
Structured Exploration via Hierarchical Variational Policy Networks

Stephan Zheng¹ Yisong Yue¹

Abstract

Reinforcement learning is challenging in environments with large state-action spaces, as exploration can be highly inefficient. Even if the dynamics are simple, the optimal policy can be combinatorially hard to discover. In this work, we propose a hierarchical approach to structured exploration to improve the sample efficiency of on-policy exploration in large state-action spaces. The key idea is to model a stochastic policy as a hierarchical latent variable model, which can learn low-dimensional structure in the state-action space, and to define exploration by sampling from the low-dimensional latent space. This approach enables lower sample complexity, while preserving the expressiveness of the policy class. To make learning tractable, we derive a joint learning and exploration strategy by combining hierarchical variational inference with actor-critic learning. The benefits of our learning approach are that it is principled, simple to implement, scalable to settings with many actions, and composable with existing deep learning approaches. We evaluate our approach on learning a deep centralized multi-agent policy, as multi-agent environments naturally have an exponentially large state-action space. We demonstrate that our approach can more efficiently learn optimal policies in challenging multi-agent games *with a large number* (~ 20) *of agents*, compared to conventional baselines.

1. Introduction

Reinforcement learning in environments with large state-action spaces is challenging, as exploration can be highly inefficient in high-dimensional spaces. Hence, even if the environment dynamics are simple, the optimal policy can be

^{*}Equal contribution ¹Caltech, Pasadena, CA. Correspondence to: Stephan Zheng <stephan@caltech.edu>.

combinatorially hard to discover. However, for many large-scale environments, the high-dimensional state-action space has (often hidden or implicit) low-dimensional structure which can be exploited.

Many natural examples are in collaborative multi-agent problems, whose state-action space is exponentially large in the number of agents, but have a low-dimensional coordination structure. For instance,




Figure 1. Equivalent solutions in a 2-hunter 2-prey game.

consider a simple variant of the Hare-Hunters problem (see Figure 1). In this game, $N = 2$ identical hunters need to capture $M = 2$ identical *static* prey within T time-steps, and exactly $H = 1$ hunter is needed to capture each prey. T is set such that no hunter can capture both preys. There are two equivalent solutions: hunter 1 captures prey 1 and hunter 2 captures prey 2, or vice versa. There are also two suboptimal choices: both hunters choose the same prey. Hence, the hunters must coordinate *over a (large) number of time-steps* to maximize their reward. This implies the solution space has low-dimensional structure that can be used to accelerate exploration and training.

In this work, we propose a principled approach to structured exploration to improve sample complexity in large state-action spaces, by learning deep probabilistic hierarchical policies with a latent structure. As a high-level intuition, consider a tabular multi-agent policy, which maps discrete (joint) states to action probabilities. For N agents with S states and A actions each, this policy has $O((S \cdot A)^N)$ weights. However, the low-dimensional coordination structure can be captured by a factorized, low-rank matrix, where the factorization can be learned and, for instance, only has $O(NK(S + A))$ weights. Similarly, our approach both 1) learns a low-dimensional “factorization” of the policy distribution and 2) defines exploration by also sampling from the low-dimensional latent space. For instance, in the multi-agent setting, we can learn a centralized multi-agent policy with a latent structure that encodes (a distribution of possible) coordination between agents and biases exploration towards policies that (likely) encode “good” coordination.

The key ideas of our approach are: 1) to utilize a shared

stochastic latent variable model that defines the structured exploration policy, and 2) to employ a principled variational method to learn the posterior distribution over the latents jointly with the optimal policy. Our approach has several desirable properties. First we do not incorporate any form of prior domain knowledge, but rather discover the coordination structure purely from empirical experience during learning. Second, our variational learning method enables fully differentiable end-to-end training of the entire policy class. Finally, by utilizing a hierarchical policy class, our approach can easily scale to large action spaces (e.g. a large number of coordinating agents). Our approach can also be seen as a deep hierarchical generalization of Thompson sampling, which is a historically popular way to capture correlations between actions (e.g. in the bandit setting (Agrawal & Goyal, 2012)).

To summarize, our contributions in this work are as follows:

- We introduce a structured probabilistic policy class that uses a hierarchy of stochastic latent variables.
- We propose an efficient and principled algorithm using variational methods to train the policy end-to-end.
- To validate our learning framework, we introduce several challenging synthetic multi-agent environments that explicitly require team coordination, and feature competitive pressures that are characteristic of many coordinated decision problems.
- We empirically verify that our approach improves sample complexity on coordination games with a large number ($N \sim 20$) of agents.
- We show that learned latent structures correlate with meaningful coordination patterns, which implies that our latent structure can recover meaningful low-dimensional structure in the state-action space.

2. Problem Setup & Approach

In this paper, we use centralized multi-agent environments to showcase the efficacy of our approach to structured exploration, as they naturally exhibit exponentially large state-action spaces. More generally, our approach can be applied to any reinforcement learning setting where one can posit the existence of a low-dimensional structure within a high dimensional state-action space.

In centralized multi-agent RL, agents sequentially interact within an environment defined by the tuple: $\mathcal{E} \equiv (\mathbf{S}, \mathbf{A}, \mathbf{r}, f_P)$. Each agent i starts in an initial state s_0^i , and at each time t observes a state $\mathbf{s}_t \in \mathbf{S}$ and executes an action a_t^i chosen by a (stochastic) policy $a_t^i \sim P(a_t^i | \mathbf{s}_t)$. Each agent then receives a reward $r^i(\mathbf{s}_t, \mathbf{a}_t)$, and the environment transitions to a new state \mathbf{s}_{t+1} with probability $f_P(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t)$.

