

# LEARNABLE BEHAVIOR CONTROL: BREAKING ATARI HUMAN WORLD RECORDS VIA SAMPLE-EFFICIENT BEHAVIOR SELECTION

Jiajun Fan<sup>1\*</sup>, Yuzheng Zhuang<sup>2</sup>, Yuecheng Liu<sup>2</sup>, Jianye Hao<sup>2</sup>, Bin Wang<sup>2</sup>

Jiangcheng Zhu<sup>3</sup>, Hao Wang<sup>4</sup>, Shutao Xia<sup>1†</sup>

<sup>1</sup> Tsinghua Shenzhen International Graduate School, Tsinghua University

<sup>2</sup> Huawei Noah's Ark Lab, <sup>3</sup> Huawei Cloud, <sup>4</sup> Zhejiang University

<sup>1</sup> fanjj21@mails.tsinghua.edu.cn, xiast@sz.tsinghua.edu.cn, <sup>4</sup> haohaow@zju.edu.cn,

<sup>2,3</sup> {zhuangyuzheng, liuyuecheng1, haojianye, wangbin158, zhujiangcheng}@huawei.com

## ABSTRACT

The exploration problem is one of the main challenges in deep reinforcement learning (RL). Recent promising works tried to handle the problem with population-based methods, which collect samples with diverse behaviors derived from a population of different exploratory policies. Adaptive policy selection has been adopted for behavior control. However, the behavior selection space is largely limited by the predefined policy population, which further limits behavior diversity. In this paper, we propose a general framework called **Learnable Behavioral Control (LBC)** to address the limitation, which a) enables a significantly enlarged behavior selection space via formulating a *hybrid behavior mapping* from all policies; b) constructs a unified *learnable process* for behavior selection. We introduce LBC into distributed off-policy actor-critic methods and achieve behavior control via optimizing the selection of the behavior mappings with bandit-based meta-controllers. Our agents have achieved 10077.52% mean human normalized score and surpassed 24 human world records within 1B training frames in the Arcade Learning Environment, which demonstrates our significant state-of-the-art (SOTA) performance without degrading the sample efficiency.

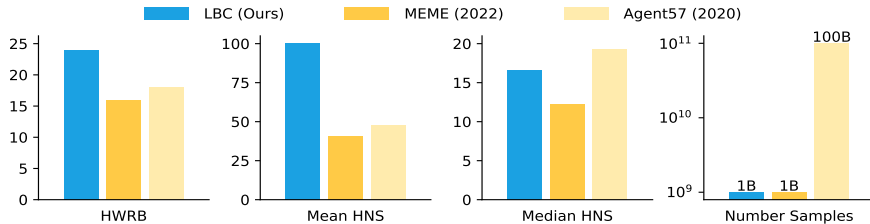


Figure 1: Performance on the 57 Atari games. Our method achieves the highest mean human normalized scores (Badia et al., 2020a), is the first to breakthrough 24 human world records (Toromanoff et al., 2019), and demands the least training data.

## 1 INTRODUCTION

Reinforcement learning (RL) has led to tremendous progress in a variety of domains ranging from video games (Mnih et al., 2015) to robotics (Schulman et al., 2015; 2017). However, efficient exploration remains one of the significant challenges. Recent prominent works tried to address the problem with population-based training (Jaderberg et al., 2017, PBT) wherein a population of policies with different degrees of exploration is jointly trained to keep both the long-term and short-term exploration capabilities throughout the learning process. A set of actors is created to acquire

\*Work done as a research intern at Huawei Noah's Ark Lab.

†Corresponding authors.

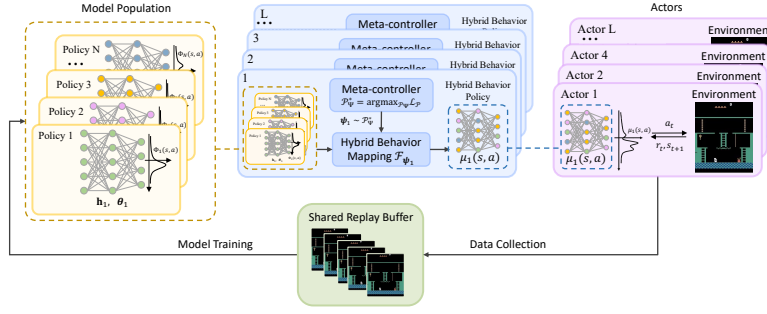


Figure 2: A General Architecture of Our Algorithm.

diverse behaviors derived from the policy population (Badia et al., 2020b;a). Despite the significant improvement in the performance, these methods suffer from the aggravated high sample complexity due to the joint training on the whole population while keeping the diversity property. To acquire diverse behaviors, NGU (Badia et al., 2020b) uniformly selects policies in the population regardless of their contribution to the learning progress (Badia et al., 2020b). As an improvement, Agent57 adopts an adaptive policy selection mechanism that each behavior used for sampling is periodically selected from the population according to a meta-controller (Badia et al., 2020a). Although Agent57 achieved significantly better results on the Arcade Learning Environment (ALE) benchmark, it costs tens of billions of environment interactions as much as NGU. To handle this drawback, GDI (Fan & Xiao, 2022) adaptively combines multiple advantage functions learned from a single policy to obtain an enlarged behavior space without increasing policy population size. However, the population-based scenarios with more than one learned policy has not been widely explored yet. Taking a further step from GDI, we try to enable a larger and non-degenerate behavior space by learning different combinations across a population of different learned policies.

In this paper, we attempt to further improve the sample efficiency of population-based reinforcement learning methods by taking a step towards a more challenging setting to control behaviors with significantly enlarged behavior space with a population of different learned policies. Differing from all of the existing works where each behavior is derived from a single selected learned policy, we formulate the process of getting behaviors from all learned policies as *hybrid behavior mapping*, and the behavior control problem is directly transformed into selecting appropriate mapping functions. By combining all policies, the behavior selection space increases exponentially along with the population size. As a special case that *population size degrades to one*, diverse behaviors can also be obtained by choosing different behavior mappings. This two-fold mechanism enables tremendous larger space for behavior selection. By properly parameterizing the mapping functions, our method enables a unified learnable process, and we call this general framework **Learnable Behavior Control**.

We use the Arcade Learning Environment (ALE) to evaluate the performance of the proposed methods, which is an important testing ground that requires a broad set of skills such as perception, exploration, and control (Badia et al., 2020a). Previous works use the normalized human score to summarize the performance on ALE and claim superhuman performance (Bellemare et al., 2013). However, the human baseline is far from representative of the best human player, which greatly underestimates the ability of humanity. In this paper, we introduce a more challenging baseline, *i.e.*, the human world records baseline (see Toromanoff et al. (2019); Hafner et al. (2021) for more information on Atari human world records). We summarize the number of games that agents can outperform the human world records (*i.e.*, HWRB, see Figs. 1) to claim a real superhuman performance in these games, inducing a more challenging and fair comparison with human intelligence. Experimental results show that the sample efficiency of our method also outperforms the concurrent work MEME Kapturowski et al. (2022), which is 200x faster than Agent57. In summary, our contributions are as follows:

1. **A data-efficient RL framework named LBC.** We propose a general framework called **Learnable Behavior Control (LBC)**, which enables a significantly enlarged behavior selection space without increasing the policy population size via formulating a hybrid behavior mapping from all policies, and constructs a unified learnable process for behavior selection.
2. **A family of LBC-based RL algorithms.** We provide a family of LBC-based algorithms by combining LBC with existing distributed off-policy RL algorithms, which shows the generality and scalability of the proposed method.
3. **The state-of-the-art performance with superior sample efficiency.** From Figs. 1, our method has achieved 10077.52% mean human normalized score (HNS) and surpassed 24

human world records within 1B training frames in the Arcade Learning Environment (ALE), which demonstrates our state-of-the-art (SOTA) sample efficiency.

## 2 BACKGROUND

### 2.1 REINFORCEMENT LEARNING

RL can be formulated as a Markov Decision Process (Howard, 1960, MDP) defined by  $(\mathcal{S}, \mathcal{A}, p, r, \gamma, \rho_0)$ . Considering a discounted episodic MDP, the initial state  $s_0$  is sampled from the initial distribution  $\rho_0(s) : \mathcal{S} \rightarrow \mathbb{P}(\mathcal{S})$ , where we use  $\mathbb{P}$  to represent the probability distribution. At each time  $t$ , the agent chooses an action  $a_t \in \mathcal{A}$  according to the policy  $\pi(a_t|s_t) : \mathcal{S} \rightarrow \mathbb{P}(\mathcal{A})$  at state  $s_t \in \mathcal{S}$ . The environment receives  $a_t$ , produces the reward  $r_t \sim r(s, a) : \mathcal{S} \times \mathcal{A} \rightarrow \mathbf{R}$  and transfers to the next state  $s_{t+1}$  according to the transition distribution  $p(s' | s, a) : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{P}(\mathcal{S})$ . The process continues until the agent reaches a terminal state or a maximum time step. Define the discounted state visitation distribution as  $d_{\rho_0}^{\pi}(s) = (1 - \gamma)\mathbb{E}_{s_0 \sim \rho_0} [\sum_{t=0}^{\infty} \gamma^t \mathbf{P}(s_t = s | s_0)]$ . Define return  $G_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k}$  wherein  $\gamma \in (0, 1)$  is the discount factor. The goal of reinforcement learning is to find the optimal policy  $\pi^*$  that maximizes the expected sum of discounted rewards  $G_t$ :

$$\pi^* := \operatorname{argmax}_{\pi} \mathbb{E}_{s_t \sim d_{\rho_0}^{\pi}} \mathbb{E}_{\pi} \left[ G_t = \sum_{k=0}^{\infty} \gamma^k r_{t+k} | s_t \right], \quad (1)$$

### 2.2 BEHAVIOR CONTROL FOR REINFORCEMENT LEARNING

In value-based methods, a behavior policy can be derived from a state-action value function  $Q_{\theta, \mathbf{h}}^{\pi}(s, a)$  via  $\epsilon$ -greedy. In policy-based methods, a behavior policy can be derived from the policy logits  $\Phi_{\theta, \mathbf{h}}$  (Li et al., 2018) via Boltzmann operator. For convenience, we define that a behavior policy can be derived from the learned policy model  $\Phi_{\theta, \mathbf{h}}$  via a behavior mapping, which normally maps a single policy model to a behavior, *e.g.*,  $\epsilon$ -greedy( $\Phi_{\theta, \mathbf{h}}$ ). In PBT-based methods, there would be a set of policy models  $\{\Phi_{\theta_1, \mathbf{h}_1}, \dots, \Phi_{\theta_N, \mathbf{h}_N}\}$ , each of which is parameterized by  $\theta_i$  and trained under its own hyper-parameters  $\mathbf{h}_i$ , wherein  $\theta_i \in \Theta = \{\theta_1, \dots, \theta_N\}$  and  $\mathbf{h}_i \in \mathbf{H} = \{\mathbf{h}_1, \dots, \mathbf{h}_N\}$ .

