

# Episodic Policy Gradient Training

Hung Le, Majid Abdolshah, Thommen K. George, Kien Do, Dung Nguyen, Svetha Venkatesh

Applied AI Institute, Deakin University, Geelong, Australia

{thai.le, m.abdolshah, thommen.karimpanalgeorge, k.do, dung.nguyen, svetha.venkatesh}@deakin.edu.au

## Abstract

We introduce a novel training procedure for policy gradient methods wherein episodic memory is used to optimize the hyperparameters of reinforcement learning algorithms on-the-fly. Unlike other hyperparameter searches, we formulate hyperparameter scheduling as a standard Markov Decision Process and use episodic memory to store the outcome of used hyperparameters and their training contexts. At any policy update step, the policy learner refers to the stored experiences, and adaptively reconfigures its learning algorithm with the new hyperparameters determined by the memory. This mechanism, dubbed as Episodic Policy Gradient Training (EPGT), enables an episodic learning process, and jointly learns the policy and the learning algorithm’s hyperparameters within a single run. Experimental results on both continuous and discrete environments demonstrate the advantage of using the proposed method in boosting the performance of various policy gradient algorithms.

## Introduction

The current success of deep reinforcement learning relies on the ability to use gradient-based optimizations for policy and value learning (Mnih et al. 2015; Silver et al. 2017). Approaches such as *policy gradient* (PG) methods have achieved remarkable results in various domains including games (Mnih et al. 2016; Schulman et al. 2017; Wu et al. 2017; Fujimoto, Hoof, and Meger 2018), robotics (Kohl and Stone 2004; Peters and Schaal 2006) or even natural language processing (Ziegler et al. 2019). However, the excellent performance of PG methods is heavily dependent on tuning the algorithms’ hyperparameters (Duan et al. 2016; Zhang et al. 2021). Applying a PG method to new environments often requires different hyperparameter settings and thus retuning (Henderson et al. 2018). The large amount of hyperparameters severely prohibits machine learning practitioners from fully utilizing PG methods in different reinforcement learning environments.

As a result, there is a huge demand for automating hyperparameter selection for policy gradient algorithms, and it remains a critical part of the Automated Machine Learning (AutoML) movement (Hutter, Kotthoff, and Vanschoren

2019). Automatic hyperparameter tuning has been well explored for supervised learning. Simple methods such as grid search and random search are effective although computationally expensive (Bergstra and Bengio 2012; Larochelle et al. 2007). Other complex methods such as Bayesian Optimization (BO (Snoek, Larochelle, and Adams 2012)) and Evolutionary Algorithms (EA (Fiszelew et al. 2007)) can efficiently search for optimal hyperparameters. Yet, they still need multiple training runs, have difficulty scaling to high-dimensional settings (Rana et al. 2017) or require extensive parallel computation (Jaderberg et al. 2017). Recent attempts introduce online hyperparameter scheduling, that jointly optimizes the hyperparameters and parameters in single run overcoming local optimality of training with fixed hyperparameters and showing great potential for supervised and reinforcement learning (Jaderberg et al. 2017; Xu, van Hasselt, and Silver 2018; Paul, Kurin, and Whiteson 2019; Parker-Holder, Nguyen, and Roberts 2020).

However, one loophole remains. These approaches do not model the *context of training* in the optimization process, and the problem is often treated as a stateless bandit or greedy optimization (Paul, Kurin, and Whiteson 2019; Parker-Holder, Nguyen, and Roberts 2020). Ignoring the context prevents the use of episodic experiences that can be critical in optimization and planning. As an example, we humans often rely on past outcomes of our actions and their contexts to optimize decisions (e.g. we may use past experiences of traffic to not return home from work at 5pm). Episodic memory plays a major role in human brains, facilitating recreation of the past and supporting decision making via recall of episodic events (Tulving 2002). We are motivated to use such a mechanism in training wherein, for instance, the hyperparameters that helped overcome a past local optimum in the loss surface can be reused when the learning algorithm falls into a similar local optimum. This is equivalent to optimizing hyperparameters based on training contexts. Patterns of bad or good training states previously explored can be reused, and we refer to this process as selecting hyperparameters. To implement this mechanism we use episodic memory. Compared to other learning methods, the use of episodic memory is non-parametric, fast and sample-efficient, and quickly directs the agents towards good behaviors (Lengyel and Dayan 2008; Kumaran, Hassabis, and McClelland 2016; Blundell et al. 2016).




Figure 1: Hyper-RL structure. The hyper-state (green circle) is captured from the PG models’ parameters and gradients at every Hyper-RL step (1). Given the hyper-states, the hyper-agent takes hyper-actions, choosing hyperparameters for the PG method to update the models (2). The update lasts  $U$  steps. After the last update step (3), the RL agent starts environment phase with the current policy, collecting an empirical return  $G$  after  $T$  environment steps (4).  $G$  is used as the hyper-reward for the last policy update step (blue diamond) (5). Other update steps (red diamond) are assigned with hyper-reward 0.

This problem of formulating methods that can take the training context into consideration and using them as episodic experiences in optimizing hyperparameters remains unsolved. The first challenge is to effectively represent the training context of PG algorithms that often involve a large number of neural network parameters. The second challenge is sample-efficiency. Current performant hyperparameter searches (Jaderberg et al. 2017; Parker-Holder, Nguyen, and Roberts 2020) often necessitate parallel interactions with the environments, which is expensive and not always feasible in real-world applications. Ideally, hyperparameter search methods should not ask for additional observations that the PG algorithms already collect. If so, it must be solved as efficiently as possible to allow efficient training of PG algorithms.

We address both these issues with a novel solution, namely Episodic Policy Gradient Training (EPGT)—a PG training scheme that allows on-the-fly hyperparameter optimization based on episodic experiences. The idea is to formulate hyperparameter scheduling as a Markov Decision Process (MDP), dubbed as Hyper-RL. In the Hyper-RL, an agent (hyper-agent) acts to optimize hyperparameters for the PG algorithms that optimize the policy for the agent of the main RL (RL-agent). The two agents operate alternately: the hyper-agent acts to reconfigure the PG algorithms with different hyperparameters, which ultimately changes the policy of the RL agent (update phase); the RL agent then acts to collect returns (environment phase), which serves as the rewards for the hyper-agent. To build the Hyper-RL, we propose mechanisms to model its state, action and reward. In particular, we model the training context as the state of the Hyper-RL by using neural networks to compress the parameters and gradients of PG models (policy/value networks) into low-dimensional state vectors. The action in the Hyper-RL corresponds to the choice of hyperparameters and the reward is derived from the RL agent’s reward.