We define the joint state and actions as $\mathbf{s}_t = \{s_t^i \in \mathbf{S}\}$ and $\mathbf{a}_t = \{a_t^i \in \mathbf{A}\}$, where $i \in \mathcal{I}$ indexes the agents. Note that the rewards for each agent r^i can depend on the full joint state and actions.

In this work, we restrict to fully cooperative MDPs that are fully observable, deterministic and episodic.¹ Each agent can see the full state \mathbf{s} , f_P is deterministic and each episode $\tau = (\mathbf{s}_t, \mathbf{a}_t)_{0 \leq t \leq T}$ ends when the agent encounters a terminal state and the MDP resets.

In the fully cooperative case, the goal for each agent is to learn its optimal policy $P^*(a_t^i | \mathbf{s}_t, \boldsymbol{\theta})$ that maximizes the total reward $R(\tau) = \sum_i R^i(\tau) = \sum_i \sum_t r^i(\mathbf{s}_t, \mathbf{a}_t)$:

$$\max_{\boldsymbol{\theta}} J(\boldsymbol{\theta}) = \max_{\boldsymbol{\theta}} \sum_{i \in \mathcal{I}} J^i(\boldsymbol{\theta}), \quad (1)$$

$$J^i(\boldsymbol{\theta}) = \mathbb{E} [R^i(\tau) | \mathbf{a}_t \sim P(\mathbf{a}_t | \mathbf{s}_t; \boldsymbol{\theta})] \quad (2)$$

To optimize, we can apply gradient descent with policy gradient estimators $\hat{g}_{\boldsymbol{\theta}}$ (Williams, 1992)

$$\begin{aligned} g_{\boldsymbol{\theta}} &= \nabla_{\boldsymbol{\theta}} J(\boldsymbol{\theta}) = \mathbb{E} [\nabla_{\boldsymbol{\theta}} \log P(\mathbf{a}_t | \mathbf{s}_t; \boldsymbol{\theta}) R(\tau) | \mathbf{a}_t, \mathbf{s}_t] \\ &\approx \frac{1}{M} \sum_{k=1}^M \sum_t \nabla_{\boldsymbol{\theta}} \log P(\mathbf{a}_t^k | \mathbf{s}_t^k; \boldsymbol{\theta}) R(\tau^k), \end{aligned} \quad (3)$$

where we sample M rollouts τ^k by sampling actions from the policy that is being learned.

Challenges in Exploration. A central issue in reinforcement learning is the *exploration-exploitation* trade-off: how can the policy sample rollouts and learn efficiently? In particular, when the state-action space is exponentially large (as in multi-agent RL), discovering a good policy becomes intractable as the number of agents N grows. Hence, exploration in large state-action spaces poses a significant challenge. We consider the general policy gradient setting where exploration is driven by the stochastics of the policy (similar to Thompson Sampling (Agrawal & Goyal, 2012) in the bandit setting). Thus, it is important for the policy to maintain *expressive low-dimensional* distributions that can make exploration more efficient.

2.1. Hierarchical Variational Policy

We now formulate (1) using a stochastic hierarchical policy class that enables structured exploration. Our approach builds upon two complementary approaches:

- Encode structured exploration using a latent variable that “factorizes” the state-action space. In the multi-agent setting, the factorization is per-agent and the

¹More generally, multi-agent problems can be generalized along many dimensions, e.g. one can learn decentralized policies in partial-information settings. For an overview, see (Busoniu et al., 2008).

latent variable encode coordination encodes coordination between agents.

- Use a variational approach to derive and optimize a lower bound on the objective (1).

Hierarchical Latent Model. For simplicity, we assume a factorization of the state-action space into state-action spaces for individual agents in the multi-agent RL setting. More generally, one could consider other factorization structures as well. To encode coordination between agents, we assume the individual per-agent sub-policies have shared structure, encoded by a latent variable $\lambda_t \in \mathbb{R}^n$ for all t , where n is the dimension of the latent space. This leads to a hierarchical policy model $P(\mathbf{a}_t, \lambda_t | \mathbf{s}_t)$, as shown in Figure 2. We first write the joint policy for a single time-step as:

$$\begin{aligned} P(\mathbf{a}_t | \mathbf{s}_t) &= \int d\lambda_t P(\mathbf{a}_t, \lambda_t | \mathbf{s}_t) = \int d\lambda_t \prod_{i=1}^N P(a_t^i, \lambda_t | \mathbf{s}_t) \\ &= \int d\lambda_t \prod_{i=1}^N P(a_t^i | \lambda_t, \mathbf{s}_t) P(\lambda_t | \mathbf{s}_t), \end{aligned} \quad (4)$$

where we introduced the conditional priors $P(\lambda_t | \mathbf{s}_t)$. The latent variables λ_t introduce dependencies among the \mathbf{a}_t , hence this policy is more flexible compared to standard fully factorized policies (Ranganath et al., 2015). Note that this approach supports centralized learning and decentralized execution, by sharing a random seed amongst agents to sample λ_t and actions \mathbf{a}_t during execution.

Computing the integral in the optimal policy (4) is hard, because the unknown distribution $P(a_t^i | \lambda_t, \mathbf{s}_t)$ can be highly complex. Hence, to make learning (4) tractable, we will use a variational approach.