The behavior control in population-based methods is normally achieved in two steps: i) select a policy model  $\Phi_{\theta, \mathbf{h}}$  from the population. ii) applying a behavior mapping to the selected policy model. When the behavior mapping is rule-based for each actor (*e.g.*,  $\epsilon$ -greedy with rule-based  $\epsilon$ ), the behavior control can be transformed into the policy model selection (See Proposition 1). Therefore, the optimization of the selection of the policy models becomes one of the critical problems in achieving effective behavior control. Following the literature on PBRL, NGU adopts a uniform selection, which is unoptimized and inefficient. Built upon NGU, Agent57 adopts a meta-controller to adaptively selected a policy model from the population to generate the behavior for each actor, which is implemented by a non-stationary multi-arm bandit algorithm. However, the policy model selection requires maintaining a large number of different policy models, which is particularly data-consuming since each policy model in the population holds heterogeneous training objectives.

To handle this problem, recent notable work GDI-H<sup>3</sup> (Fan & Xiao, 2022) enables to obtain an enlarged behavior space via adaptively controls the temperature of the softmax operation over the weighted advantage functions. However, since the advantage functions are derived from the same target policy under different reward scales, the distributions derived from them may tend to be similar (*e.g.*, See App. N), thus would lead to degradation of the behavior space. Differing from all of the existing works where each behavior is derived from a single selected learned policy, in this paper, we try to handle this problem via three-fold: i) we bridge the relationship between the learned policies and each behavior via a hybrid behavior mapping, ii) we propose a general way to build a non-degenerate large behavior space for population-based methods in Sec. 4.1, iii) we propose a way to optimize the hybrid behavior mappings from a population of different learned models in Proposition. 2.

## 3 LEARNABLE BEHAVIOR CONTROL

In this section, we first formulate the behavior control problem and decouple it into two sub-problems: behavior space construction and behavior selection. Then, we discuss how to construct the behavior

space and select behaviors based on the formulation. By integrating both, we can obtain a general framework to achieve behavior control in RL, called learnable behavior control (LBC).

### 3.1 BEHAVIOR CONTROL FORMULATION

**Behavior Mapping** Define *behavior mapping*  $\mathcal{F}$  as a mapping from some policy model(s) to a behavior. In previous works, a behavior policy is typically obtained using a single policy model. In this paper, as a generalization, we define two kinds of  $\mathcal{F}$  according to how many policy models they take as input to get a behavior. The first one, *individual behavior mapping*, is defined as a mapping from a single model to a behavior that is widely used in prior works, *e.g.*,  $\epsilon$ -greedy and Boltzmann Strategy for discrete action space and Gaussian Strategy for continuous action space; And the second one, *hybrid behavior mapping*, is defined to map all policy models to a single behavior, *i.e.*,  $\mathcal{F}(\Phi_{\theta_1, \mathbf{h}_1}, \dots, \Phi_{\theta_N, \mathbf{h}_N})$ . The hybrid behavior mapping enables us to get a hybrid behavior by combining all policies together, which provides a greater degree of freedom to acquire a larger behavior space. For any behavior mapping  $\mathcal{F}_\psi$  parameterized by  $\psi$ , there exists a family of behavior mappings  $\mathcal{F}_\Psi = \{\mathcal{F}_\psi | \psi \in \Psi\}$  that hold the same parametrization form with  $\mathcal{F}_\psi$ , where  $\Psi \subseteq \mathbf{R}^k$  is a parameter set that contains all possible parameter  $\psi$ .

**Behavior Formulation** As described above, in our work, a behavior can be acquired by applying a behavior mapping  $\mathcal{F}_\psi$  to some policy model(s). For the individual behavior mapping case, a behavior can be formulated as  $\mu_{\theta, \mathbf{h}, \psi} = \mathcal{F}_\psi(\Phi_{\theta, \mathbf{h}})$ , which is also the most used case in previous works. As for the hybrid behavior mapping case, a behavior is formulated as  $\mu_{\Theta, \mathbf{H}, \psi} = \mathcal{F}_\psi(\Phi_{\Theta, \mathbf{H}})$ , wherein  $\Phi_{\Theta, \mathbf{H}} = \{\Phi_{\theta_1, \mathbf{h}_1}, \dots, \Phi_{\theta_N, \mathbf{h}_N}\}$  is a policy model set containing all policy models.

**Behavior Control Formulation** Behavior control can be decoupled into two sub-problems: 1) which behaviors can be selected for each actor at each training time, namely the *behavior space construction*. 2) how to select proper behaviors, namely the *behavior selection*. Based on the behavior formulation, we can formulate these sub-problems:

**Definition 3.1** (Behavior Space Construction). *Considering the RL problem that behaviors  $\mu$  are generated from some policy model(s). We can acquire a family of realizable behaviors by applying a family of behavior mappings  $\mathcal{F}_\Psi$  to these policy model(s). Define the set that contains all of these realizable behaviors as the behavior space, which can be formulated as:*

$$\mathbf{M}_{\Theta, \mathbf{H}, \Psi} = \begin{cases} \{\mu_{\theta, \mathbf{h}, \psi} = \mathcal{F}_\psi(\Phi_{\theta, \mathbf{h}}) | \theta \in \Theta, \mathbf{h} \in \mathbf{H}, \psi \in \Psi\}, & \text{for individual behavior mapping} \\ \{\mu_{\Theta, \mathbf{H}, \psi} = \mathcal{F}_\psi(\Phi_{\Theta, \mathbf{H}}) | \psi \in \Psi\}, & \text{for hybrid behavior mapping} \end{cases} \quad (2)$$

**Definition 3.2** (Behavior Selection). *Behavior selection can be formulated as finding a optimal selection distribution  $\mathcal{P}_{\mathbf{M}_{\Theta, \mathbf{H}, \Psi}}^*$  to select the behaviors  $\mu$  from behavior space  $\mathbf{M}_{\Theta, \mathbf{H}, \Psi}$  and maximizing some optimization target  $\mathcal{L}_{\mathcal{P}}$ , wherein  $\mathcal{L}_{\mathcal{P}}$  is the optimization target of behavior selection:*

$$\mathcal{P}_{\mathbf{M}_{\Theta, \mathbf{H}, \Psi}}^* := \underset{\mathcal{P}_{\mathbf{M}_{\Theta, \mathbf{H}, \Psi}}}{\operatorname{argmax}} \mathcal{L}_{\mathcal{P}} \quad (3)$$

### 3.2 BEHAVIOR SPACE CONSTRUCTION

In this section, we further simplify the equation 2, and discuss how to construct the behavior space.

**Assumption 1.** *Assume all policy models share the same network structure, and  $\mathbf{h}_i$  can uniquely index a policy model  $\Phi_{\theta_i, \mathbf{h}_i}$ . Then,  $\Phi_{\theta, \mathbf{h}}$  can be abbreviated as  $\Phi_{\mathbf{h}}$ .*

Unless otherwise specified, in this paper, we assume Assumption 1 holds. Under Assumption 1, the behavior space defined in equation 2 can be simplified as,

$$\mathbf{M}_{\mathbf{H}, \Psi} = \begin{cases} \{\mu_{\mathbf{h}, \psi} = \mathcal{F}_\psi(\Phi_{\mathbf{h}}) | \mathbf{h} \in \mathbf{H}, \psi \in \Psi\}, & \text{for individual behavior mapping} \\ \{\mu_{\mathbf{H}, \psi} = \mathcal{F}_\psi(\Phi_{\mathbf{H}}) | \psi \in \Psi\}, & \text{for hybrid behavior mapping} \end{cases} \quad (4)$$

According to equation 4, four core factors need to be considered when constructing a behavior space: the network structure  $\Phi$ , the form of behavior mapping  $\mathcal{F}$ , the hyper-parameter set  $\mathbf{H}$  and the parameter set  $\Psi$ . Many notable representation learning approaches have explored how to design the

network structure (Chen et al., 2021; Irie et al., 2021), but it is not the focus of our work. In this paper, we do not make any assumptions about the model structure, which means it can be applied to any model structure. Hence, there remains three factors, which will be discussed below.

For cases that behavior space is constructed with *individual behavior mappings*, there are two things to be considered if one want to select a specific behavior from the behavior space: the policy model  $\Phi_{\mathbf{h}}$  and behavior mapping  $\mathcal{F}_{\psi}$ . Prior methods have tried to realize behavior control via selecting a policy model  $\Phi_{\mathbf{h}_i}$  from the population  $\{\Phi_{\mathbf{h}_1}, \dots, \Phi_{\mathbf{h}_N}\}$  (See Proposition 1). The main drawback of this approach is that only one policy model is considered to generate behavior, leaving other policy models in the population unused. In this paper, we argue that we can tackle this problem via *hybrid behavior mapping*, wherein the hybrid behavior is generated based on all policy models.