We propose to solve the Hyper-RL through episodic memory. As an episodic memory provides a direct binding from experiences (state-action) to final outcome (return), it enables fast utilization of past experiences and accelerates the searching of near-optimal policy (Lengyel and Dayan 2008). Unlike other memory forms augmenting RL agents with stronger working memory to cope with partial obser-

vations (Hung et al. 2019; Le, Tran, and Venkatesh 2020) or contextual changes within an episode (Le and Venkatesh 2020), episodic memory persists across agent lifetime to maintain a global value estimation. In our case, the memory estimates the value of a state-action pair in the Hyper-RL by nearest neighbor memory lookup (Pritzel et al. 2017). To store learning experience, we use a novel weighted average nearest neighbor writing rule that quickly propagates the value inside the memory by updating multiple memory slots per memory write. Our episodic memory is designed to cope with noisy and sparse rewards in the Hyper-RL.

Our key contribution is to provide a new formulation for online hyperparameter search leveraging context of previous training experiences, and demonstrate that episodic memory is a feasible way to solve this. This is also the first time episodic memory is designed for hyperparameter optimization. Our rich set of experiments shows that EPGT works well with various PG methods and diverse hyperparameter types, achieving higher rewards without significant increase in computing resources. Our solution has desirable properties, it is (i) computationally cheap and run once without parallel computation, (ii) flexible to handle many hyperparameters and PG methods, and (iii) shows consistent/significant performance gains across environments and PG methods.

## Methods

### Hyperparameter Reinforcement Learning (Hyper-RL)

In this paper, we address the problem of online hyperparameter search. We argue that in order to choose good values, hyperparameter search (HS) methods should be aware of the past training states. This intuition suggests that we should treat the HS problem as a standard MDP. Put in the context of HS for RL, our HS algorithm becomes a Hyper-RL algorithm besides the main RL algorithm. In Hyper-RL, the hyper-agent makes decisions at each policy update step to configure the PG algorithm with suitable hyperparameters  $\psi$ . The ultimate goal of the Hyper-RL is the same as the main RL’s: to maximize the return of the RL agent.

To construct the Hyper-RL, we define its state  $s^\psi$ , action  $a^\psi$  and reward  $r^\psi$ . Hereafter, we refer to them as hyper-state, hyper-action and hyper-reward to avoid confusion with

the main RL’s  $s$ ,  $a$  and  $r$ . Fig. 1 illustrates the operation of Hyper-RL. In the update phase, the Hyper-RL runs for  $U$  steps. At each step, taking the hyper-state captured from the PG models’ parameters and gradients, the hyper-agent outputs hyper-actions, producing hyperparameters for PG algorithms to update the policy/value networks accordingly. After the last update (blue diamond), the resulting policy will be used by the RL agent to perform the environment phase, collecting returns after  $T$  environment interactions. The returns will be used in the PG methods, and utilized as hyper-reward for the last policy update step. Below we detail the hyper-action, hyper-reward and hyper-state.

**Hyper-action** A hyper-action  $\mathbf{a}^\psi$  defines the values for the hyperparameters  $\psi$  of interest. For simplicity, we assume the hyper-action is discrete by quantizing the range of each hyperparameter into  $B$  discrete values. A hyper-action  $\mathbf{a}^\psi$  selects a set of discrete values, each of which is assigned to a hyperparameter (see more Appendix A.2).

**Hyper-reward** The hyper-reward  $r^\psi$  is computed based on the empirical return that the RL agent collects in the environment phase after hyperparameters are selected and used to update the policy. The return is  $G = \mathbb{E}_{s_{t:t+T}, a_{t:t+T}} \left[ \sum_{k=0}^T \gamma^k r_{t+k} \right]$  where  $t$  and  $T$  are the environment step and learning horizon, respectively. Since there can be  $U$  consecutive policy update steps in the update phase, the last update step in the update phase receives hyper-reward  $G$  while others get zero hyper-reward, making the Hyper-RL, in general, a sparse reward problem. That is,

$$\mathbf{r}_i^\psi = \begin{cases} G & \text{if } i = U \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

To define the objective for the Hyper-RL, we treat the update phase as a learning episode. Each learning episode can last for multiple of  $U$  update steps and for each step  $i$  in the episode, we aim to maximize the hyper-return  $G_i^\psi = \sum_{j \geq i}^{Un} \mathbf{r}_j^\psi$  where  $n \in \mathbb{N}^+$ . In this paper,  $n$  is simply set to 1 and thus,  $G_i^\psi = G$ .

**Hyper-state** A hyper-state  $\mathbf{s}^\psi$  should capture the current training state, which may include the status of the trained model, the loss function or the amount of parameter update. We fully capture  $\mathbf{s}^\psi$  if we know exactly the loss surface and the current value of the optimized parameters, which can result in perfect hyperparameter choices. This, however, is infeasible in practice, thus we only model observable features of the hyper-state space. The Hyper-RL is then partially observable and noisy. In the following, we propose a method to represent the hyper-state efficiently.

**Hyper-state representation** Our hypothesis is that one signature feature of the hyper-state is the current value of optimized parameters  $\theta$  and the derivatives of the PG method’s objective function w.r.t  $\theta$ . We maintain a list of the last  $N_{order}$  first-order derivatives:  $\{\nabla_{\theta_n}\}_{n=1}^{N_{order}}$ , which preserves information of high-order derivatives (e.g. a second-order derivative can be estimated by the difference between two consecutive first-order derivatives). Let us denote the parameters and their derivatives, often in tensor form, as

---

#### Algorithm 1: Episodic Policy Gradient Training.

---

**Require:** A parametric policy function  $\pi_\theta$  of the main RL algorithm  $PG_\psi(\pi_\theta, G)$  where  $\psi$  is the set of hyperparameters for training  $\pi_\theta$  and  $G$  the empirical return collected by function  $Agent(\pi_\theta)$ .