Hierarchical Variational Lower Bound. We next derive a tractable learning algorithm using variational methods. Instead of directly optimizing (1), we cast it as a probabilistic inference problem, as in (Levine & Koltun, 2013; Vlassis et al., 2009), and instead optimize a lower bound. To do so, we assume that the total reward R^i for each i to be non-negative and bounded. Hence, we can view the total reward $R(\tau)$ as a *random variable*, whose *unnormalized* distribution is defined as $P(R|\tau) = R$. We can then rewrite (1) as a maximum likelihood problem:

$$\max_{\theta} \mathbb{E}_{\pi(\mathbf{s}_t; \theta)} [R] = \max_{\theta} \mathbb{E}_{\pi(\mathbf{s}_t; \theta)} [P(R|\tau)] \quad (5)$$

Hence, the RL objective is equivalent to a maximal likelihood problem:

$$\begin{aligned} \max_{\theta} \mathbb{E}_{\pi(\mathbf{s}_t; \theta)} [P(R|\tau)] &= \max_{\theta} \int d\tau P(R|\tau) P(\tau; \theta) \\ &\Leftrightarrow \max_{\theta} \int d\tau \log P(R|\tau) P(\tau; \theta). \end{aligned} \quad (6)$$

We introduce a latent variable λ_t in the probability of a rollout τ , using (4):

$$\begin{aligned} P(\tau; \theta) &= \\ P(\mathbf{s}_0) \int d\lambda_{0:T} \prod_{t=0}^T P(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t) P(\mathbf{a}_t, \lambda_t | \mathbf{s}_t; \theta). \end{aligned} \quad (7)$$

Here, we defined

$$\int d\lambda_{0:T} \equiv \prod_{t=0}^T \int d\lambda_t, \quad (8)$$

where computing the policy distribution $P(\mathbf{a}_t, \lambda_t | \mathbf{s}_t; \theta)$ is intractable, which makes the maximization in Equation (6) hard. Hence, we derive a lower bound on the log-likelihood $\log P(R|\tau) P(\tau; \theta)$ in Equation (6), using a variational approach. Specifically, we use an approximate *factorized* variational distribution Q_R that is weighted by the reward R :

$$\begin{aligned} Q_R(\lambda_{0:T} | \tau; \phi) &= \\ P(R|\tau) P(\mathbf{s}_0) \prod_{t=0}^T P(\mathbf{s}_{t+1} | \mathbf{s}_t, \mathbf{a}_t) Q(\lambda_t | \mathbf{s}_t; \phi), \end{aligned} \quad (9)$$

where ϕ are the parameters for the variational distribution Q_R . Using Jensen's inequality (Hoffman et al., 2013) and (4) to factorize $P(\mathbf{a}_t, \lambda_t | \mathbf{s}_t; \theta)$, we can derive:

$$\begin{aligned} \log P(R|\tau) P(\tau; \theta) &\geq \int d\lambda_{0:T} Q_R(\lambda_{0:T} | \tau; \phi) \log \left(\frac{P(R|\tau) P(\tau; \theta)}{Q_R(\lambda_{0:T} | \tau; \phi)} \right) \\ &= \int d\lambda_{0:T} Q_R(\lambda_{0:T} | \tau; \phi) \\ &\quad \times \sum_{t=0}^T \left(\log P(\mathbf{a}_t | \lambda_t, \mathbf{s}_t; \theta) + \log \frac{P(\lambda_t | \mathbf{s}_t)}{Q(\lambda_t | \mathbf{s}_t; \phi)} \right) \\ &= \text{ELBO}(Q_R, \theta, \phi), \end{aligned} \quad (10)$$

where the right-hand side is called the evidence lower bound (ELBO), which we can maximize as a proxy for (5). For more details on the derivation, see the Appendix.

The standard choice for the prior $P(\lambda_t | \mathbf{s}_t)$ is to use maximum-entropy standard-normal priors: $P(\lambda_t | \mathbf{s}_t) = \mathcal{N}(\mathbf{0}, \mathbf{1})$. We can then optimize (10) using e.g. stochastic gradient ascent. Formally, the policy gradient is:

$$\begin{aligned} g_{\theta} &\approx g_{\theta, Q} = \nabla_{\theta} \text{ELBO}(Q_R, \theta, \phi) \\ &= \int d\tau d\lambda_{0:T} Q_R(\lambda_{0:T} | \tau; \phi) \\ &\quad \times \sum_{t'=0}^T \nabla_{\theta} \log P(\mathbf{a}_{t'} | \lambda_{t'}, \mathbf{s}_{t'}; \theta), \end{aligned} \quad (11)$$

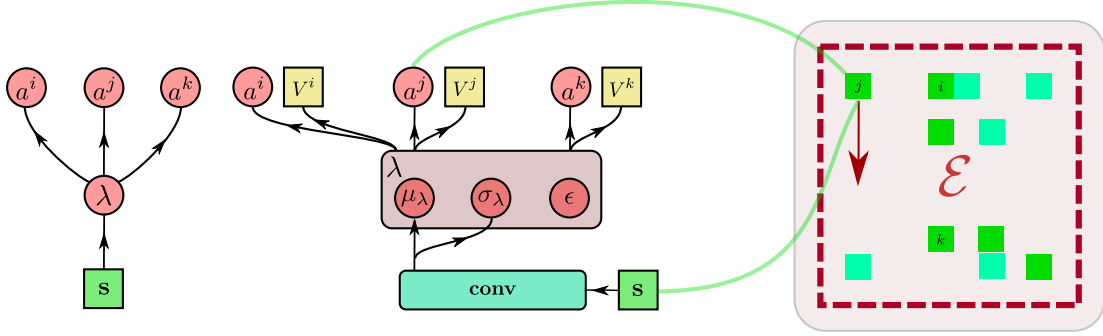


Figure 2. Structured latent variable model of a centralized multi-agent policy (*actor*, left) and instance of the actor-critic interacting in the environment \mathcal{E} (right). The policy contains two stacked layers of stochastic latent variables (red), and deterministically receives states and computes actions (green). The actions for individual agents are independent conditional on latent variables λ , which implies a “factorization” of the state-action space. On the right, a neural network instance of the actor-critic uses the reparameterization trick and receives the environment state, samples actions from the policy (for all agents) and computes the value function V^\bullet .

which is an approximation of the true policy gradient (3), and can be estimated using roll-outs τ^k of the policy P^π :

$$g_{\theta, Q} \approx \hat{g}_{\theta, Q} = \frac{1}{M} \sum_{k=1}^M \sum_{t=0}^T \nabla_{\theta} \log P(\mathbf{a}_t^k | \lambda_t^k, \mathbf{s}_t^k; \theta) R(\tau^k). \quad (12)$$

During a rollout τ^k , we sample $\lambda \sim Q$, observe rewards $R \sim P(R|\tau)$ and transitions $s_{t+1} \sim P(s_{t+1}|\cdot)$, and use these to compute (12). We can similarly compute $g_{\phi, Q} = \nabla_{\phi} \text{ELBO}(Q_R, \theta, \phi)$, the gradient for the variational posterior Q_R .