In this paper, we only consider the case that all of the  $N$  policy models are used for behavior generating, *i.e.*,  $\mu_{\mathbf{H},\psi} = \mathcal{F}_{\psi}(\Phi_{\mathbf{H}})$ . Now there is only one thing to be considered, *i.e.*, the behavior mapping function  $\mathcal{F}_{\psi}$ , and the behavior control problem will be transformed into the optimization of the behavior mapping (See Proposition 2). We also do not make any assumptions about the form of the mapping. As an example, one could acquire a hybrid behavior from all policy models via network distillation, parameter fusion, mixture models, etc.

### 3.3 BEHAVIOR SELECTION

According to equation 4, each behavior can be indexed by  $\mathbf{h}$  and  $\psi$  for individual behavior mapping cases, and when the  $\psi$  is not learned for each actor, the behavior selection can be cast to the selection of  $\mathbf{h}$  (see Proposition 1). As for the hybrid behavior mapping cases, since each behavior can be indexed by  $\psi$ , the behavior selection can be cast into the selection of  $\psi$  (see Proposition 2). Moreover, according to equation 3, there are two keys in behavior selection: **1)** Optimization Target  $\mathcal{L}_{\mathcal{P}}$ . **2)** The optimization algorithm to learn the selection distribution  $\mathcal{P}_{\mathbf{M}_{\mathbf{H}},\Psi}$  and maximize  $\mathcal{L}_{\mathcal{P}}$ . In this section, we will discuss them sequentially.

**Optimization Target** Two core factors have to be considered for the optimization target: the diversity-based measurement  $V_{\mu}^{\text{TD}}$  (Eysenbach et al., 2019) and the value-based measurement  $V_{\mu}^{\text{TV}}$  (Parker-Holder et al., 2020). By integrating both, the optimization target can be formulated as:

$$\begin{aligned}\mathcal{L}_{\mathcal{P}} &= \mathcal{R}_{\mu \sim \mathcal{P}_{\mathbf{M}_{\mathbf{H}},\Psi}} + c \cdot \mathcal{D}_{\mu \sim \mathcal{P}_{\mathbf{M}_{\mathbf{H}},\Psi}} \\ &= \mathbb{E}_{\mu \sim \mathcal{P}_{\mathbf{M}_{\mathbf{H}},\Psi}} [V_{\mu}^{\text{TV}} + c \cdot V_{\mu}^{\text{TD}}],\end{aligned}\quad (5)$$

wherein,  $\mathcal{R}_{\mu \sim \mathcal{P}_{\mathbf{M}_{\mathbf{H}},\Psi}}$  and  $\mathcal{D}_{\mu \sim \mathcal{P}_{\mathbf{M}_{\mathbf{H}},\Psi}}$  is the expectation of value and diversity of behavior  $\mu$  over the selection distribution  $\mathcal{P}_{\mathbf{M}_{\mathbf{H}},\Psi}$ . When  $\mathcal{F}_{\psi}$  is *unlearned* and *deterministic* for each actor, behavior selection for *each actor* can be simplified into the selection of the policy model:

**Proposition 1** (Policy Model Selection). *When  $\mathcal{F}_{\psi}$  is a deterministic and individual behavior mapping for each actor at each training step (wall-clock), e.g., **Agent57**, the behavior for each actor can be uniquely indexed by  $\mathbf{h}$ , so equation 5 can be simplified into*

$$\mathcal{L}_{\mathcal{P}} = \mathbb{E}_{\mathbf{h} \sim \mathcal{P}_{\mathbf{H}}} [V_{\mu_{\mathbf{h}}}^{\text{TV}} + c \cdot V_{\mu_{\mathbf{h}}}^{\text{TD}}], \quad (6)$$

where  $\mathcal{P}_{\mathbf{H}}$  is a selection distribution of  $\mathbf{h} \in \mathbf{H} = \{\mathbf{h}_1, \dots, \mathbf{h}_N\}$ . For each actor, the behavior is generated from a selected policy model  $\Phi_{\mathbf{h}_i}$  with a pre-defined behavior mapping  $\mathcal{F}_{\psi}$ .

In Proposition 1, the behavior space size is controlled by the policy model population size (*i.e.*,  $|\mathbf{H}|$ ). However, maintaining a large population of different policy models is data-consuming. Hence, we try to control behaviors via optimizing the selection of behavior mappings:

**Proposition 2** (Behavior Mapping Optimization). *When all the policy models are used to generate each behavior, e.g.,  $\mu_{\psi} = \mathcal{F}_{\psi}(\Phi_{\theta,\mathbf{h}})$  for single policy model cases or  $\mu_{\psi} = \mathcal{F}_{\psi}(\Phi_{\theta_1,\mathbf{h}_1}, \dots, \Phi_{\theta_N,\mathbf{h}_N})$  for  $N$  policy models cases, each behavior can be uniquely indexed by  $\mathcal{F}_{\psi}$ , and equation 5 can be simplified into:*

$$\mathcal{L}_{\mathcal{P}} = \mathbb{E}_{\psi \sim \mathcal{P}_{\Psi}} [V_{\mu_{\psi}}^{\text{TV}} + c \cdot V_{\mu_{\psi}}^{\text{TD}}], \quad (7)$$

where  $\mathcal{P}_{\Psi}$  is a selection distribution of  $\psi \in \Psi$ .

In Proposition 2, the behavior space is majorly controlled by  $|\Psi|$ , which could be a continuous parameter space. Hence, a larger behavior space can be enabled.

**Selection Distribution Optimization** Given the optimization target  $\mathcal{L}_{\mathcal{P}}$ , we seek to find the optimal behavior selection distribution  $\mathcal{P}_{\mu}^*$  that maximizes  $\mathcal{L}_{\mathcal{P}}$ :

$$\begin{aligned} \mathcal{P}_{\mathbf{M}, \Psi}^* &:= \operatorname{argmax}_{\mathcal{P}_{\mathbf{M}, \Psi}} \mathcal{L}_{\mathcal{P}} \stackrel{(1)}{=} \operatorname{argmax}_{\mathcal{P}_{\mathbf{H}}} \mathcal{L}_{\mathcal{P}} \\ &\stackrel{(2)}{=} \operatorname{argmax}_{\mathcal{P}_{\Psi}} \mathcal{L}_{\mathcal{P}}, \end{aligned} \quad (8)$$

where (1) and (2) hold because we have Proposition 1 and 2, respectively. This optimization problem can be solved with existing optimizers, *e.g.*, evolutionary algorithm (Jaderberg et al., 2017), multi-arm bandits (MAB) (Badia et al., 2020a), etc.

## 4 LBC-BM: A BOLTZMANN MIXTURE BASED IMPLEMENTATION FOR LBC

In this section, we provide an example of improving the behavior control of off-policy actor-critic methods (Espenholt et al., 2018) via optimizing the behavior mappings as Proposition 2. We provide a practical design of hybrid behavior mapping, inducing an implementation of LBC, which we call **Boltzmann Mixture based LBC**, namely **LBC-BM**. By choosing different  $\mathbf{H}$  and  $\Psi$ , we can obtain a family of implementations of **LBC-BM** with different behavior spaces (see Sec. 5.4).

### 4.1 BOLTZMANN MIXTURE BASED BEHAVIOR SPACE CONSTRUCTION

In this section, we provide a general hybrid behavior mapping design including three sub-processes:

**Generalized Policy Selection** In Agent57, behavior control is achieved by selecting a single policy from the policy population at each iteration. Following this idea, we generalize the method to the case where multiple policies can be selected. More specifically, we introduce a importance weights vector  $\omega$  to describe how much each policy will contribute to the generated behavior,  $\omega = [\omega_1, \dots, \omega_N], \omega_i \geq 0, \sum_{i=1}^N \omega_i = 1$ , where  $\omega_i$  represents the importance of  $i$ th policy in the population (*i.e.*,  $\Phi_{\mathbf{h}_i}$ ). In particular, if  $\omega$  is a one-hot vector, *i.e.*,  $\exists i \in \{1, 2, \dots, N\}, \omega_i = 1; \forall j \in \{1, 2, \dots, N\} \neq i, \omega_j = 0$ , then the policy selection becomes a single policy selection as Proposition 1. Therefore, it can be seen as a generalization of single policy selection, and we call this process *generalized policy selection*.

**Policy-Wise Entropy Control** In our work, we propose to use entropy control (which is typically rule-based controlled in previous works) to make a better trade-off between exploration and exploitation. For a policy model  $\Phi_{\mathbf{h}_i}$  from the population, we will apply a entropy control function  $f_{\tau_i}(\cdot)$ , *i.e.*,  $\pi_{\mathbf{h}_i, \tau_i} = f_{\tau_i}(\Phi_{\mathbf{h}_i})$ , where  $\pi_{\mathbf{h}_i, \tau_i}$  is the new policy after entropy control, and  $f_{\tau_i}(\cdot)$  is parameterized by  $\tau_i$ . Here we should note that the entropy of all the policies from the population is controlled in a policy-wise manner. Thus there would be a set of entropy control functions to be considered, which is parameterized by  $\tau = [\tau_1, \dots, \tau_N]$ .