- 1: Initialize the episodic memory  $M = \emptyset$
  - 2: **for**  $episode = 1, 2, \dots$  **do** {loop over learning episodes}
  - 3:   Initialize a buffer  $D = \emptyset$  {storing hyper-state, action, and reward within a learning episode}
  - 4:   **for**  $i = 1, \dots, U$  **do** {loop over policy updates}
  - 5:     Compute  $\phi(\mathbf{s}_i^\psi)$ . Select  $\mathbf{a}_i^\psi$  by  $\epsilon$ -greedy with  $Q(\mathbf{s}_i^\psi, \mathbf{a}^\psi) = M.read(\phi(\mathbf{s}_i^\psi), \mathbf{a}^\psi)$  (Eq. 3)
  - 6:     Convert  $\mathbf{a}_i^\psi$  to the hyperparameter values  $\psi_i$  and update  $\theta \leftarrow PG_{\psi_i}(\pi_\theta, G)$
  - 7:     Compute  $\mathbf{r}_i^\psi$  (Eq. 1). Add  $(\phi(\mathbf{s}_i^\psi), \mathbf{a}_i^\psi, \mathbf{r}_i^\psi)$  to  $D$
  - 8:     **if**  $i == U$  **then**  $G = Agent(\pi_\theta)$
  - 9:   **end for**
  - 10:   Update episodic memory with  $M.update(D)$  (Eq. 4)
  - 11: **end for**
- 

$\theta = \{W_m^0\}_{m=1}^M$  and  $\nabla_{\theta_n} = \{W_m^n\}_{m=1}^M$  where  $M$  is the number of layers in the policy/value network.  $\{\theta, \nabla_{\theta_n}\}$  can be denoted jointly  $\{W_m^n\}_{n=0, m=1}^M$  or  $\{W_m^n\}$  for short (see Appendix B.4 for dimension details of  $W_m^n$ ).

Merely using  $\{W_m^n\}$  to represent the learning state is still challenging since the number of parameters is enormous as it often is in the case of recent PG methods. To make the hyper-state tractable, we propose to use linear transformation to map the tensors to lower-dimensional features and concatenate them to create the state vector  $\mathbf{s}^\psi = [s_m^n]_{n=0, m=1}^{N_{order}, M}$ . Here,  $s_m^n$  is the feature of  $W_m^n$ , computed as

$$s_m^n = \text{vec}(W_m^n C_m^n) \quad (2)$$

where  $C_m^n \in \mathbb{R}^{d^{nm} \times d}$  is the transformation matrix,  $d^{nm}$  the last dimension of  $W_m^n$  ( $d^{nm} \gg d$ ) and  $\text{vec}(\cdot)$  the vectorize operator, flattening the input tensor. To make our representation robust, we propose to learn the transformation  $C_m^n$  as described in the next section.

**Learning to represent hyper-state and memory key** We map  $\mathbf{s}^\psi$  to its embedding by using a feed-forward neural network  $\phi$ , resulting in the state embedding  $\phi(\mathbf{s}^\psi) \in \mathbb{R}^h$ .  $\phi(\mathbf{s}^\psi)$  later will be stored as the key of the episodic memory. We can just use random  $\phi$  and  $C_m^n$  for simplicity. However, to encourage  $\phi(\mathbf{s}^\psi)$  to store meaningful information of  $\mathbf{s}^\psi$ , we propose to reconstruct  $\mathbf{s}^\psi$  from  $\phi(\mathbf{s}^\psi)$  via another decoder network  $\omega$  and minimize the following reconstruction error  $\mathcal{L}_{rec} = \|\omega(\phi(\mathbf{s}^\psi)) - \mathbf{s}^\psi\|_2^2$ . Similar to (Blundell et al. 2016), we employ latent-variable probabilistic models such as VAE to learn  $C_m^n$  and update the encoder-decoder networks. Thanks to using  $C_m^n$  projection to lower dimensional space, the hyper-state distribution becomes simpler and potential for VAE reconstruction. Notably, the VAE is jointly trained online with the RL agent and the episodic memory (more details in Appendix A.3).




Figure 2: Performance on (a) MountainCarContinuous (log scale) and (b) BipedalWalker over env. steps. In each plot, average return is on the left with mean and std. over 10 runs. The right is smoothed (taking average over a window of 100 steps) learning rate  $\alpha$  found by the baselines (first 3 runs).

### Episodic Control for Solving the Hyper-RL

Theoretically, given the hyper-state, hyper-action and hyper-reward clearly defined in the previous section, we can use any RL algorithm to solve the Hyper-RL problem. However, in practice, the hyper-reward is usually sparse and the number of steps of the Hyper-RL is usually much smaller than that of the main RL algorithm ( $U \ll T$ ). It means parametric methods (e.g. DQN) which require a huge number of update steps are not suitable for learning a good approximation of the Hyper-RL’s Q-value function  $Q(\mathbf{s}_i^\psi, \mathbf{a}_i^\psi)$ .

To quickly estimate  $Q(\mathbf{s}_i^\psi, \mathbf{a}_i^\psi)$ , we maintain an episodic memory that lasts across learning episodes and stores the outcomes of selecting hyperparameters from a given hyper-state. We hypothesize that the training process involves hyper-states that share similarities, which is suitable for episodic recall using KNN memory lookup. Concretely, the episodic memory  $M$  binds the learning experience  $(\phi(\mathbf{s}_i^\psi), \mathbf{a}_i^\psi)$ —the key, where  $\phi$  is an embedding function, to the approximated expected hyper-return  $\tilde{G}_i^\psi$ —the value. We index the memory using key  $(\phi(\mathbf{s}_i^\psi), \mathbf{a}_i^\psi)$  to access the value, a.k.a  $M[\phi(\mathbf{s}_i^\psi), \mathbf{a}_i^\psi] = \tilde{G}_i^\psi$ . Computing and updating the  $Q(\mathbf{s}_i^\psi, \mathbf{a}_i^\psi)$  corresponds to two memory operators: read and update. The read  $(\phi(\mathbf{s}_i^\psi), \mathbf{a}_i^\psi)$  takes the hyper-state embedding plus hyper-action and returns the hyper-state-action value  $Q(\mathbf{s}_i^\psi, \mathbf{a}_i^\psi)$ . The update (D) takes a buffer  $D$  containing observations  $(\mathbf{s}_i^\psi, \mathbf{a}_i^\psi, \mathbf{r}_i^\psi)_{i=1}^U$ , and updates the content of the memory  $M$ . The details of the two operators are as follows.