Actor-Critic and Bias-Variance. Estimating policy gradients g_{θ} using empirical rewards can suffer from high variance. It is useful to consider more general objectives F^i :

$$J^i(\theta) = \mathbb{E}[F^i(\tau) | \mathbf{a}_t \sim P^{\pi_t}(\mathbf{a}_t | \mathbf{s}_t; \theta)], \quad (13)$$

such that the variance in \hat{g} is reduced.² In practice, we find that using (12) with more general F , such as generalized advantages (Schulman et al., 2015), performs quite well.

3. Experimental Validation

3.1. Multi-agent Environments

To validate our approach, we created two grid-world games, depicted in Figure 2, inspired by the classic Predator-Prey and Stag-Hunt games (Shoham & Leyton-Brown, 2008). In both, the world is periodic and the initial positions of the hunters and prey are randomized. We consider two instances for both games: either the prey are moving or fixed.

Hare-Hunters. Predator-Prey is a classic test environment for multi-agent learning, where 4 predators try to capture a

²Note that if the total reward R is bounded, we can define R -weighted probabilities by shifting and normalizing R . In general, the derivation applies if F is similarly bounded.

prey by boxing it in. We consider a variation defined by the settings (N, M, H, T) : N hunters and M prey. Each prey can be captured by exactly H hunters: to capture the prey, a hunter gets next to it, after which the hunter is frozen. Once a prey has had H hunters next to it, it is frozen and cannot be captured by another hunter. The terminal rewards are:

$$R^i = \begin{cases} 1, & \text{if all prey are captured } H \text{ times before the} \\ & \text{time limit } T \\ 0, & \text{otherwise} \end{cases} \quad (14)$$

The challenge of the game is for the agents to inactivate all prey within a finite time T . Due to the time limit, the optimal strategy is for the agents to distribute targets efficiently, which can be challenging due to the combinatorially large number of possible hunter-to-prey assignments.

Stag-Hunters. The Stag-Hunt is another classic multi-agent game designed to study coordination. In this game, hunters have a choice: either they capture a hare for low reward, or, together with another hunter, capture a stag for a high reward. We extend this to the multi-agent (N, M, H, T) -setting: N hunters hunt M prey ($M/2$ stags and $M/2$ hares). Each stag has H hit-points, while hares and hunters have 1 hit-point. Capturing is as in Hare-Hunters. The spatial domain is similar to the Hare-Hunters game and we also use a time limit T . The terminal reward is now defined as:

$$R^i = \begin{cases} 1, & \text{if } i \text{ captured a live stag that became inactive} \\ & \text{before the time limit } T \\ 0.1, & \text{if } i \text{ captured a live hare} \\ & \text{before the time limit } T \\ 0, & \text{otherwise} \end{cases} \quad (15)$$

Preys are	Frozen		Moving		Frozen		Moving	
Samples (x100k)	5	10	5	10	5	10	5	10
Hare-Hunters	10-vs-10				20-vs-20			
Cloned	85.0	178.3	465.0	912.5	20.0	70.0	706.7	1401.7
Shared	65.0	155.0	457.5	923.8	65.0	105.0	491.4	962.9
HLPN	240.0	580.0	662.7	1381.8	200.0	393.3	1260.0	2344.0
Stag-Hunters	10-vs-10				20-vs-20			
Cloned	1229.2	2482.9	2079.4	4171.2	3224.5	6219.9	5934.5	11429.2
Shared	1214.5	2423.7	2005.7	4144.3	3150.7	6379.8	6344.4	12196.8
HLPN	1515.2	3047.3	2275.7	4610.7	3799.3	7158.1	6880.7	13358.6

Table 1. Total terminal reward (averaged over 5 best runs) for N agents for set # of training samples. Our HLPN approach outperforms baselines by up to 10x.

The challenge for the agents here is to discover that choosing to capture the same prey can yield substantially higher reward, but this requires coordinating with another hunter.

3.2. Neural Coordination Model

For experiments, we instantiated our policy class (as in Figure 2) with deep neural networks. For simplicity, we only used *reactive policies* without memory, although it is straightforward to apply our approach using policies with memory (e.g. LSTMs). The model takes a joint state \mathbf{s}_t as input and computes features $\phi(\mathbf{s})$ using a 2-layer convolutional neural network. To compute the latent variable $\lambda \in \mathbb{R}^d$, we use the reparametrization trick (Kingma & Welling, 2013) to learn the variational distribution (e.g. $Q(\lambda|\mathbf{s})$), sampling λ via $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ and distribution parameters μ, σ (omitting t):

$$\begin{aligned} \mu(\mathbf{s}) &= W_\mu \phi(\mathbf{s}) + b_\mu, & \log \sigma(\mathbf{s})^2 &= W_\sigma \phi(\mathbf{s}) + b_\sigma, \\ \lambda &= \mu(\mathbf{s}) + \sigma(\mathbf{s}) \odot \epsilon. \end{aligned} \quad (16)$$

Given λ , the model then computes the policies $P(a^i|\lambda, \mathbf{s})$ and value functions $V^i(\mathbf{s})$ as (omitting t):

$$\begin{aligned} P(a^i|\lambda, \mathbf{s}) &= \text{softmax}(W_\pi^i[\lambda \phi(\mathbf{s})] + b_\pi^i), \\ V^i(\mathbf{s}) &= W_V^i \phi(\mathbf{s}) + b_V^i, \end{aligned} \quad (17)$$

where $\text{softmax}(\mathbf{x}) = \exp \mathbf{x} / \sum_j \exp x^j$. In this way, the model can be trained end-to-end.

