior of the agent in all the *training rooms* (cf. Appendix G). For instance, in a grid world environment composed of 9 rooms with similar shapes, we only train one latent policy per exit direction  $\{\leftarrow, \rightarrow, \uparrow, \downarrow\}$  instead of  $9 \cdot 4 = 36$ . For training in a room  $R$ , we let  $\mathbf{I}_R$  uniformly distribute the agent's possible entry positions. Adversaries' initial positions are randomly set




Figure 10: Two rooms of  $20 \times 20$  cells (9 rooms in Fig. 12). Demonstration with 9 rooms: <https://youtu.be/crowN8-GaRg>.

by  $\mathbf{I}_R$  but may vary according to the function  $\mathcal{I}_R$  in the high-level model (unknown at training time). Objectives  $\mathbb{O}_R^d$  specify reaching the target exit while avoiding adversaries before the episode ends.

**Synthesis.** To estimate the latent entrance function, we explore the high-level environment through random execution of the low-level latent policies. We further consider *Masked Autoencoders* (MADEs, (Germain et al. 2015)), which allow to learn complex distributions from a dataset. With the data collected via this exploration, we train a MADE to learn  $\bar{\mathcal{I}}_R$  for any room  $R$ . To learn those latent entrance functions, consistently with WAE-MDPs, we use the same kind of MADE as the one introduced by (Delgrange et al. 2023) for estimating the probability of the latent transition function. We finally construct  $\bar{\mathcal{M}}_{\Pi}^{\mathcal{G}}$  (cf. Sect. 5) and apply the synthesis procedure to obtain a two-level controller  $\pi = \langle \tau, \Pi \rangle$ . Tab. 1 reports the values of  $\pi$  obtained for various environment sizes.

**Hyperparameter search.** To train our WAE-DQN agent, we ran 4 environments in parallel per training instance and used a replay buffer of capacity  $7.5 \cdot 10^5$ . We performed a grid search to find the best parameters for our WAE-DQN algorithm. Tab. 3 presents the range of hyperparameters used. In particular, we found that prioritized experience replay does not improve the results in our environments significantly. We used a batch size of 128 for the WAE-MDP.

For the MADE modeling the latent entrance function, we used a dataset of size 25600, and the training was split into 100 epochs (i.e., the model performed 100 passes through the entire dataset) with a learning rate of  $10^{-3}$ . We used a batch size of 32 or 64, and two hidden layers, either with 64 or 128 neurons.

For generating the set of low-level policies  $\Pi$ , we used the hyperparameters that worked the best for each specific direction. We used the same parameters for the DQN training instances shown in Figure 6.

## H.1 Grid World Environment

We provide additional details on the state space of our environment. The agent has  $L_P$  life points, decrementing upon adversary contact or timeout. Collecting power-ups (appearing randomly) shortly makes the agent invincible. The state space comprises two components: (i) A 4-dimensional bitmap  $\mathbf{M} \in \{0, 1\}^{N \times l \times m \times n}$ , where each layer in  $k \in \{1, \dots, l\}$  corresponds to an item type on the grid; entry  $\mathbf{M}_{R,k,i,j}$  is 1 iff room  $R$  has item  $k$  in cell  $(i, j)$ ; (ii) step, power-up, and life-point counters  $\langle a, b, c \rangle$ . Figure 10 shows examples of rooms composed of  $20 \times 20$  cells.

**DRL components.** We use CNNs (LeCun et al. 1989) to process bitmaps  $\mathbf{M}$  and a sparse reward signal  $rew(s, a, s') = r^* \cdot \mathbb{1}\{s \in T\} - r^* \cdot \mathbb{1}\{s \in B\}$ , where  $r^* > 0$  is an arbitrary reward (or conversely, a penalty) obtained upon reaching the target  $T$  (or an undesirable state in  $B$ ). To guide the agent, we add a *potential-based reward shaping* (Ng et al. 1999; Wiewiora 2003) based on the  $L_1$  distance to the target. The resulting reward function is  $rew_{\Phi}(s, a, s') = \gamma \Phi(s') - \Phi(s) + rew(s, a, s')$  where

$$\Phi(s) = 1 - \frac{\min \{|x(t_1) - x(s)| + |y(t_2) - y(s)| : t_1, t_2 \in T\}}{N \cdot (m + n)}, \quad (22)$$