**Memory reading** Similarly to (Pritzel et al. 2017), we estimate the state-action value of any  $\mathbf{s}_i^\psi$ - $\mathbf{a}_i^\psi$  pair by:

$$\begin{aligned}
 Q(\mathbf{s}_i^\psi, \mathbf{a}_i^\psi) &= \text{read}(\mathbf{s}_i^\psi, \mathbf{a}_i^\psi) \\
 &= \frac{\sum_{k=1}^{|\mathcal{N}(i)|} \text{Sim}(i, k) M[\phi(\mathbf{s}_k^\psi), \mathbf{a}_i^\psi]}{\sum_{k=1}^{|\mathcal{N}(i)|} \text{Sim}(i, k)}
 \end{aligned} \quad (3)$$

where  $\mathcal{N}(i)$  denotes the neighbor set of the embedding  $\phi(\mathbf{s}_i^\psi)$  in  $M$  and  $\phi(\mathbf{s}_k^\psi)$  the  $k$ -th nearest neighbor.  $\mathcal{N}(i)$  includes  $\phi(\mathbf{s}_i^\psi)$  if it exists in  $M$ .  $\text{Sim}(i, k)$  is a kernel measuring the similarity between  $\phi(\mathbf{s}_i^\psi)$  and  $\phi(\mathbf{s}_k^\psi)$ .

**Memory update** To cope with noisy observations from the Hyper-RL, we propose to use weighted average to write the hyper-return to the memory slots. Unlike max writing rule (Blundell et al. 2016) that always stores the best return, our writing propagates the average return inside the memory, which helps cancel out the noise of the Hyper-RL. In particular, for each observed transition in a learning episode (stored in the buffer  $D$ ), we compute the hyper-return  $G_i^\psi$ . The hyper-return is then used to update the memory such that the action value of  $\phi(\mathbf{s}_i^\psi)$ ’s neighbors is adjusted towards  $G_i^\psi$  with speeds relative to the distances (Le et al. 2021):

$$M[\phi(\mathbf{s}_k^\psi), \mathbf{a}_i^\psi] \leftarrow M[\phi(\mathbf{s}_k^\psi), \mathbf{a}_i^\psi] + \beta \frac{\Delta_{ik} \text{Sim}(i, k)}{\sum_{k=1}^{|\mathcal{N}(i)|} \text{Sim}(i, k)} \quad (4)$$

where  $\phi(\mathbf{s}_k^\psi)$  is the  $k$ -th nearest neighbor of  $\phi(\mathbf{s}_i^\psi)$  in  $\mathcal{N}(i)$ ,  $\Delta_{ik} = G_i^\psi - M[\phi(\mathbf{s}_k^\psi), \mathbf{a}_i^\psi]$ , and  $0 < \beta < 1$  the writing rate. If the key  $(\phi(\mathbf{s}_i^\psi), \mathbf{a}_i^\psi)$  is not in  $M$ , we also add  $(\phi(\mathbf{s}_i^\psi), \mathbf{a}_i^\psi, G_i^\psi)$  to the memory. When the stored tuples exceed memory capacity  $N_{mem}$ , the earliest added tuple will be removed.

Under this formulation,  $M[\phi(\mathbf{s}_i^\psi), \mathbf{a}_i^\psi]$  is an approximation of the expected hyper-return collected by taking the hyper-action  $\mathbf{a}_i^\psi$  at the hyper-state  $\mathbf{s}_i^\psi$  (see Appendix C for proof). As we update several neighbors at one write, the hyper-return propagation inside the episodic memory is faster and helps to handle the sparsity of the Hyper-RL. Unless stated otherwise, we use the same neighbor size  $|\mathcal{N}(i)|$  for both reading and writing process, denoted as  $K$  for short.

**Integration with PG methods** Our episodic control

mechanisms can be used to estimate the hyper-state-action-value of the Hyper-RL. The hyper-agent uses that value to select the hyper-action through  $\epsilon$ -greedy policy and schedule the hyperparameters of PG methods. Algo. 1, Episodic Policy Gradient Training (EPGT), depicts the use of our episodic control with a generic PG method. Our code can be found at <https://github.com/thaihungle/EPGT>

## Experimental Results

Across experiments, we examine EPGT with different PG methods including A2C, ACKTR and PPO. We benchmark EPGT against the original PG methods with tuned hyperparameters and 4 recent hyperparameter search methods. The experimental details can be found in the Appendix B.

### Why Episodic Control?

In this section, we validate the choice of episodic control to solve the proposed Hyper-RL problem. As such, we choose A2C as the PG method and examine EPGT, random hyper-action (RND) and DQN (Mnih et al. 2015) as 3 methods to schedule the learning rate ( $\alpha$ ) for A2C. We also compare with A2C using different fixed- $\alpha$  within the search range (default, min and max learning rates). We test on 2 environments: Mountain Car Continuous (MCC) and Bipedal Walker (BW) with long and short learning rate search ranges ( $[4 \times 10^{-5}, 10^{-2}]$  and  $[2.8 \times 10^{-4}, 1.8 \times 10^{-3}]$ , respectively).