Training. We used A3C (Mnih et al., 2016) with KL-controlled policy gradients (12), generalized advantage as F (Schulman et al., 2015), and policy-entropy regularization. The loss for the value function at each state \mathbf{s}_t is the standard L_2 -loss between the observed total rewards for each agent i and its value estimate:

$$\alpha \sum_t \sum_i (V^i(\mathbf{s}_t) - R^i(\mathbf{s}_t, \mathbf{a}_t))^2. \quad (18)$$

In addition, in line with other work using actor-critic methods, we found that adding a small entropy regularization on the policy can sometimes positively influence performance, but this does not seem to be always required for our testbeds. The entropy regularization is:

$$H(P) = -\beta \sum_t \sum_{\mathbf{a}_t} P(\mathbf{a}_t|\mathbf{s}_t; \theta) \log P(\mathbf{a}_t|\mathbf{s}_t; \theta). \quad (19)$$

A3C additionally defines training minibatches in terms of a fixed number of environment steps L : a smaller L gives faster training with higher variance and vice versa.

For all experiments, we performed a hyper-parameter search and report the best 5 runs seen (see the Appendix for details).

Baselines. We compared against two natural baselines:

- **Shared** (shared actor-critic): agents share a deterministic hidden layer, but maintain individual weights θ^i for their (stochastic) policy $P(a|\lambda, \mathbf{s}; \theta^i)$ and value function $V^i(\mathbf{s}; \theta)$. The key difference is that this model does not sample from the shared hidden layer.
- **Cloned** (actor-critic): each agent uses an identical policy and value function with shared weights. There is shared information between the agents, and actions are sampled according to the agents' own policies.

4. Quantitative Analysis

We now validate the efficacy of our approach by showing our method scales to environments with a large number of agents. We ran experiments for both Hare-Hunters and Stag-Hunters for $N = M = 10, 20$ in a spatial domain of 50×50 grid cells. We refer to our approach as HLPN (Hierarchical Latent Policy Network).

Sample complexity. In Table 1 we show the achieved rewards after a fixed number of training samples, and Figure

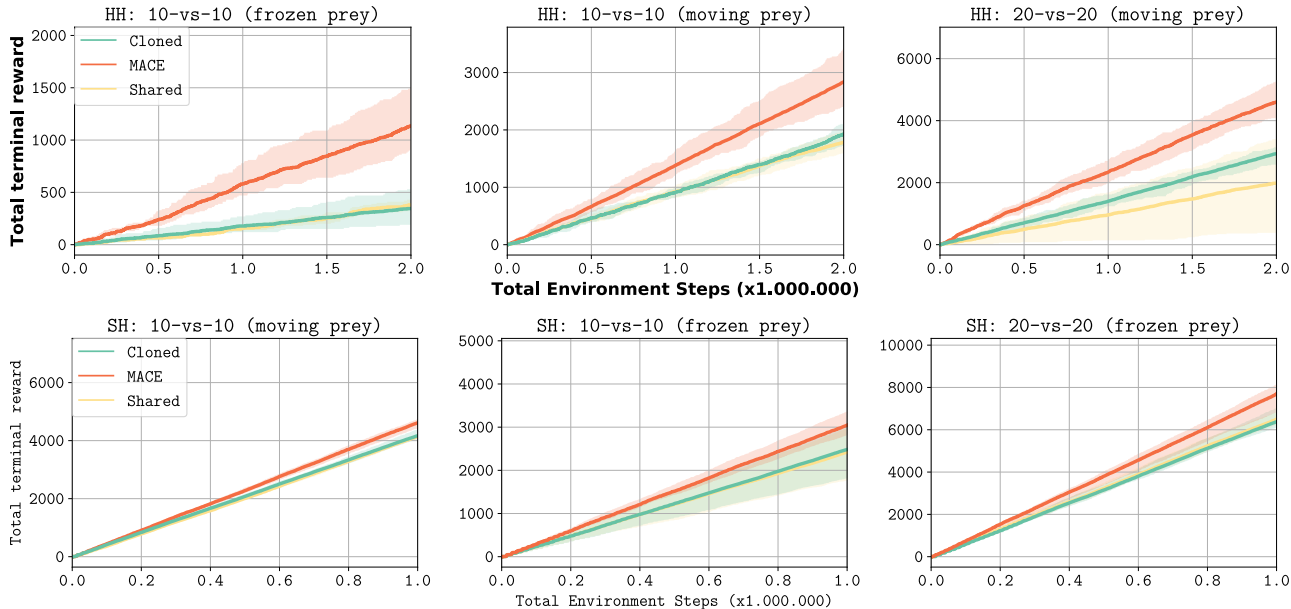


Figure 3. Train-time cumulative terminal reward for N agents in $(N, M, 1, T)$ Hare-Hunters (upper, $T = 2000$) and Stag-Hunters (lower, $T = 1000$) on a 50×50 gridworld, for **10-vs-10** or **20-vs-20** agents; randomly moving or fixed preys. Average, minimal and maximal rewards for the best 5 runs for each model are shown. Our HLPN approach accumulates increasingly higher rewards compared to the baselines, by 1) achieving higher terminal reward per episode and 2) finishing episodes faster (see Figure 4). For 10-10 Stag-Hunters with frozen prey, average reward per-episode is 4.64 (Cloned), 6.22 (Shared), 6.61 (HLPN) after 1 million samples. For more, see the Appendix.