and  $x(s), y(s)$  respectively return the Euclidean coordinates along the horizontal and vertical axes corresponding to state  $s \in \mathcal{S}$ . Intuitively,  $|\Phi(s) - 1|$  reflects the normalized distance of state  $s$  to the targets  $T$ . When the agent gets closer (resp. further) to  $T$  when executing an action, the resulting reward is positive (resp. negative). Our DQN implementation uses state-of-the-art extensions and improvements from (Hessel et al. 2018). Nevertheless, as demonstrated in Fig. 6, while DQN reduces contact with adversaries, the two-level nature of the decisions required to reach a target hinders learning the high-level objective.

**DQN and WAE-DQN experiments.** We provide a more detailed version of Figure 6 in Figure 11, where the WAE-DQN performance is specified per direction. Precisely, we trained five different instances of the algorithm per policy with different






Figure 11: A more detailed version of Figure 6, where the WAE-DQN performance is specified per direction. We train five different instances of the algorithm per policy with different random seeds, where the solid line corresponds to the mean and the shaded interval to the standard deviation. To train the DQN agent, we set a time limit five times longer than that used for training rooms with the WAE-DQN agents. Note that the DQN agent is equipped with 3 life points, while the WAE-DQN agents are limited to one.

random seeds, where the solid line is the mean and the shaded interval is the standard deviation. To train the DQN agent, we set a time limit five times longer than that used for training rooms with the WAE-DQN agents. Furthermore, the DQN agent is equipped with 3 life points, while the WAE-DQN agents are limited to one.

## H.2 ViZDoom Environment

A top view of the environment and samples of visual inputs processed by the agent are provided in Fig. 13. Each game frame consists of an image of size  $60 \times 45$ . Note that we do not stack frames in the agent’s observation. Indeed, the agent additionally processes velocities, its angle in the map, and its health, which makes the observation Markovian. We used CNNs to process the game frames with the same architecture as in (Kempka et al. 2016).

As mentioned in Sect. 6, this environment is very challenging due to the nature of its observation space. In our experiments, we found that, instead of using the discrete latent states directly, it was beneficial to use the continuous relaxation of discrete random variables (Maddison et al. 2017) learned by WAE-MDPs (see (Delgrange et al. 2023)) as latent states to explore the environment. This allowed for a smoother optimization. Continuous relaxations rely on a temperature parameter, which intuitively controls the continuity of the latent space. When annealed to zero (the “zero-temperature limit”), the latent space is guaranteed to be discrete. We used the latent space at its zero-temperature limit to verify the values in the latent model learned.

**Reward function.** Denote by  $health(s)$  the health of the agent in state  $s \in \mathcal{S}$  and define  $hit(s, a, s')$  as the function that returns a constant  $C > 0$  when  $a \in \mathcal{A}$  is the “shoot” action and an enemy is hit in  $s'$ ,  $-C$  if no enemy is hit, and 0 if  $a$  is any other action. The base reward function of the agent is given by  $rew(s, a, s') = health(s') - health(s) + r^* \cdot \mathbb{1}\{s \in T\} + hit(s, a, s')$ , where, as for the grid world environment,  $r^* > 0$  is an arbitrary reward obtained upon reaching the target set  $T$ . A state  $s$  is labeled as “bad” when, whatever the action  $a$  played, the probability that  $health(s') - health(s)$  turns negative is non-zero. To guide the agent, we use exactly the same reward shaping scheme as the one we defined for the grid world environment (Eq. 22).

## I Broader Impact

Our work presents primarily theoretical and fundamental results, enhancing the reliability of RL solutions. Our claims are also illustrated experimentally with an experimental environment (involving an agent moving within a grid world amid moving adversaries). Specifically, our approach focuses on providing performance (“reach-”) and safety (“avoid”) guarantees with RL policies. We believe our work may have positive societal impacts in the long-term, including (i) *safety-critical applications*: prevent failures (in, e.g., autonomous driving, healthcare, robotics); (ii) *trust and wide adoption*: builds and improves confidence in RL solutions; (iii) *avoiding harmful behavior*: mitigates unintended, risky actions; and (iv) *performance compliance*: check whether performance standard are met (e.g., in industry).