**Behavior Distillation from Multiple Policies** Different from previous methods where only one policy is used to generate the behavior, in our approach, we combine  $N$  policies  $[\pi_{\mathbf{h}_1, \tau_1}, \dots, \pi_{\mathbf{h}_N, \tau_N}]$ , together with their importance weights  $\omega = [\omega_1, \dots, \omega_N]$ . Specially, in order to make full use of these policies according to their importance, we introduce a *behavior distillation function*  $g$  which takes both the policies and importance weights as input, *i.e.*,  $\mu_{\mathbf{H}, \tau, \omega} = g(\pi_{\mathbf{h}_1, \tau_1}, \dots, \pi_{\mathbf{h}_N, \tau_N}, \omega)$ . The distillation function  $g(\cdot, \omega)$  can be implemented in different ways, *e.g.*, knowledge distillation (supervised learning), parameters fusion, etc. In conclusion, the behavior space can be constructed as,

$$\mathbf{M}_{\mathbf{H}, \Psi} = \{g(f_{\tau_1}(\Phi_{\mathbf{h}_1}), \dots, f_{\tau_N}(\Phi_{\mathbf{h}_N}), \omega_1, \dots, \omega_N) \mid \psi \in \Psi\} \quad (9)$$

wherein  $\Psi = \{\psi = (\tau_1, \dots, \tau_N, \omega_1, \dots, \omega_N)\}$ ,  $\mathbf{H} = \{\mathbf{h}_1, \dots, \mathbf{h}_N\}$ . Note that this is a general approach which can be applied to different tasks and algorithms by simply selecting different entropy control function  $f_{\tau_i}(\cdot)$  and behavior distillation function  $g(\cdot, \omega)$ . As an example, for Atari task, we model the policy as a Boltzmann distribution, *i.e.*,  $\pi_{\mathbf{h}_i, \tau_i}(a|s) = e^{\tau_i \Phi_{\mathbf{h}_i}(a|s)} \sum_{a'} e^{\tau_i \Phi_{\mathbf{h}_i}(a'|s)}$ , where  $\tau_i \in (0, \infty)$ . The entropy can thus be controlled by controlling the temperature. As for the behavior distillation function, we are inspired by the behavior design of GDI, which takes a weighted sum of two softmax distributions derived from two advantage functions. We can further extend this approach

to the case to do a combination across different policies, *i.e.*,  $\mu_{\mathbf{H},\tau,\omega}(a|s) = \sum_{i=1}^N \omega_i \pi_{\mathbf{h}_i,\tau_i}(a|s)$ . This formula is actually a form of *mixture model*, where the importance weights play the role of mixture weights of the mixture model. Then the behavior space becomes,

$$\mathbf{M}_{\mathbf{H},\Psi} = \{\mu_{\mathbf{H},\psi} = \sum_{i=1}^N \omega_i \text{softmax}_{\tau_i}(\Phi_{\mathbf{h}_i}) | \psi \in \Psi\} \quad (10)$$

## 4.2 MAB BASED BEHAVIOR SELECTION

According to Proposition 2, the behavior selection over behavior space 10 can be simplified to the selection of  $\psi$ . In this paper, we use MAB-based meta-controller to select  $\psi \in \Psi$ . Since  $\Psi$  is a continuous multidimensional space, we discretize  $\Psi$  into  $K$  regions  $\{\Psi_1, \dots, \Psi_K\}$ , and each region corresponds to an arm of MAB. At the beginning of a trajectory  $i$ ,  $l$ -th actor will use MAB to sample a region  $\Psi_k$  indexed by arm  $k$  according to  $\mathcal{P}_{\Psi} = \text{softmax}(\text{Score}_{\Psi_k}) = \frac{e^{\text{Score}_{\Psi_k}}}{\sum_j e^{\text{Score}_{\Psi_j}}}$ . We adopt UCB

score as  $\text{Score}_{\Psi_k} = V_{\Psi_k} + c \cdot \sqrt{\frac{\log(1 + \sum_{j \neq k}^K N_{\Psi_j})}{1 + N_{\Psi_k}}}$  to tackle the reward-diversity trade-off problem in equation 7 (Garivier & Moulines, 2011).  $N_{\Psi_k}$  means the number of the visit of  $\Psi_k$  indexed by arm  $k$ .  $V_{\Psi_k}$  is calculated by the expectation of the undiscounted episodic returns to measure the value of each  $\Psi_k$ , and the UCB item is used to avoid selecting the same arm repeatedly and ensure sufficient diverse behavior mappings can be selected to boost the behavior diversity. After an  $\Psi_k$  is sampled, a  $\psi$  will be uniformly sampled from  $\Psi_k$ , corresponding to a behavior mapping  $\mathcal{F}_{\psi}$ . With  $\mathcal{F}_{\psi}$ , we can obtain a behavior  $\mu_{\psi}$  according to equation 10. Then, the  $l$ -th actor acts  $\mu_{\psi}$  to obtain a trajectory  $\tau_i$  and the undiscounted episodic return  $G_i$ , then  $G_i$  is used to update the reward model  $V_{\Psi_k}$  of region  $\Psi_k$  indexed by arm  $k$ . As for the nonstationary problem, we are inspired from GDI, which ensembles several MAB with different learning rates and discretization accuracy. We can extend to handle the nonstationary problem by jointly training a population of bandits from very exploratory to purely exploitative (*i.e.*, different  $c$  of the UCB item, similar to the policy population of Agent57). Moreover, we will periodically replace the members of the MAB population to ease the nonstationary problem further. More details of implementations of MAB can be found in App. E. Moreover, the mechanism of the UCB item for behavior control has not been widely studied in prior works, and we will demonstrate how it boosts behavior diversity in App. K.3.

## 5 EXPERIMENT

In this section, we design our experiment to answer the following questions:

- Whether our methods can outperform prior SOTA RL algorithms in both sample efficiency and final performance in Atari 1B Benchmarks (See Sec. 5.2 and Figs. 3)?
- Can our methods adaptively adjust the exploration-exploration trade-off (See Figs. 4)?
- How to enlarge or narrow down the behavior space? What is the performance of methods with different behavior spaces (See Sec. 5.4)?
- How much performance will be degraded without proper behavior selection (See Figs. 5)?

### 5.1 EXPERIMENTAL SETUP

#### 5.1.1 EXPERIMENTAL DETAILS

We conduct our experiments in ALE (Bellemare et al., 2013). The standard pre-processing settings of Atari are identical to those of Agent57 (Badia et al., 2020a), and related parameters have been concluded in App. I. We employ a separate evaluation process to record scores continuously. We record the undiscounted episodic returns averaged over five seeds using a windowed mean over 32 episodes. To avoid any issues that aggregated metrics may have, App. J provides full learning curves for all games and detailed comparison tables of raw and normalized scores. Apart from the mean and median HNS, we also report how many human worlds records our agents have broken to emphasize the superhuman performance of our methods. For more experimental details, see App. H.

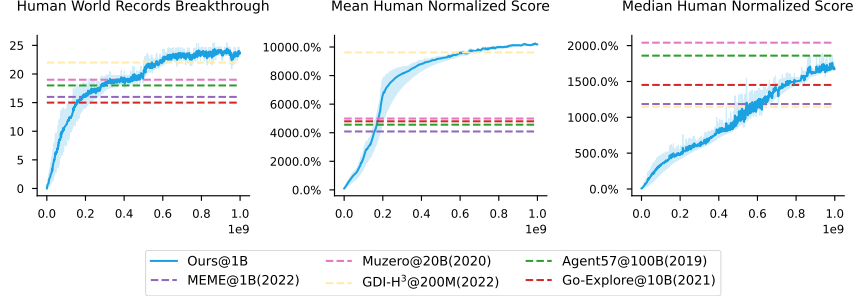


Figure 3: The learning curves in Atari. Curves are smoothed with a moving average over 5 points.

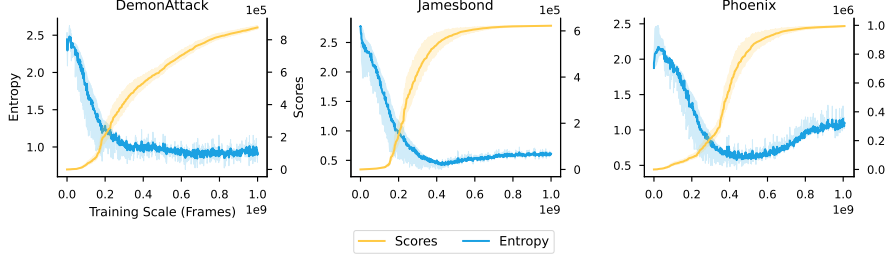


Figure 4: Behavior entropy and scores curve across training for different games where we achieved unprecedented performance. The names of the axes are the same as that of the leftmost figure.

### 5.1.2 IMPLEMENTATION DETAILS

We jointly train three policies, and each policy can be indexed by the hyper-parameters  $\mathbf{h}_i = (\gamma_i, \mathcal{RS}_i)$ , wherein  $\mathcal{RS}_i$  is a reward shaping method (Badia et al., 2020a), and  $\gamma_i$  is the discounted factor. Each policy model  $\Phi_{\mathbf{h}_i}$  adopts the dueling network structure (Wang et al., 2016), where  $\Phi_{\mathbf{h}_i} = A_{\mathbf{h}_i} = Q_{\mathbf{h}_i} - V_{\mathbf{h}_i}$ . More details of the network structure can be found in App. L. To correct for harmful discrepancy of off-policy learning, we adopt V-Trace (Espeholt et al., 2018) and ReTrace (Munos et al., 2016) to learn  $V_{\mathbf{h}_i}$  and  $Q_{\mathbf{h}_i}$ , respectively. The policy is learned by policy gradient (Schulman et al., 2017). Based on equation 10, we could build a behavior space with a hybrid mapping as  $\mathbf{M}_{\mathbf{H}, \Psi} = \{\mu_{\mathbf{H}, \psi} = \sum_{i=1}^3 \omega_i \text{softmax}_{\tau_i}(\Phi_{\mathbf{h}_i})\}$ , wherein  $\mathbf{H} = \{\mathbf{h}_1, \mathbf{h}_2, \mathbf{h}_3\}$ ,  $\Psi = \{\psi = (\tau_1, \omega_1, \tau_2, \omega_2, \tau_3, \omega_3) | \tau_i \in (0, \tau^+), \sum_{j=1}^3 \omega_j = 1\}$ . The behavior selection is achieved by MAB described in Sec. 4.2, and more details can see App. E. Finally, we could obtain an implementation of LBC- $\mathcal{BM}$ , which is our **main algorithm**. The target policy for  $A_1^\pi$  and  $A_2^\pi$  in GDI-H<sup>3</sup> is the same, while in our work the target policy for  $A_i^{\pi_i}$  is  $\pi_i = \text{softmax}(A_i)$ .