Figure 4. Train-time episode lengths during 1 million steps for **10-vs-10** Hare-Hunters (left) and Stag-Hunters (right), with fixed preys. Our HLPN approach (orange) finishes an episode successfully before the time limit more often than the baselines (Cloned (blue) and Shared (yellow)).

3 shows the corresponding learning curves. We see that HLPN achieves up to $10\times$ reward compared to the baselines. Figure 4 shows the corresponding distribution of training episode lengths. We see that HLPN solves game instances more than 20% faster than baselines in 50% (10%) of Hare-Hunters (Stag-Hunters) episodes. In particular, HLPN learns to coordinate for higher reward more often: it achieves the highest average reward per-episode (e.g. for 10-10 Stag-Hunters with frozen prey, average rewards are 4.64 (Cloned), 6.22 (Shared), 6.61 (HLPN)). Hence, HLPN coordinates successfully more often to capture the stags. Together, these results show HLPN enables more efficient learning.

Using the ELBO. A salient difference between (12) and (3) is the KL-regularization, which stems from the derivation of the ELBO. Since we use a more general objective F , c.f. (13), we also investigated the impact of using the KL-regularized policy gradient (12) versus the standard (3). We ran several instances of the above experiments both with and without KL-regularization. We found that without KL-regularization, training is unstable and prone to mode collapse: the variance σ of the variational distribution can go to 0, leading to essentially 0 achieved reward for any reasonable hyperparameter settings.

Impact of dynamics and T . Inspecting training performance, we see the relative difficulty of capturing moving or

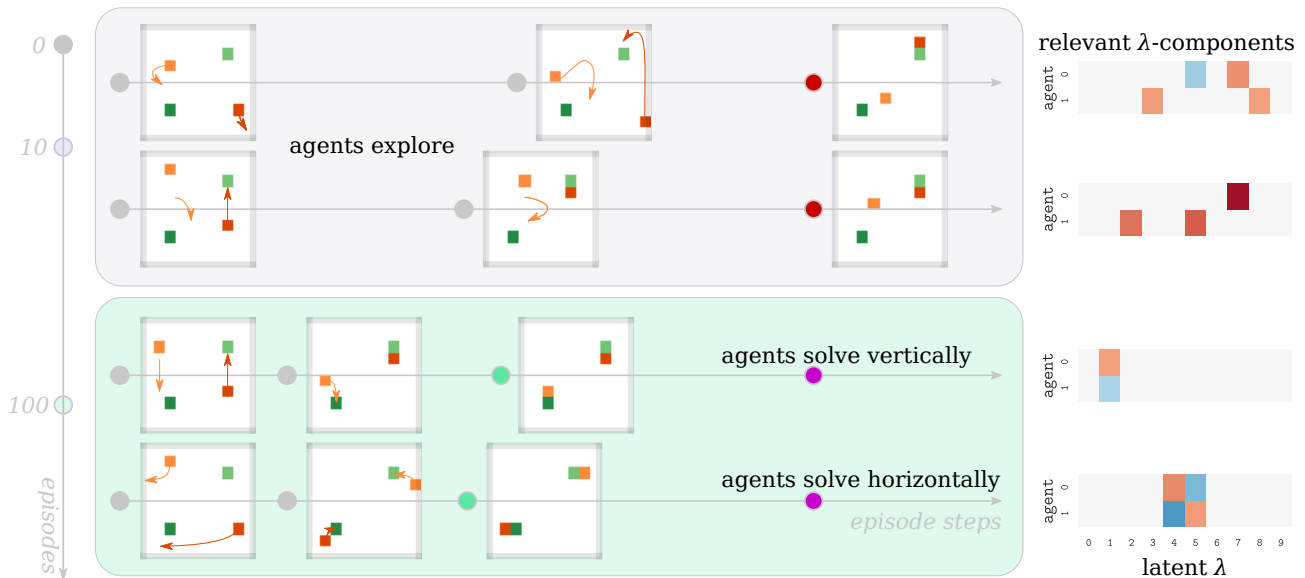


Figure 5. Predators (red) and prey (green) during training for 2v2 Hare-Hunters for 100 episodes. Arrows show where agents move to in the next frame. Top: at the start, predators explore via λ , but do not succeed before the time limit T (red dot). Bottom: after convergence agents succeed consistently (green dot) before the time limit (purple dot) and λ encodes the two strategies from Figure 1. Highlighted λ -components correlate with rollout under a 2-sided t -test at $\alpha = 0.1$ significance.




Figure 6. Visualization of our HLPN approach for a (10, 10, 1, 1000) Hare-Hunters game in a 30×30 world. Left: components of the latent code λ that significantly correlate with sampled actions (computed as in Figure 5). Right: Three episode snapshots: at the start, middle and end. Red: predators; green: prey. Arrows indicate where agents move to in the next snapshot. The hunters solve the game (green dot) before the time limit $T = 1000$ (purple dot), by distributing targets amongst themselves.

randomly moving prey. Capturing moving prey is easier to learn than capturing fixed preys, as comparing rewards in Table 1 shows. This shows a feature of the game dynamics: the expected distance between a hunter and an uncaptured prey are lower when the preys are randomly moving, resulting in an easier game. Comparing Hare-Hunters and Stag-Hunters, we also see the impact of the time limit T . Since we use terminal rewards only, as T gets larger, the reward becomes very sparse and models need more samples to discover good policies.