Fig. 2 demonstrates the learning curves and learning rate schedules found by EPGT, RND and DQN. In MCC, the search range is long, which makes RND performance unstable, far lower than the fixed- $\alpha$  A2Cs. DQN also struggles to learn good  $\alpha$  schedule for A2C since the number of trained environment steps is only 20,000, which corresponds to only 4,000 steps in the Hyper-RL. This might not be enough to train DQN’s value network and leads to slower learning. On the contrary, EPGT helps A2C achieve the best performance faster than any other baseline. In BW, thanks to shorter search range and large number of training steps, RND and DQN show better results, yet still underperform the best fixed- $\alpha$  A2C. By contrast, EPGT outperforms the best fixed- $\alpha$  A2C by a significant margin, which confirms the benefit of episodic dynamic hyperparameter scheduling.

Besides performance plots, we visualize the selected values of learning rates over training steps for the first 3 runs of each baseline. Interestingly, DQN finds more consistent values, often converging to extreme learning rates, indicating that the DQN mostly selects the same action for any state, which is unreasonable. EPGT, on the other hand, prefers moderate learning rates, which keep changing depending on the state. Compared to random schedules by RND, those found by EPGT have a pattern, either gradually decreasing (MCC) or increasing (BW). In terms of running time, EPGT runs slightly slower than A2C without any scheduler, yet much faster than DQN (see Appendix’s Table 5).

### EPGT vs. Online Hyperparameter Search Methods

Our main baselines are existing methods for dynamic tuning of hyperparameters of policy gradient algorithms, which can be divided into 2 groups: (i) sequential HOOF (Paul, Kurin,

Model	HalfCheetah	Hopper	Ant	Walker
TMG <sup>♠</sup>	1,568	378	950	492
HOOF <sup>♠</sup>	1,523	350	952	467
HOOF <sup>◇</sup>	1,427±293	452±40.7	954±8.57	674±195
EPGT	<b>2,530±1268</b>	<b>603±187</b>	1,083±126	<b>888±425</b>

Table 1: EPGT vs sequential search (A2C as the PG). Bold denotes statistically better results in terms of Cohen effect size  $> 0.5$ . We train agents for 5 million steps and report the mean (and std. if applicable) over 10 runs. <sup>♠</sup> is from Paul, Kurin, and Whiteson (2019) (no std. reported) and <sup>◇</sup> is our run.

Model	BW	LLC	Hopper	IDP
PBT <sup>°</sup>	223	159	1492	8,893
PB2 <sup>°</sup>	276	<b>235</b>	2,346	8,893
PB2 <sup>◇</sup>	280	223	2,156	9,253
EPGT	<b>282</b>	<b>235</b>	<b>3,253</b>	<b>9,322</b>

Table 2: EPGT vs parallel search (PPO as the PG). Following Parker-Holder, Nguyen, and Roberts (2020), we train the PPO agents for 1 million steps and report the best median over 10 runs. <sup>°</sup> denotes the numbers reported in Parker-Holder, Nguyen, and Roberts (2020), and <sup>◇</sup> is our run.

and Whiteson 2019) and Meta-gradient (Xu, van Hasselt, and Silver 2018) and (ii) parallel PBT (Jaderberg et al. 2017) and PB2 (Parker-Holder, Nguyen, and Roberts 2020). We follow the same experimental setting (PG configuration and environment version) and apply our EPGT to the same set of optimized hyperparameters, keeping other hyperparameters as in other baselines. We also rerun the baselines HOOF and PB2 using our codebase to ensure fair comparison. For the first group, the PG method is A2C and only the learning rate is optimized, while for the second group, the PG method is PPO and we optimize 4 hyperparameters (learning rate  $\alpha$ , batch size  $b$ , GAE  $\lambda$  and PPO clip  $\epsilon$ ).

Table 1 reports the mean test performance of EPGT against Tuned Meta-gradient (TMG) and HOOF on 4 Mujoco environments. EPGT demonstrates better results in all 4 tasks where HalfCheetah, Hopper and Walker observe significant gain. Notably, compared to HOOF, EPGT exhibits higher mean and variance, indicating that EPGT can find distinctive solutions, breaking the local optimum bottleneck of other baselines.

Table 2 compares EPGT with PBT and PB2 on corresponding environments and evaluation metrics. In the four tasks used in PB2 paper, EPGT achieves better median best score for 3 tasks while maintaining competitive performance in LLC task. We note that EPGT is jointly trained with the PG methods in a single run and thus, achieves this excellent performance without parallel interactions with the environments as PB2 or PBT. Learning curves of our runs for the above tasks are in Appendix B.3.

### EPGT vs. Grid-Search/Manual Tuning

**Atari** We now examine EPGT on incremental sets of hyperparameters. We adopt 6 standard Atari games and train 2 PG

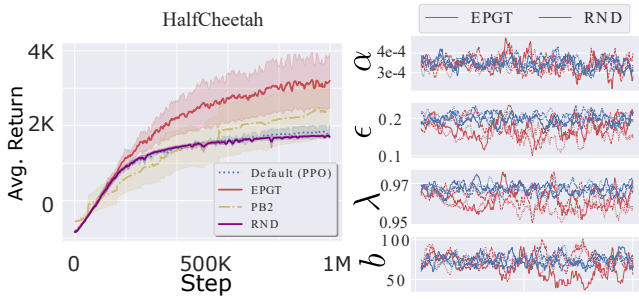


Figure 3: Performance on the representative HalfCheetah task over env. steps. The left is testing return over training iterations (mean  $\pm$  std. over 10 runs) and the right hyperparameters schedule for PG methods found by EPGT and RND in the first 3 runs.

methods: ACKTR and PPO for 20 million steps per game. For ACKTR, we apply EPGT to schedule the trust region radius  $\delta$ , step size  $\eta$  and the value loss coefficient  $l_v$ . For PPO, the optimized hyperparameters are learning rate  $\alpha$ , trust region clip  $\epsilon$  and batch size  $b$ . We form 3 hyperparameter sets for each PG method. For each set, we further perform grid search near the default hyperparameters and record the best tuned results. We compare these results with EPGT’s and report the relative improvement on human normalized score (Appendix Fig. 9). The results indicate that, for all hyperparameter sets, EPGT on average show gains up to more than 10% over tuned PG methods. For certain games, the performance gain can be more than 30%.