## 5.2 SUMMARY OF RESULTS

**Results on Atari Benchmark** The aggregated results across games are reported in Figs. 3. Among the algorithms with superb final performance, our agents achieve the best mean HNS and surpass the most human world records across 57 games of the Atari benchmark with relatively minimal training frames, leading to the best learning efficiency. Noting that Agent57 reported the maximum scores across training as the final score, and if we report our performance in the same manner, our median is 1934%, which is **higher** than Agent57 and demonstrates our superior performance.

**Discussion of Results** With LBC, we can understand the mechanisms underlying the performance of GDI-H<sup>3</sup> more clearly: **i)** GDI-H<sup>3</sup> has a high-capacity behavior space and a meta-controller to optimize the behavior selection **ii)** only a single target policy is learned, which enables stable learning and fast converge (See the case study of KL divergence in App. N). Compared to GDI-H<sup>3</sup>, to ensure the behavior space will not degenerate, LBC maintains a population of diverse policies and, as a price, sacrifices some sample efficiency. Nevertheless, LBC can **continuously** maintain a significantly larger behavior space with hybrid behavior mapping, which enables RL agents to continuously explore and get improvement.

### 5.3 CASE STUDY: BEHAVIOR CONTROL

To further explore the mechanisms underlying the success of behavior control of our method, we adopt a case study to showcase our control process of behaviors. As shown in Figs. 4, in most tasks,



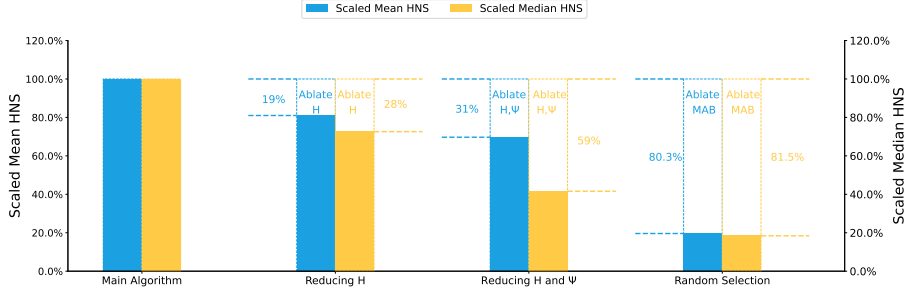


Figure 5: Ablation Results. All the results are scaled by the main algorithm to improve readability.

our agents prefer exploratory behaviors first (i.e., high stochasticity policies with high entropy), and, as training progresses, the agents shift into producing experience from more exploitative behaviors. On the verge of peaking, the entropy of the behaviors could be maintained at a certain level (task-wise) instead of collapsing swiftly to zero to avoid converging prematurely to sub-optimal policies.

#### 5.4 ABLATION STUDY

In this section, we investigate several properties of our method. For more details, see App. K.

**Behavior Space Decomposition** To explore the effect of different behavior spaces, we decompose the behavior space of our main algorithm via reducing  $\mathbf{H}$  and  $\Psi$ :

**1) Reducing  $\mathbf{H}$ .** When we set all the policy models of our main algorithm the same, the behavior space transforms from  $\mathcal{F}(\Phi_{h_1}, \Phi_{h_2}, \Phi_{h_3})$  into  $\mathcal{F}(\Phi_{h_1}, \Phi_{h_1}, \Phi_{h_1})$ .  $\mathbf{H}$  degenerates from  $\{h_1, h_2, h_3\}$  into  $\{h_1\}$ . We can obtain a control group with a smaller behavior space by reducing  $\mathbf{H}$ .

**2) Reducing  $\mathbf{H}$  and  $\Psi$ .** Based on the control group reducing  $\mathbf{H}$ , we can further reduce  $\Psi$  to further narrow down the behavior space. Specially, we can directly adopt a individual behavior mapping to build the behavior space as  $\mathbf{M}_{\mathbf{H}, \Psi} = \{\mu_\psi = \text{softmax}_\tau(\Phi_{h_1})\}$ , where  $\Psi$  degenerates from  $\{\omega_1, \omega_2, \omega_3, \tau_1, \tau_2, \tau_3\}$  to  $\{\tau\}$  and  $\mathbf{H} = \{h_1\}$ . Then, we can obtain a control group with the smallest behavior space by reducing  $\mathbf{H}$  and  $\Psi$ .

The performance of these methods is illustrated in Figs. 5, and from left to right, the behavior space of the first three algorithms decreases in turn (According to Corollary 4 in App. C). It is evident that narrowing the behavior space via reducing  $\mathbf{H}$  or  $\Psi$  will degrade the performance. On the contrary, the performance can be boosted by enlarging the behavior space, which could be a promising way to improve the performance of existing methods.

**Behavior Selection** To highlight the importance of an appropriate behavior selection, we replace the meta-controller of our main algorithm with a random selection. The ablation results are illustrated in Figs. 5, from which it is evident that, with the same behavior space, not learning an appropriate selection distribution of behaviors will significantly degrade the performance. We conduct a t-SNE analysis in App. K.3 to demonstrate that our methods can acquire more diverse behaviors than the control group with pre-defined behavior mapping. Another ablation study that removed the UCB item has been conducted in App. K.3 to demonstrate the behavior diversity may be boosted by the UCB item, which can encourage the agents to select more different behavior mappings.

## 6 CONCLUSION

We present the first deep reinforcement learning agent to break 24 human world records in Atari using only 1B training frames. To achieve this, we propose a general framework called LBC, which enables a significantly enlarged behavior selection space via formulating a hybrid behavior mapping from all policies, and constructs a unified learnable process for behavior selection. We introduced LBC into off-policy actor-critic methods and obtained a family of implementations. A large number of experiments on Atari have been conducted to demonstrate the effectiveness of our methods empirically. Apart from the full results, we do detailed ablation studies to examine the effectiveness of the proposed components. While there are many improvements and extensions to be explored going forward, we believe that the ability of LBC to enhance the control process of behaviors results in a powerful platform to propel future research.

## ACKNOWLEDGMENTS

This work is majorly supported by the National Key R&D Program of China (Grant Number 2021ZD0110400). In addition, this work is partly supported by the National Natural Science Foundation of China under Grant 62171248, and the R&D Program of Shenzhen under Grant JCYJ20220818101012025.

We are very grateful for the careful reading and insightful reviews of meta-reviewers and reviewers.

## REPRODUCIBILITY STATEMENT

Open-sourced code will be implemented with Mindspore (MS, 2022) and released on our website.

## REFERENCES

- Mindspore, 2022. URL <https://www.mindspore.cn/>. Software available from <https://www.mindspore.cn/>.
- Adrià Puigdomènech Badia, Bilal Piot, Steven Kapturowski, Pablo Sprechmann, Alex Vitvitskyi, Zhaohan Daniel Guo, and Charles Blundell. Agent57: Outperforming the atari human benchmark. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 507–517. PMLR, 2020a. URL <http://proceedings.mlr.press/v119/badia20a.html>.
- Adrià Puigdomènech Badia, Pablo Sprechmann, Alex Vitvitskyi, Zhaohan Daniel Guo, Bilal Piot, Steven Kapturowski, Olivier Tieleman, Martín Arjovsky, Alexander Pritzel, Andrew Bolt, and Charles Blundell. Never give up: Learning directed exploration strategies. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020b. URL <https://openreview.net/forum?id=Sye57xStvB>.
- Marc G. Bellemare, Yavar Naddaf, Joel Veness, and Michael Bowling. The arcade learning environment: An evaluation platform for general agents. *J. Artif. Intell. Res.*, 47:253–279, 2013. doi: 10.1613/jair.3912. URL <https://doi.org/10.1613/jair.3912>.
- Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Misha Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34:15084–15097, 2021.
- Lasse Espeholt, Hubert Soyer, Rémi Munos, Karen Simonyan, Volodymyr Mnih, Tom Ward, Yotam Doron, Vlad Firoiu, Tim Harley, Iain Dunning, Shane Legg, and Koray Kavukcuoglu. IMPALA: scalable distributed deep-rl with importance weighted actor-learner architectures. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1406–1415. PMLR, 2018. URL <http://proceedings.mlr.press/v80/espeholt18a.html>.
- Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you need: Learning skills without a reward function. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=SJx63jRqFm>.
- Jiajun Fan and Changnan Xiao. Generalized data distribution iteration. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (eds.), *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pp. 6103–6184. PMLR, 2022. URL <https://proceedings.mlr.press/v162/fan22c.html>.
- Aurélien Garivier and Eric Moulines. On upper-confidence bound policies for switching bandit problems. In Jyrki Kivinen, Csaba Szepesvári, Esko Ukkonen, and Thomas Zeugmann (eds.), *Algorithmic Learning Theory - 22nd International Conference, ALT 2011, Espoo, Finland, October 5-7, 2011. Proceedings*, volume 6925 of *Lecture Notes in Computer Science*, pp. 174–188.