5. Model Inspection

Beyond training benefits, we now demonstrate empirical evidence that suggest efficacy and meaningfulness of our ap-

proach to structured exploration. We start by inspecting the behavior of the latent variable λ for a simple $N = M = 2$ Hare-Hunters game, which enables semantic inspection of the learned policies, as in Figure 5. We make a number of observations. First, λ is relevant: many components are statistically significantly correlated with the agents' actions. This suggests the model does indeed use the latent λ : it (partly) controls the coordination between agents.³ Second, the latent λ shows strong correlation during all phases of training. This suggests that the model indeed is performing a form of structured exploration. Third, the components of λ are correlated with semantic meaningful behavior. We show a salient example in the bottom 2 rows in Figure 5:

³This is parallel to the discussion in (Chen et al., 2016), which investigates the effectiveness of latent codes λ .

the correlated components of λ are disjoint and each component correlates with both agents. The executed policies are exactly the two equivalent ways to assign 2 hunters to 2 preys, as illustrated in Figure 1.

Coordination with a large N . In the large $N = M = 10$ case, the collective dynamics generalize the $N = M = 2$ case. There are now redundancies in multi-agent hunter-prey assignments that are analogous to the $N = M = 2$ case that are prohibitively complex to analyze due to combinatorial complexity. However our experiments strongly suggest (see e.g. Figure 6) the latent code is again correlated with agents’ behavior during all phases of training, hence λ induces meaningful multi-agent coordination.

6. Related Work

Deep Structured Inference. Recent works have focused on learning structured representations using expressive distributions, which enable more powerful probabilistic inference. For instance, (Johnson et al., 2016) has proposed combining neural networks with graphical models, while (Ranganath et al., 2015) learn hierarchical latent distributions. Our work builds upon these approaches to learn structured policies in the reinforcement learning setting. In the multi-agent setting, the RL problem has also been considered as an inference problem in e.g. (Liu et al., 2015; Wu et al., 2013; Liu et al., 2016).

Variational methods in RL. (Neumann, 2011; Furnston & Barber, 2010) discuss variational approaches for RL problems, but did not consider end-to-end trainable models. (Levine & Koltun, 2013) used variational methods for guided policy search. (Houthoofd et al., 2016) learned exploration policies via information gain using variational methods. However, these only consider 1 agent.

Coordination in RL. Multi-agent coordination has been studied in the RL community (e.g. (Guestrin et al., 2002; Kapetanakis & Kudenko, 2002; Chalkiadakis & Boutilier, 2003), for instance, as a method to reduce the instability of multiple agents learning simultaneously using RL. The benefit of coordination was already demonstrated in simple multi-agent settings in e.g. (Tan, 1993). The shared latent variable λ of our structured policy can also be interpreted as a learned *correlation device* (see (Bernstein, 2005) for an example in the decentralized setting), which can be used to e.g. break ties between alternatives or induce coordination between agents. More generally, they can be used to achieve correlated equilibria (Greenwald & Hall, 2003), a more general solution concept than Nash equilibria. However, previous methods learned hand-crafted models and do not scale well to complex state spaces and many agents. In contrast, our method learns coordination end-to-end via on-policy methods, *learns* the multi-agent exploration policy

and scales well to many agents via its simple hierarchical structure.

Communication Models. Recently, end-to-end learning of communication models has been studied for shared broadcast channels (Sukhbaatar et al., 2016), sequential communication (Peng et al., 2017), heuristic multi-agent exploration (Usunier et al., 2016) and bit-channels (Foerster et al., 2016). These works show that communication protocols can be learned through back-propagation or heuristic stabilization methods, but often do not scale well to a large number of agents. Our hierarchical approach is complementary, learns via variational methods, and can scale to large N .

Multi-task learning. Hierarchical models have been studied for multi-task learning, e.g. (Daume III, 2014) learns latent hierarchies via EM in a supervised learning setting. Instead, we study flexible end-to-end trainable latent hierarchies in the reinforcement learning setting.

7. Discussion & Conclusion

We proposed a hierarchical deep policy network for conducting structured exploration in complex state-action spaces. We further derive a variational learning approach to train our policy in the actor-critic framework. We validate our approach on a collection of challenging multi-agent predator-prey style games, and demonstrate significant improvements over conventional baselines. We finally show that the latent variables in our hierarchical policy capture meaningful coordination between the agents, which suggests that the low-dimensional latent variable bottleneck in the hierarchical structure does indeed enable more efficient exploration.

In a sense, we studied the simplest setting that can benefit from structured exploration, in order to isolate the contribution of our work. Our hierarchical model and variational approach are a simple way to implement multi-agent coordination, and easily combine with existing actor-critic methods. Moving forward, there are many ways to expand on our work. Firstly, for complex (partial-information) environments, instead of using reactive policies with simple priors $P \sim \mathcal{N}(0, 1)$, memoryfull policies with flexible priors (Chen et al., 2016) may be needed. Secondly, within the multi-agent setting, our approach is complementary to richer forms of communication between agents. Our hierarchical structure can be interpreted as a broadcast channel, where agents are passive receivers of the message λ . Richer communication protocols could be encoded by policies with more complex inter-agent structure. Finally, there are other interesting combinatorial structures beyond multi-agent settings, and it would be interesting to investigate how to learn over those structures as well.