**Mujoco** Here, we conduct experiments on 6 Mujoco environments: HalfCheetah, Hooper, Walker2d, Swimmer, Ant and Humanoid. For the last two challenging tasks, we train with 10M steps while the others 1M steps. The set of optimized hyperparameters are  $\{\alpha, \epsilon, \lambda, b\}$ . The baseline Default (PPO) has fixed hyperparameters, which are well-tuned by previous works, and PB2 uses the same hyperparameter search range as our method. Random hyper-action (RND) baseline is included to see the difference between random and episodic policy in Hyper-RL formulation.

On 6 Mujoco tasks, on average, EPGT helps PPO earn more than 583 score while PB2 fails to clearly outperform the tuned PPO (see more in Appendix Fig. 12). Fig. 3 (left) illustrates the result on HalfCheetah where performance gap between EPGT and other baselines can be clearly seen. Despite using the same search range, PB2 and RND show lower average return. We include the hyperparameters used by EPGT and RND throughout training in Fig. 3 (right). Overall, EPGT’s schedules do not diverge much from the default values, which are already well-tuned. However, we can see a pattern of using smaller hyperparameters during middle phase of training, which aligns with the moments when there are changes in the performance.

### Ablation Studies

In this section, we describe the hyperparameter selection for EPGT used in above the experiments. We note that although EPGT introduces several hyperparameters, it is efficient to

pick reasonable values and keep using them across tasks.

**Learning to represent the hyper-state** The hyper-state is captured by projecting the model’s weights and their gradients to a low-dimensional vectors using  $C_m^n$ . The state is further transformed to the memory’s key using the mapping network  $\phi$ . Here, we validate the choice of using VAE to learn  $C_m^n$  and  $\phi$  by comparing it with random mapping. We use PG A2C and test the two EPGT variants on Mountain Car (MC). Fig. 4 (a, left) demonstrates that EPGT with VAE training learns fastest and achieves the best convergence. EPGT with random projections can learn fast but shows similar convergence as the original A2C.

We visualize the final representations  $\phi(\mathbf{s}^\psi)$  by using t-SNE and use colors to denote the corresponding average values  $\hat{V}(\mathbf{s}^\psi) = \sum_a M(\phi(\mathbf{s}^\psi), \mathbf{a}^\psi)$  in Fig. 4 (a, right). The upper figure is randomly projected hyper-states and the lower VAE-trained ones at 5,000 environment step. From both figures, we can see that similar-value states tend to lie together, which validates the hypothesis on existing similar training contexts. Compared to the random ones, the representations learned by VAE exhibit clearer clusters. Cluster separation is critical for nearest neighbor memory access in episodic control, and thus explains why VAE-trained EPGT outperforms random EPGT significantly. Notably, training the VAE is inexpensive. Empirical results demonstrates that with reasonable hyper-state sizes, the VAE converges quickly (see Appendix Fig. 6).

**Writing rule** To verify the contribution of our proposed writing rule, we test different number of writing neighbor size ( $K_w$ ). In this experiment, the reading size is fixed to 3 and different from the writing size. Fig. 4 (b) shows the learning curves of EPGT using different  $K_w$  against the original PPO. When  $K_w = 1$ , our rule becomes single-slot writing as in (Blundell et al. 2016; Pritzel et al. 2017), which even underperforms using default hyperparameters. By contrast, increasing  $K_w = 3$  boosts EPGT’s performance dramatically, on average improving PPO by around 500 score in Alien game. Increasing  $K_w$  further seems not helpful since it may create noise in writing. Thus, we use  $K_w = 3$  in all of our experiments. Others showing our average writing rule is better than traditional max rule and examining different numbers of general neighbor size  $K$  are in Appendix B.4.

**Order of representation** Finally, we examine EPGT’s performance with different order of representation ( $N_{order}$ ).  $N_{order} = 0$  means the hyper-state only includes the parameters  $\theta$ . Increasing  $N_{order}$  gives more information, providing better state representations. That holds true in MsPacman game when we increase the order from 0 to 2 as shown in Fig. 4 (c). However, when  $N_{order}$  is set to 4, the performance drops since the hyper-states now are in a very high dimension (16K) and VAE does not work well in this case. Hence, we use  $N_{order} = 2$  for all of our experiments.

### Related Works

**Hyperparameter search** Automatic hyperparameter tuning generally requires multiple training runs. Parallel search methods such as grid or random search (Bergstra and Bengio 2012; Larochelle et al. 2007) perform multiple runs concurrently and pick the hyperparameters that achieve best re-

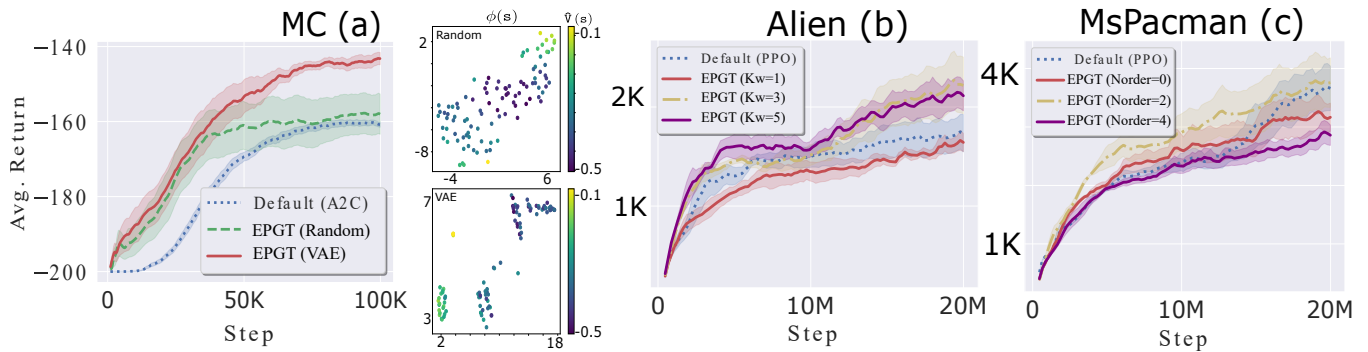


Figure 4: (a) Performance (left) and hyper-state representations  $\phi(s^\psi)$  (right) on Mountain Car (MC) using PG A2C where t-SNE is used to project  $\phi(s^\psi)$  to 2d space, showing the quality of representation learned by VAE (below) versus Random mapping (above). Performance on Alien (b) and MsPacman (c) using PG PPO with different  $K_w$  and  $N_{order}$ , respectively. The curves are mean and std. over 5 runs.