- Springer, 2011. doi: 10.1007/978-3-642-24412-4\_16. URL [https://doi.org/10.1007/978-3-642-24412-4\\_16](https://doi.org/10.1007/978-3-642-24412-4_16).
- Danijar Hafner, Timothy P. Lillicrap, Mohammad Norouzi, and Jimmy Ba. Mastering atari with discrete world models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=0oabwyZbOu>.
- Matteo Hessel, Joseph Modayil, Hado van Hasselt, Tom Schaul, Georg Ostrovski, Will Dabney, Dan Horgan, Bilal Piot, Mohammad Gheshlaghi Azar, and David Silver. Rainbow: Combining improvements in deep reinforcement learning. In Sheila A. McIlraith and Kilian Q. Weinberger (eds.), *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pp. 3215–3222. AAAI Press, 2018. URL <https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/17204>.
- Ronald A Howard. *Dynamic programming and markov processes*. John Wiley, 1960.
- Kazuki Irie, Imanol Schlag, Róbert Csordás, and Jürgen Schmidhuber. Going beyond linear transformers with recurrent fast weight programmers. *Advances in Neural Information Processing Systems*, 34:7703–7717, 2021.
- Max Jaderberg, Valentin Dalibard, Simon Osindero, Wojciech M. Czarnecki, Jeff Donahue, Ali Razavi, Oriol Vinyals, Tim Green, Iain Dunning, Karen Simonyan, Chrisantha Fernando, and Koray Kavukcuoglu. Population based training of neural networks. *CoRR*, abs/1711.09846, 2017. URL <http://arxiv.org/abs/1711.09846>.
- Steven Kapturowski, Georg Ostrovski, John Quan, Rémi Munos, and Will Dabney. Recurrent experience replay in distributed reinforcement learning. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=r1lyTjAqYX>.
- Steven Kapturowski, Víctor Campos, Ray Jiang, Nemanja Rakićević, Hado van Hasselt, Charles Blundell, and Adrià Puigdomènech Badia. Human-level atari 200x faster. *CoRR*, abs/2209.07550, 2022. URL <https://arxiv.org/abs/2209.07550>.
- Jinke Li, Ruonan Rao, and Jun Shi. Learning to trade with deep actor critic methods. In *2018 11th International Symposium on Computational Intelligence and Design (ISCID)*, volume 2, pp. 66–71. IEEE, 2018.
- Marlos C. Machado, Marc G. Bellemare, Erik Talvitie, Joel Veness, Matthew J. Hausknecht, and Michael Bowling. Revisiting the arcade learning environment: Evaluation protocols and open problems for general agents. *J. Artif. Intell. Res.*, 61:523–562, 2018. doi: 10.1613/jair.5699. URL <https://doi.org/10.1613/jair.5699>.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dhharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep reinforcement learning. *Nat.*, 518(7540):529–533, 2015. doi: 10.1038/nature14236. URL <https://doi.org/10.1038/nature14236>.
- Rémi Munos, Tom Stepleton, Anna Harutyunyan, and Marc G. Bellemare. Safe and efficient off-policy reinforcement learning. In Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, pp. 1046–1054, 2016. URL <https://proceedings.neurips.cc/paper/2016/hash/c3992e9a68c5ae12bd18488bc579b30d-Abstract.html>.
- Jack Parker-Holder, Aldo Pacchiano, Krzysztof Marcin Choromanski, and Stephen J. Roberts. Effective diversity in population based reinforcement learning. In Hugo

- Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/d1dc3a8270a6f9394f88847d7f0050cf-Abstract.html>.
- Sepp Hochreiter; Jürgen Schmidhuber. Long short-term memory. *Neural Computation.*, 1997.
- Simon Schmitt, Matteo Hessel, and Karen Simonyan. Off-policy actor-critic with shared experience replay. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pp. 8545–8554. PMLR, 2020. URL <http://proceedings.mlr.press/v119/schmitt20a.html>.
- John Schulman, Sergey Levine, Pieter Abbeel, Michael I. Jordan, and Philipp Moritz. Trust region policy optimization. In Francis R. Bach and David M. Blei (eds.), *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015*, volume 37 of *JMLR Workshop and Conference Proceedings*, pp. 1889–1897. JMLR.org, 2015. URL <http://proceedings.mlr.press/v37/schulman15.html>.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *CoRR*, abs/1707.06347, 2017. URL <http://arxiv.org/abs/1707.06347>.
- Richard S. Sutton and Andrew G. Barto. *Reinforcement learning - an introduction*. Adaptive computation and machine learning. MIT Press, 1998. ISBN 978-0-262-19398-6. URL <https://www.worldcat.org/oclc/37293240>.
- Marin Toromanoff, Émilie Wirbel, and Fabien Moutarde. Is deep reinforcement learning really superhuman on atari? *CoRR*, abs/1908.04683, 2019. URL <http://arxiv.org/abs/1908.04683>.
- Ziyu Wang, Tom Schaul, Matteo Hessel, Hado van Hasselt, Marc Lanctot, and Nando de Freitas. Dueling network architectures for deep reinforcement learning. In Maria-Florina Balcan and Kilian Q. Weinberger (eds.), *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pp. 1995–2003. JMLR.org, 2016. URL <http://proceedings.mlr.press/v48/wangf16.html>.
- Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Mach. Learn.*, 8:229–256, 1992. doi: 10.1007/BF00992696. URL <https://doi.org/10.1007/BF00992696>.

## Appendix

**Overview of the Appendix** For ease of reading, we submit the main body of our paper together with the appendix as supplemental material. Below we will briefly introduce the structure of our appendix for easy reading. Readers can jump from the corresponding position in the body to the corresponding content in the appendix, or jump from the following table of contents to the corresponding content of interest in the appendix.

- In App. A, we briefly summarize some common notations in this paper for the convenience of readers.
- In App. B, we summarize and recall the background knowledge used in this paper.
- In App. C, we provide theoretical proofs.
- In App. D, we provide several implementations of LBC-based behavior sapce design to facilitate future research.
- In App. E, we provide a detailed implementation of our MAB-based meta-controller.
- In App. F, we provide a detailed implementation of our core framework.
- In App. G, we use our framework to introduce an LBC-based version of well-known RL methods, which leads to a better understanding of their original counterparts.
- In App. H, we provide relevant details of our experiments.
- In App. I, we summarize the hyper-parameters used in our experiments.
- In App. J, we provides full learning curves for all games and detailed comparison tables of raw and normalized scores.
- In App. K, we provides the design of the ablation study and the overall results.
- In App. L, we summarize our model architecture in detail to for reproducibility.

## A SUMMARY OF NOTATION AND ABBREVIATION

In this section, we briefly summarize some common notations in this paper for the convenience of readers, which are concluded in Tab. 1.

Table 1: Summary of Notation

Symbol	Description	Symbol	Description
$s$	State	$\mathcal{S}$	Set of all states
$a$	Action	$\mathcal{A}$	Set of all actions
$\mathbb{P}$	Probability distribution	$\mu$	Behavior policy
$\pi$	Policy	$G_t$	Cumulative discounted reward at $t$
$d_{\rho_0}^\pi$	States visitation distribution of $\pi$	$V_\pi$	State value function of $\pi$
$Q_\pi$	State-action value function of $\pi$	$A_\pi$	Advantage function of $\pi$
$\gamma$	Discount-rate parameter	$\delta_t$	Temporal-difference error at $t$
$\mathcal{F}$	Behavior mapping	$\Phi$	Policy models
$\theta$	Parameters of the policy network	$\Theta$	Set of $\theta$
$\mathbf{h}$	hyper-parameters of policy models	$\mathbf{H}$	Set of $h$
$\psi$	Parameters to index $\mathcal{F}$	$\Psi$	Set of $\psi$
$\mathbf{M}_{\Theta, \mathbf{H}, \Psi}$	Behavior policy sets/space parameterized by $\theta, \mathbf{h}, \psi$	$\mu_{\theta, \mathbf{h}, \psi}$	A behavior policy of $\mathbf{M}_{\Theta, \mathbf{H}, \Psi}$
$ \mathbf{H} $	Size of set $\mathbf{H}$	$ \Psi $	Size of set $\Psi$
$\mathcal{P}_{\mathbf{M}_{\Theta, \mathbf{H}, \Psi}}$	Behavior selection distribution over $\mathbf{M}_{\Theta, \mathbf{H}, \Psi}$	$\mathbb{1}_{\psi=\psi_0}$	One-point distribution $\mathbf{P}(\psi = \psi_0) = 1$
$V_\mu^{\text{TV}}$	Some measurement on the value of policy $\pi$	$V_\mu^{\text{TD}}$	Some measurement on the diversity of policy $\pi$
$\mathbf{M}_{\Theta, \mathbf{H}}$	Subspace of $\mathbf{M}_{\Theta, \mathbf{H}, \Psi}$	$\mathcal{RS}$	Reward shaping

## B BACKGROUND

Most of the previous work has introduced the background knowledge of RL in detail. In this section, we only summarize and recall the background knowledge used in this paper. If you are interested in the relevant content, we recommend you read the relevant material (Badia et al., 2020a; Sutton & Barto, 1998; Badia et al., 2020b; Kapturowski et al., 2019; Espeholt et al., 2018). The relevant notations and abbreviations have been documented in App. A.

### B.1 POLICY GRADIENT

Policy gradient (Williams, 1992, PG) is a policy-based RL methods, which updates the policy via optimizing the following target:

$$\mathcal{J}(\theta) = \mathbb{E}_{\pi} \left[ \sum_{t=0}^{\infty} \log \pi_{\theta}(a_t | s_t) G(\tau) \right]$$

The policy gradient of Actor-Critic (Sutton & Barto, 1998, AC) methods is updated as follows:

$$\nabla_{\theta} \mathcal{J}(\theta) = \mathbb{E}_{\pi} \left[ \sum_{t=0}^{\infty} \Phi_t \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \right]$$

wherein  $\Phi_t$  could be  $Q^{\pi}(s_t, a_t)$  or  $A^{\pi}(s_t, a_t)$ .

### B.2 VTRACE

V-Trace is an off-policy technique proposed by IMPALA (Espeholt et al., 2018) to correct for the discrepancy between target policy and behavior policy. V-Trace estimates  $V(s_t)$  by

$$V^{\tilde{\pi}}(s_t) = \mathbb{E}_{\mu} [V(s_t) + \sum_{k \geq 0} \gamma^k c_{[t:t+k-1]} \rho_{t+k} \delta_{t+k}^V V],$$

wherein  $\delta_t^V V \stackrel{def}{=} r_t + \gamma V(s_{t+1}) - V(s_t)$  and  $\rho_t = \min\{\frac{\pi_t}{\mu_t}, \bar{\rho}\}$ .