References

- Agrawal, Shipra and Goyal, Navin. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on Learning Theory*, pp. 39–1, 2012.
- Bernstein, Daniel S. Bounded policy iteration for decentralized pomdps. In *International Joint Conference on Artificial Intelligence*. Citeseer, 2005.
- Busoniu, Lucian, Babuska, Robert, and De Schutter, Bart. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, And Cybernetics-Part C: Applications and Reviews*, 38 (2), 2008, 2008.
- Chalkiadakis, Georgios and Boutilier, Craig. Coordination in Multiagent Reinforcement Learning: A Bayesian Approach. In *International Joint Conference on Autonomous Agents and Multiagent Systems*, AAMAS '03, New York, NY, USA, 2003.
- Chen, Xi, Kingma, Diederik P., Salimans, Tim, Duan, Yan, Dhariwal, Prafulla, Schulman, John, Sutskever, Ilya, and Abbeel, Pieter. Variational Lossy Autoencoder. *arXiv:1611.02731 [cs, stat]*, November 2016.
- Daume III, Hal. Bayesian Multitask Learning with Latent Hierarchies. *arXiv:1408.2032 [cs, stat]*, August 2014.
- Foerster, Jakob N., Assael, Yannis M., de Freitas, Nando, and Whiteson, Shimon. Learning to Communicate with Deep Multi-Agent Reinforcement Learning. *arXiv:1605.06676 [cs]*, May 2016. arXiv: 1605.06676.
- Furmston, T. and Barber, D. Variational methods for reinforcement learning. *Journal of Machine Learning Research*, 9:241–248, December 2010. ISSN 1533-7928.
- Greenwald, Amy and Hall, Keith. Correlated q-learning. In *International Conference on Machine Learning (ICML-03)*, pp. 242–249, 2003.
- Guestrin, Carlos, Lagoudakis, Michail, and Parr, Ronald. Coordinated reinforcement learning. In *ICML*, 2002.
- Hoffman, Matthew D., Blei, David M., Wang, Chong, and Paisley, John. Stochastic variational inference. *Journal of Machine Learning Research*, 14:1303–1347, 2013.
- Houthoofd, Rein, Chen, Xi, Duan, Yan, Schulman, John, De Turck, Filip, and Abbeel, Pieter. VIME: Variational Information Maximizing Exploration. *arXiv:1605.09674 [cs, stat]*, May 2016.
- Johnson, M. J., Duvenaud, D., Wiltschko, A. B., Datta, S. R., and Adams, R. P. Composing graphical models with neural networks for structured representations and fast inference. *ArXiv e-prints*, March 2016.
- Kapetanakis, Spiros and Kudenko, Daniel. Reinforcement Learning of Coordination in Cooperative Multi-agent Systems. In *National Conference on Artificial Intelligence*, pp. 326–331, Menlo Park, CA, USA, 2002.
- Kingma, Diederik P and Welling, Max. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Levine, Sergey and Koltun, Vladlen. Variational Policy Search via Trajectory Optimization. In Burges, C. J. C., Bottou, L., Welling, M., Ghahramani, Z., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 26*, pp. 207–215. Curran Associates, Inc., 2013.
- Liu, Miao, Amato, Christopher, Liao, Xuejun, Carin, Lawrence, and How, Jonathan P. Stick-breaking policy learning in decpmdps. In *IJCAI*, pp. 2011–2018, 2015.
- Liu, Miao, Amato, Christopher, Anesta, Emily P, Griffith, John Daniel, and How, Jonathan P. Learning for decentralized control of multiagent systems in large, partially-observable stochastic environments. 2016.
- Mnih, Volodymyr, Badia, Adrià Puigdomènech, Mirza, Mehdi, Graves, Alex, Lillicrap, Timothy P., Harley, Tim, Silver, David, and Kavukcuoglu, Koray. Asynchronous Methods for Deep Reinforcement Learning. *arXiv:1602.01783 [cs]*, February 2016. arXiv: 1602.01783.
- Neumann, Gerhard. Variational inference for policy search in changing situations. In *International Conference on Machine Learning, ICML 2011*, June 2011.
- Peng, Peng, Yuan, Quan, Wen, Ying, Yang, Yaodong, Tang, Zhenkun, Long, Haitao, and Wang, Jun. Multiagent Bidirectionally-Coordinated Nets for Learning to Play StarCraft Combat Games. *arXiv:1703.10069 [cs]*, March 2017. arXiv: 1703.10069.
- Ranganath, Rajesh, Tran, Dustin, and Blei, David M. Hierarchical Variational Models. *arXiv:1511.02386 [cs, stat]*, November 2015.
- Schulman, John, Moritz, Philipp, Levine, Sergey, Jordan, Michael, and Abbeel, Pieter. High-Dimensional Continuous Control Using Generalized Advantage Estimation. *arXiv:1506.02438 [cs]*, June 2015.
- Shoham, Yoav and Leyton-Brown, Kevin. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, New York, NY, USA, 2008. ISBN 9780521899437.
- Sukhbaatar, Sainbayar, Szlam, Arthur, and Fergus, Rob. Learning Multiagent Communication with Backpropagation. *arXiv:1605.07736 [cs]*, May 2016. arXiv: 1605.07736.
- Tan, Ming. Multi-Agent Reinforcement Learning: Independent vs. Cooperative Agents. In *International Conference on Machine Learning*, pp. 330–337, 1993.
- Usunier, Nicolas, Synnaeve, Gabriel, Lin, Zeming, and Chintala, Soumith. Episodic Exploration for Deep Deterministic Policies: An Application to StarCraft Micromanagement Tasks. *arXiv:1609.02993 [cs]*, September 2016. arXiv: 1609.02993.
- Vlassis, Nikos, Toussaint, Marc, Kontes, Georgios, and Piperidis, Savas. Learning model-free robot control by a Monte Carlo EM algorithm. *Autonomous Robots*, 27(2):123–130, August 2009. ISSN 0929-5593, 1573-7527.
- Williams, Ronald J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992.
- Wu, Feng, Zilberstein, Shlomo, and Jennings, Nicholas R. Monte-carlo expectation maximization for decentralized pomdps. In *International Joint Conference on Artificial Intelligence*, pp. 397–403. AAAI Press, 2013.