sult. These methods are simple yet expensive. Sequential search approaches reduce the number of runs by consecutively executing experiments using a set of candidate hyperparameters and utilize the evaluation result to guide the subsequent choice of candidates (Hutter, Hoos, and Leyton-Brown 2011). Bayesian Optimization approaches (Brochu, Cora, and De Freitas 2010) exploit the previous experimental results to update the posterior of a Bayesian model of hyperparameters. They have been widely used in hyperparameter tuning for various machine learning algorithms including deep learning (Snoek, Larochelle, and Adams 2012). Recently, to speed up the process, distributed versions of BO are also introduced to evaluate in parallel batches of hyperparameter settings (Chen et al. 2018).

However, these approaches still suffer from the issue of computational inefficiency, demanding high computing resources and training time. If applied to RL, they require more environment interactions, which leads to sample inefficiency. The hyperparameters found by these methods are usually fixed, which can be suboptimal (Luketina et al. 2016). Inspired by biological evolution, population-based methods initially start as random search then select best performing hyperparameter instances to generate subsequent hyperparameter candidates (Young et al. 2015).

Recent works propose using evolutionary algorithms to jointly learn the weights and hyperparameters of neural networks under supervised training (Li et al. 2019). In BPT (Jaderberg et al. 2017) as an example, multiple training are executed asynchronously and evaluated periodically. Underperforming models are replaced by better ones whose hyperparameters evolve to explore better configurations. This approach allows hyperparameter scheduling on-the-fly but still requires a large number of parallel runs and are thus unsuitable for machines with small computational budget.

**On-the-fly hyperparameter search for reinforcement learning** Early works on gradient-based hyperparameter search focus on learning rate adjustment (Sutton 1992). The approach has been recently extended to RL by using the meta-gradient of the return function to adjust the hyperparameters such as discount factor or bootstrapping parameter

(Xu, van Hasselt, and Silver 2018). Hence, in this approach, the return needs to be a differentiable function w.r.t the hyperparameters, which cannot extend to any hyperparameter type such as “clip” or policy gradient algorithm.

HOOF (Paul, Kurin, and Whiteson 2019) is an alternative to meta-gradient methods wherein hyperparameter optimization is done via random search and weighted important sampling. The method relies on off-policy estimate of the value of the policy, which is known to have high variance and thus requires enforcing additional KL constraint. The search is also limited to some specific hyperparameters. Population-based approaches have been applied to RL hyperparameter search. These methods become more efficient by utilizing off-policy PG’s samples (Tang and Choromanski 2020) and small-size population (Parker-Holder, Nguyen, and Roberts 2020), showing better results than PBT or BO in RL domains. However, they still suffer from the inherited expensive computation issue of population-based training. All of these prior works do not formulate hyperparameter search as a MDP, bypassing the context of training, which is addressed in this paper.

## Discussion

We introduced Episodic Policy Gradient Training (EPGT), a new approach for online hyperparameter search using episodic memory. Unlike prior works, EPGT formulates the problem as a Hyper-RL and focuses on modeling the training state to utilize episodic experiences. Then, an episodic control with improved writing mechanisms is employed to search for optimal hyperparameters on-the-fly. Our experiments demonstrate that EPGT can augment various PG algorithms to optimize different types of hyperparameters, achieving better results.

## Acknowledgements

This research was partially funded by the Australian Government through the Australian Research Council (ARC). Prof Venkatesh is the recipient of an ARC Australian Laureate Fellowship (FL170100006).