Figure 12: Environment for  $N = 9$  rooms of  $20 \times 20$  cells. The agent is depicted in yellow (top left), adversaries in red, power-ups as cherries, and the goal at the bottom right.

Table 3: Hyperparameter range used for (WAE-)DQN. If “optimization scheme” is set to “round robin”, separate optimizers are used for the policy, the state embedding function, the WAE-MDP minimizer, and the WAE-MDP maximizer. If “concatenate” is used, the policy, the state embedding function, and the WAE-MDP minimizer share the same optimizer. In that case, all those components are optimized at once by concatenating their loss functions. For details about the WAE-MDP parameters, see (Delgrange et al. 2023).

Parameter	Range	Grid World	VizDoom
Common to DQN and WAE-DQN			
Activation	{ReLU, LeakyReLU, ELU, tanh, sigmoid}, SiLU	LeakyReLU	sigmoid
# Hidden layers per network	{1, 2, 3}	3	2
# Neurons per layer	{128, 256, 512}	128	256
CNN filters		3 → 3 → 3	7 → 4 (strides 2)
CNN kernels		64 → 32 → 16	32 → 32
Optimization scheme	{round robin, concatenate}	round robin	concatenate
DQN			
Use Boltzmann exploration	{Yes, No}	Yes	Yes
Boltzmann temperature	{0.25, 0.10, 0.5, 0.75, 1, 10, 100}	0.75	0.5
Use $\epsilon$ -greedy exploration (decay to $\epsilon = 0.1$ )	{Yes, No}	No	No
Target update period	{1, 250, 500, 1000}	250	250
Target update scale ( $\alpha$ in Algorithm 1)	{ $10^{-4}$ , $5 \cdot 10^{-4}$ , $10^{-3}$ , $5 \cdot 10^{-3}$ , 1}	1	1
Reward scaling	{1, 10, 25, 100}	100	1
Learning rate	{ $6.25 \cdot 10^{-5}$ , $10^{-4}$ , $2.5 \cdot 10^{-4}$ , $10^{-3}$ }	$6.25 \cdot 10^{-5}$	$6.25 \cdot 10^{-5}$
Batch size	{32, 64, 128}	64	64
Use double $Q$ -networks (van Hasselt et al. 2016)	{Yes, No}	Yes	Yes
Categorical network (Bellemare et al. 2017)	{Yes, No}	No	Yes
WAE-MDP			
Latent state size (power of 2)	{12, 13, 14, 15}	13	14
State embedding function temperature	{ $1/3$ , $1/2$ , $2/3$ , $3/4$ , 0.999}	$1/3$	0.999
Transition function temperature	{ $1/3$ , $1/2$ , $2/3$ , $3/4$ , 0.999}	$1/3$	$1/2$
Steady-state regularizer scale factor	{0.01, 0.1, 1, 10, 25, 50, 75}	50	0.01
Transition regularizer scale factor	{ $10^{-2}$ , $10^{-1}$ , 1, 10, 25, 50, 75}	50	0.1
Minimizer learning rate	{ $10^{-4}$ , $5 \cdot 10^{-4}$ , $10^{-3}$ }	$10^{-4}$	/
Maximizer learning rate	{ $10^{-4}$ , $5 \cdot 10^{-4}$ , $10^{-3}$ }	$10^{-4}$	$10^{-3}$
State embedding function learning rate	{ $10^{-4}$ , $5 \cdot 10^{-4}$ , $10^{-3}$ }	$10^{-4}$	/
# critic updates	{3, 5, 10, 15}	5	3
State reconstruction function	{none, $L_2$ , binary cross entropy (for M)}	$L_2$	No reconstruction



(b) Samples of frames processed by the agent.

(a) Top view of the ViZDoom environment (8 rooms). The agent is located at the top leftmost room (depicted by a green point). Red points depict enemies on the map. The goal of the agent is to reach the goal position \*. Exit positions are depicted by colored rectangles: in blue, the agent needs to exit to the right, in yellow to the left, in green below, and in red above. Note that enemies spawn at random positions every 250 “ticks” (4 game ticks correspond to 1 step in the environment).

Figure 13: ViZDoom environment. Demonstration of a synthesized controller: <https://youtu.be/BAVLmsWEaQY>.