## References

- Bengio, Y. 2000. Gradient-based optimization of hyperparameters. *Neural computation*, 12(8): 1889–1900.
- Bergstra, J.; and Bengio, Y. 2012. Random search for hyperparameter optimization. *Journal of machine learning research*, 13(2).
- Blundell, C.; Uria, B.; Pritzel, A.; Li, Y.; Ruderman, A.; Leibo, J. Z.; Rae, J.; Wierstra, D.; and Hassabis, D. 2016. Model-free episodic control. *arXiv preprint arXiv:1606.04460*.
- Brochu, E.; Cora, V. M.; and De Freitas, N. 2010. A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*.
- Chen, Y.; Huang, A.; Wang, Z.; Antonoglou, I.; Schrittwieser, J.; Silver, D.; and de Freitas, N. 2018. Bayesian optimization in alphago. *arXiv preprint arXiv:1812.06855*.
- Duan, Y.; Chen, X.; Houthoofd, R.; Schulman, J.; and Abbeel, P. 2016. Benchmarking deep reinforcement learning for continuous control. In *International conference on machine learning*, 1329–1338. PMLR.
- Fiszelew, A.; Britos, P.; Ochoa, A.; Merlino, H.; Fernández, E.; and García-Martínez, R. 2007. Finding optimal neural network architecture using genetic algorithms. *Advances in computer science and engineering research in computing science*, 27: 15–24.
- Fujimoto, S.; Hoof, H.; and Meger, D. 2018. Addressing function approximation error in actor-critic methods. In *International Conference on Machine Learning*, 1587–1596. PMLR.
- González, J.; Dai, Z.; Hennig, P.; and Lawrence, N. 2016. Batch Bayesian optimization via local penalization. In *Artificial intelligence and statistics*, 648–657. PMLR.
- Henderson, P.; Islam, R.; Bachman, P.; Pineau, J.; Precup, D.; and Meger, D. 2018. Deep reinforcement learning that matters. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- Hung, C.-C.; Lillicrap, T.; Abramson, J.; Wu, Y.; Mirza, M.; Carnevale, F.; Ahuja, A.; and Wayne, G. 2019. Optimizing agent behavior over long time scales by transporting value. *Nature communications*, 10(1): 1–12.
- Hutter, F.; Hoos, H. H.; and Leyton-Brown, K. 2011. Sequential model-based optimization for general algorithm configuration. In *International conference on learning and intelligent optimization*, 507–523. Springer.
- Hutter, F.; Kotthoff, L.; and Vanschoren, J. 2019. *Automated machine learning: methods, systems, challenges*. Springer Nature.
- Jaderberg, M.; Dalibard, V.; Osindero, S.; Czarnecki, W. M.; Donahue, J.; Razavi, A.; Vinyals, O.; Green, T.; Dunning, I.; Simonyan, K.; et al. 2017. Population based training of neural networks. *arXiv preprint arXiv:1711.09846*.
- Klein, A.; Falkner, S.; Bartels, S.; Hennig, P.; and Hutter, F. 2017. Fast bayesian optimization of machine learning hyperparameters on large datasets. In *Artificial Intelligence and Statistics*, 528–536. PMLR.
- Kohl, N.; and Stone, P. 2004. Policy gradient reinforcement learning for fast quadrupedal locomotion. In *IEEE International Conference on Robotics and Automation, 2004. Proceedings. ICRA'04. 2004*, volume 3, 2619–2624. IEEE.
- Kumaran, D.; Hassabis, D.; and McClelland, J. L. 2016. What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends in cognitive sciences*, 20(7): 512–534.
- Larochelle, H.; Erhan, D.; Courville, A.; Bergstra, J.; and Bengio, Y. 2007. An empirical evaluation of deep architectures on problems with many factors of variation. In *Proceedings of the 24th international conference on Machine learning*, 473–480.
- Le, H.; George, T. K.; Abdolshah, M.; Tran, T.; and Venkatesh, S. 2021. Model-Based Episodic Memory Induces Dynamic Hybrid Controls. In *Thirty-Fifth Conference on Neural Information Processing Systems*.
- Le, H.; Tran, T.; and Venkatesh, S. 2019. Learning to remember more with less memorization. In *Proceedings of the 7th International Conference on Learning Representations*.
- Le, H.; Tran, T.; and Venkatesh, S. 2020. Self-Attentive Associative Memory. In *Proceedings of the 37th International Conference on Machine Learning*.
- Le, H.; and Venkatesh, S. 2020. Neurocoder: Learning General-Purpose Computation Using Stored Neural Programs. *arXiv preprint arXiv:2009.11443*.
- Lengyel, M.; and Dayan, P. 2008. Hippocampal contributions to control: the third way. In *Advances in neural information processing systems*, 889–896.
- Li, A.; Spyra, O.; Perel, S.; Dalibard, V.; Jaderberg, M.; Gu, C.; Budden, D.; Harley, T.; and Gupta, P. 2019. A generalized framework for population based training. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 1791–1799.
- Luketina, J.; Berglund, M.; Greff, K.; and Raiko, T. 2016. Scalable gradient-based tuning of continuous regularization hyperparameters. In *International conference on machine learning*, 2952–2960. PMLR.
- Mnih, V.; Badia, A. P.; Mirza, M.; Graves, A.; Lillicrap, T.; Harley, T.; Silver, D.; and Kavukcuoglu, K. 2016. Asynchronous methods for deep reinforcement learning. In *International conference on machine learning*, 1928–1937. PMLR.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533.
- Parker-Holder, J.; Nguyen, V.; and Roberts, S. J. 2020. Provably efficient online hyperparameter optimization with population-based bandits. *Advances in Neural Information Processing Systems*, 33.
- Paul, S.; Kurin, V.; and Whiteson, S. 2019. Fast Efficient Hyperparameter Tuning for Policy Gradient Methods. In *Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc,*



F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.

Peters, J.; and Schaal, S. 2006. Policy gradient methods for robotics. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2219–2225. IEEE.

Pritzel, A.; Uria, B.; Srinivasan, S.; Badia, A. P.; Vinyals, O.; Hassabis, D.; Wierstra, D.; and Blundell, C. 2017. Neural episodic control. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 2827–2836. JMLR. org.

Rana, S.; Li, C.; Gupta, S.; Nguyen, V.; and Venkatesh, S. 2017. High dimensional Bayesian optimization with elastic Gaussian process. In *International conference on machine learning*, 2883–2891. PMLR.

Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.

Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. 2017. Mastering the game of go without human knowledge. *nature*, 550(7676): 354–359.

Snoek, J.; Larochelle, H.; and Adams, R. P. 2012. Practical Bayesian optimization of machine learning algorithms. In *Proceedings of the 25th International Conference on Neural Information Processing Systems-Volume 2*, 2951–2959.

Sutton, R. S. 1992. Adapting bias by gradient descent: An incremental version of delta-bar-delta. In *AAAI*, 171–176. San Jose, CA.

Tang, Y.; and Choromanski, K. 2020. Online hyperparameter tuning in off-policy learning via evolutionary strategies. *arXiv preprint arXiv:2006.07554*.

Tulving, E. 2002. Episodic memory: From mind to brain. *Annual review of psychology*, 53(1): 1–25.

Wu, Y.; Mansimov, E.; Liao, S.; Grosse, R.; and Ba, J. 2017. Scalable trust-region method for deep reinforcement learning using Kronecker-factored approximation. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 5285–5294.

Xu, Z.; van Hasselt, H. P.; and Silver, D. 2018. Meta-Gradient Reinforcement Learning. *Advances in Neural Information Processing Systems*, 31: 2396–2407.

Young, S. R.; Rose, D. C.; Karnowski, T. P.; Lim, S.-H.; and Patton, R. M. 2015. Optimizing deep learning hyperparameters through an evolutionary algorithm. In *Proceedings of the Workshop on Machine Learning in High-Performance Computing Environments*, 1–5.

Zhang, B.; Rajan, R.; Pineda, L.; Lambert, N.; Biedenkapp, A.; Chua, K.; Hutter, F.; and Calandra, R. 2021. On the importance of hyperparameter optimization for model-based reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, 4015–4023. PMLR.

Ziegler, D. M.; Stiennon, N.; Wu, J.; Brown, T. B.; Radford, A.; Amodei, D.; Christiano, P.; and Irving, G. 2019. Fine-tuning language models from human preferences. *arXiv preprint arXiv:1909.08593*.