### B.3 RETRACE

ReTrace (Munos et al., 2016) estimates  $Q(s_t, a_t)$  as follows:

$$Q^{\tilde{\pi}}(s_t, a_t) = \mathbb{E}_{\mu} [Q(s_t, a_t) + \sum_{k \geq 0} \gamma^k c_{[t+1:t+k]} \delta_{t+k}^Q Q],$$

where  $\delta_t^Q Q \stackrel{def}{=} r_t + \gamma Q(s_{t+1}, a_{t+1}) - Q(s_t, a_t)$ .

## C PROOF

*Proof of Proposition 1.* When  $\mathcal{F}_\psi$  is a pre-defined or rule-based mapping,  $\mathcal{F}_\psi$  of actor  $j$  at each training step (wall-clock) is deterministic, namely  $\mathcal{F}_\psi = \mathcal{F}_{\psi_j}$ , so each behavior of actor  $j$  can be uniquely indexed by  $\mathbf{h} \in \mathbf{H} = \{\mathbf{h}_1, \dots, \mathbf{h}_N\}$ , namely,

$$\mathbf{M}_{\mathbf{H}} = \{\mu_{\mathbf{h}} = \mathcal{F}_{\psi_j}(\Phi_{\mathbf{h}}) | \mathbf{h} \in \mathbf{H}\}, \quad (11)$$

where  $\psi_j$  is the same for each behavior of actor  $j$ . Hence, the selection distribution of behavior can be simplified into the selection distribution over  $\mathbf{h} \in \mathbf{H} = \{\mathbf{h}_1, \dots, \mathbf{h}_N\}$  as:

$$\mathcal{P}_{\mu \in \mathbf{M}_{\mathbf{H}}, \Psi} = \mathcal{P}_{\mu_{\mathbf{h}} \in \mathbf{M}_{\mathbf{H}}} = \mathcal{P}_{\mathbf{H}}, \quad (12)$$

where  $\mathbf{M}_{\mathbf{H}} = \{\mathcal{F}_{\psi_j}(\Phi_{\mathbf{h}}) | \mathbf{h} \in \mathbf{H}\}$  and  $\mathcal{P}_{\mathbf{H}}$  is a selection distribution of  $\mathbf{h} \in \mathbf{H} = \{\mathbf{h}_1, \dots, \mathbf{h}_N\}$  and  $N$  is the number of policy models. Substituting equation 12 into equation 5, we can obtain

$$\begin{aligned} \mathcal{L}_{\mathcal{P}} &= \mathbb{E}_{\mu \sim \mathcal{P}_{\mathbf{M}_{\mathbf{H}}, \Psi}} [V_{\mu}^{\text{TV}} + c \cdot V_{\mu}^{\text{TD}}] \\ &= \mathbb{E}_{\mu_{\mathbf{h}} \sim \mathcal{P}_{\mathbf{M}_{\mathbf{H}}}} [V_{\mu_{\mathbf{h}}}^{\text{TV}} + c \cdot V_{\mu_{\mathbf{h}}}^{\text{TD}}] \\ &= \mathbb{E}_{\mathbf{h} \sim \mathcal{P}_{\mathbf{H}}} [V_{\mu_{\mathbf{h}}}^{\text{TV}} + c \cdot V_{\mu_{\mathbf{h}}}^{\text{TD}}] \end{aligned}$$

□

**Corollary 1** (Behavior Circling in Policy Model Selection). *When the policy models overlap as  $\Phi_{\mathbf{h}_1} \approx \dots \approx \Phi_{\mathbf{h}_N}$ , all realizable behavior of actor  $j$  will overlap as  $\mathcal{F}_{\psi_j}(\Phi_{\mathbf{h}_1}) \approx \dots \approx \mathcal{F}_{\psi_j}(\Phi_{\mathbf{h}_N})$ . Behaviors from different model can not be distinguished by  $\mathbf{h}_i$ , the behavior selection via policy model selection becomes invalid.*

*Proof of Proposition 2.* When  $\theta, \mathbf{h}$  are shared among behaviors for each actor at each training step, such as  $\mu = \mathcal{F}_{\psi}(\Phi_{\theta, \mathbf{h}})$  or  $\mu = \mathcal{F}_{\psi}(\Phi_{\theta_1, \mathbf{h}_1}, \dots, \Phi_{\theta_N, \mathbf{h}_N})$ , each behavior for each behavior can be uniquely indexed by  $\psi$ , namely,

$$\mathbf{M}_{\Psi} = \{\mu_{\psi} = \mathcal{F}_{\psi}(\Phi_{1:N}) | \psi \in \Psi\},$$

where  $\Phi_{1:N}$  is the same among behaviors. Hence, the selection distribution of behavior can be simplified into the selection distribution of  $\psi \in \Psi$  as

$$\mathcal{P}_{\mu \in \mathbf{M}_{\mathbf{H}}, \Psi} = \mathcal{P}_{\mu_{\psi} \in \mathbf{M}_{\Psi}} = \mathcal{P}_{\Psi}, \quad (13)$$

where  $\mathcal{P}_{\Psi}$  is a selection distribution of  $\psi \in \Psi$ . Substituting equation 13 into equation 5, we can obtain

$$\begin{aligned} \mathcal{L}_{\mathcal{P}} &= \mathbb{E}_{\mu \sim \mathcal{P}_{\mathbf{M}_{\mathbf{H}}, \Psi}} [V_{\mu}^{\text{TV}} + c \cdot V_{\mu}^{\text{TD}}] \\ &= \mathbb{E}_{\mu_{\psi} \sim \mathcal{P}_{\mathbf{M}_{\Psi}}} [V_{\mu_{\psi}}^{\text{TV}} + c \cdot V_{\mu_{\psi}}^{\text{TD}}] \\ &= \mathbb{E}_{\psi \sim \mathcal{P}_{\Psi}} [V_{\mu_{\psi}}^{\text{TV}} + c \cdot V_{\mu_{\psi}}^{\text{TD}}] \end{aligned}$$

□

The behavior mapping optimization may be a cure for the behavior circling:

**Corollary 2** (Behavior Mapping Optimization Is An Antidote for Behavior Circling). *As for an behavior mapping optimization method, the behavior of actor  $j$  is indexed by  $\mathcal{F}_{\psi}$ . When all the policy models overlap as  $\Phi_{\mathbf{h}_1} \approx \dots \approx \Phi_{\mathbf{h}_N}$ , the realizable behavior of actor  $j$  are*

$$\mathcal{F}_{\psi_1}(\Phi_{\mathbf{h}_1}, \dots, \Phi_{\mathbf{h}_N}), \mathcal{F}_{\psi_2}(\Phi_{\mathbf{h}_1}, \dots, \Phi_{\mathbf{h}_N}), \dots, \mathcal{F}_{\psi_{\infty}}(\Phi_{\mathbf{h}_1}, \dots, \Phi_{\mathbf{h}_N}), \quad (14)$$

wherein  $\psi_i$  is a continuous parameter. Assuming  $\mathcal{F}_{\psi}$  can be uniquely indexed by  $\psi$  and  $\mathcal{F}_{\psi_i} \neq \mathcal{F}_{\psi_j}$ , there are still infinite different behaviors that can be realized by actor  $j$ .



**Proposition 3** (Comparison of Behavior Space). *Given two behavior space  $\mathbf{M}_{\Theta_1, \mathbf{H}_1, \Psi_1}$  and  $\mathbf{M}_{\Theta_2, \mathbf{H}_2, \Psi_2}$ , if  $\mathbf{M}_{\Theta_1, \mathbf{H}_1, \Psi_1}$  is a sub-space of  $\mathbf{M}_{\Theta_2, \mathbf{H}_2, \Psi_2}$ , the space  $\mathbf{M}_{\Theta_2, \mathbf{H}_2, \Psi_2}$  is not less than  $\mathbf{M}_{\Theta_1, \mathbf{H}_1, \Psi_1}$ . Furthermore, if  $\mathbf{M}_{\Theta_1, \mathbf{H}_1, \Psi_1}$  is a sub-space of  $\mathbf{M}_{\Theta_2, \mathbf{H}_2, \Psi_2}$  and  $\mathbf{M}_{\Theta_1, \mathbf{H}_1, \Psi_1}$  is not equal to  $\mathbf{M}_{\Theta_2, \mathbf{H}_2, \Psi_2}$ , the space  $\mathbf{M}_{\Theta_2, \mathbf{H}_2, \Psi_2}$  is larger than  $\mathbf{M}_{\Theta_1, \mathbf{H}_1, \Psi_1}$ .*

*Proof of Proposition 3.* Since  $\mathbf{M}_{\Theta_1, \mathbf{H}_1, \Psi_1}$  and  $\mathbf{M}_{\Theta_2, \mathbf{H}_2, \Psi_2}$  are sets. When  $\mathbf{M}_{\Theta_1, \mathbf{H}_1, \Psi_1} \subseteq \mathbf{M}_{\Theta_2, \mathbf{H}_2, \Psi_2}$ ,  $\mathbf{M}_{\Theta_1, \mathbf{H}_1, \Psi_1}$  is not larger than  $\mathbf{M}_{\Theta_2, \mathbf{H}_2, \Psi_2}$ . When  $\mathbf{M}_{\Theta_1, \mathbf{H}_1, \Psi_1} \subset \mathbf{M}_{\Theta_2, \mathbf{H}_2, \Psi_2}$ ,  $\mathbf{M}_{\Theta_1, \mathbf{H}_1, \Psi_1}$  is smaller than  $\mathbf{M}_{\Theta_2, \mathbf{H}_2, \Psi_2}$ .  $\square$

According to the behavior space construction formulation, we can draw the following Corollary:

**Corollary 3.** *Given the same policy model structure  $\Phi$  and the same form of behavior mapping  $\mathcal{F}$ . Under Assumption 1, the behavior space can be fully determined by  $\mathbf{H}$  and  $\Psi$ . For any two behavior space  $\mathbf{M}_{\mathbf{H}_1, \Psi_1}$  and  $\mathbf{M}_{\mathbf{H}_2, \Psi_2}$ , if  $\mathbf{H}_1 \subseteq \mathbf{H}_2$  and  $\Psi_1 \subseteq \Psi_2$ , the behavior space  $\mathbf{M}_{\mathbf{H}_1, \Psi_1}$  is a sub-space of  $\mathbf{M}_{\mathbf{H}_2, \Psi_2}$ . Based on that, the space  $\mathbf{M}_{\mathbf{H}_2, \Psi_2}$  is not smaller than  $\mathbf{M}_{\mathbf{H}_1, \Psi_1}$ .*

**Corollary 4.** *Given the same policy model structure  $\Phi$  and the same form of behavior mapping  $\mathcal{F}$ . Under Assumption 1, the behavior space can be fully determined by  $\mathbf{H}$  and  $\Psi$ . For any two behavior space  $\mathbf{M}_{\mathbf{H}_1, \Psi_1}$  and  $\mathbf{M}_{\mathbf{H}_2, \Psi_2}$ , if at least one of the following conditions holds:*

- $\mathbf{H}_1 \subset \mathbf{H}_2$  and  $\Psi_1 \subseteq \Psi_2$ ,
- $\mathbf{H}_1 \subseteq \mathbf{H}_2$  and  $\Psi_1 \subset \Psi_2$ ,
- $\mathbf{H}_1 \subset \mathbf{H}_2$  and  $\Psi_1 \subset \Psi_2$ ,

*the behavior space  $\mathbf{M}_{\mathbf{H}_1, \Psi_1}$  is a sub-space of  $\mathbf{M}_{\mathbf{H}_2, \Psi_2}$  and  $\mathbf{M}_{\mathbf{H}_1, \Psi_1}$  is not equal to  $\mathbf{M}_{\mathbf{H}_2, \Psi_2}$ . Based on that, the space  $\mathbf{M}_{\mathbf{H}_2, \Psi_2}$  is larger than  $\mathbf{M}_{\mathbf{H}_1, \Psi_1}$ .*

## D BEHAVIOR SPACE CONSTRUCTION FOR MORE TASKS AND ALGORITHMS VIA LBC

Following the pipeline given in equation 9, different implementations of LBC can be acquired by simply selecting different entropy control function  $f_{\tau_i}(\cdot)$  and behavior distillation function  $g(\cdot, \omega)$  according to the corresponding RL algorithms and tasks.

### D.1 SELECTION FOR ENTROPY CONTROL FUNCTION

Here we would give some examples for the selection of entropy control function  $f_{\tau_i}(\cdot)$ .

**Continuous Control Tasks** For tasks with continuous action spaces, the entropy control function can be selected as gaussian distribution, *i.e.*,

$$\mathbf{M}_{\mathbf{H}, \Psi} = \{g(\text{Normal}(\Phi_{\mathbf{h}_1}, \sigma_1), \dots, \text{Normal}(\Phi_{\mathbf{h}_N}, \sigma_N), \omega_1, \dots, \omega_N) | \psi \in \Psi\} \quad (15)$$

or uniform distribution, *i.e.*,

$$\mathbf{M}_{\mathbf{H}, \Psi} = \{g(\text{Uniform}(\Phi_{\mathbf{h}_1} - b_1/2, \Phi_{\mathbf{h}_1} + b_1/2), \dots, \text{Uniform}(\Phi_{\mathbf{h}_N} - b_N/2, \Phi_{\mathbf{h}_N} + b_N/2), \omega_1, \dots, \omega_N) | \psi \in \Psi\} \quad (16)$$

where  $\psi = (\sigma_1, \dots, \sigma_N, \omega_1, \dots, \omega_N)$  for gaussian distribution and  $\psi = (b_1, \dots, b_N, \omega_1, \dots, \omega_N)$  for uniform distribution.

### Discrete Control Tasks and Value-Based Algorithms

$$\mathbf{M}_{\mathbf{H}, \Psi} = \{g(\epsilon_1\text{-greedy}(\Phi_{\mathbf{h}_1}), \dots, \epsilon_N\text{-greedy}(\Phi_{\mathbf{h}_N}), \omega_1, \dots, \omega_N) | \psi \in \Psi\} \quad (17)$$

where  $\psi = (\epsilon_1, \dots, \epsilon_N, \omega_1, \dots, \omega_N)$ .

### Discrete Control Tasks and Policy-Based Algorithms

$$\mathbf{M}_{\mathbf{H}, \Psi} = \{g(\text{softmax}_{\tau_1}(\Phi_{\mathbf{h}_1}), \dots, \text{softmax}_{\tau_N}(\Phi_{\mathbf{h}_N}), \omega_1, \dots, \omega_N) | \psi \in \Psi\} \quad (18)$$

where  $\psi = (\tau_1, \dots, \tau_N, \omega_1, \dots, \omega_N)$ .

### D.2 SELECTION FOR BEHAVIOR DISTILLATION FUNCTION

#### Mixture Model

$$\mathbf{M}_{\mathbf{H}, \Psi} = \left\{ \sum_{i=1}^N \omega_i f_{\tau_i}(\Phi_{\mathbf{h}_i}) | \psi \in \Psi \right\} \quad (19)$$

**Knowledge Distillation** The knowledge distillation method can be seen as a derivative form of mixture model. The mixture model is simple and straightforward, but it requires more resources for model storage and inference. To address this disadvantage, we can distill the knowledge of multiple policies into a single network using knowledge distillation.

$$\mathbf{M}_{\mathbf{H}, \Psi} = \{\text{Distill}(f_{\tau_1}(\Phi_{\mathbf{h}_1}), \dots, f_{\tau_N}(\Phi_{\mathbf{h}_N}), \omega_1, \dots, \omega_N) | \psi \in \Psi\} \quad (20)$$

and the knowledge distillation process  $\text{Distill}(\cdot)$  can be realized by supervised learning.

**Parameters Fusion** Define  $\mu_f$  as the generated behavior policy which shares the same network structure with the policy models in the population, and is parameterized by  $\theta_f$ . Define  $\theta_{\mathbf{h}_i}$  as the parameters of policy  $f_{\tau_i}(\Phi_{\mathbf{h}_i}, \tau_i)$ . Then we can define the parameters fusion function  $\text{Fusion}(\cdot, \omega)$ ,

$$\mathbf{M}_{\mathbf{H}, \Psi} = \{\mu_f = \text{Fusion}(f_{\tau_1}(\Phi_{\mathbf{h}_1}), \dots, f_{\tau_N}(\Phi_{\mathbf{h}_N}), \omega_1, \dots, \omega_N) | \psi \in \Psi\} \quad (21)$$

where  $\theta_f = \sum_{i=1}^N \omega_i \theta_{\mathbf{h}_i, \tau_i}$ .

## E ADAPTIVE CONTROL MECHANISM

In this paper, we cast the behavior control into the behavior mapping optimization, which can be further simplified into the selection of  $\psi \in \Psi$ . We formalize this problem via multi-armed bandits (MAB). In this section, we describe the multi-arm bandit design of our method. For a more thorough explanation and analysis, we refer the readers to (Garivier & Moulines, 2011).

### E.1 DISCRETIZATION

Since  $\Psi$  is a continuous space, the optimization of  $\psi \in \Psi$  is a continuous optimization problem. However, MAB usually only handle discrete control tasks. Hence, we have to discretize  $\Psi$  into  $K$  regions according to the discretization accuracy  $\tau$ , wherein each arm of MAB corresponds to a region of the continuous space.

**Remark.** *The discretization accuracy  $\tau$  is related to the accuracy of the algorithm. In general, a higher discretization accuracy indicates a higher accuracy of the algorithm, but correspondingly, a higher computational complexity of the algorithm.*

**Example 1** (Example of Discretization). *As for a  $\epsilon$ -greedy behavior mapping,  $\psi = \epsilon$  and  $\Psi = \{\psi = \epsilon | \epsilon \in [0, 1]\}$ . We can set the discretization accuracy  $\tau = 0.1$ , and we can discretize  $\Psi$  into 10 regions corresponding to  $K = 10$  arms. Each arm corresponds to an interval. For example,  $k = 1$  corresponds to  $[0, 0.1]$ ;  $k = 2$  corresponds to  $[0.1, 0.2]$ ...  $k = 10$  corresponds to  $[0.9, 1.0]$ .*

### E.2 SAMPLE AND UPDATE

We adopt the Thompson Sampling (Garivier & Moulines, 2011).  $\mathcal{K} = \{1, \dots, K\}$  denote a set of arms available to the decision maker, who is interested in maximizing the expected cumulative return (Badia et al., 2020a; Parker-Holder et al., 2020). The optimal strategy for each actor is to pull the arm with the largest mean reward. At the beginning of each round, each actor will produce a sample mean from its mean reward model for each arm, and pulls the arm from which it obtained the largest sample. After observing the selected arm's reward, it updates its mean reward model.

In general, at each time  $t$ , MAB method will choose an arm  $k_t$  from all possible arms  $\mathcal{K} = \{1, \dots, K\}$  according to a sampling distribution  $\mathcal{P}_{\mathcal{K}}$ , which is normally conditioned on the sequence of previous decisions and returns. Then we will uniformly sample the parameters  $\psi$  from this discretized regions. Based on the  $\psi$ , we can obtain the corresponding behavior according to  $\mathcal{F}(\Phi_h)$  or  $\mathcal{F}(\Phi_{h_1}, \dots, \Phi_{h_N})$ . Executing the behaviors in the environment, each actor will receive a excitation/reward signal  $R_t(k_t) \in \mathbb{R}$ , which will be used to update the MAB.

### E.3 UPPER CONFIDENCE BOUND

The UCB (Garivier & Moulines, 2011) are often used to encourage MAB to try more the arms with a low frequency of use. Let's first define the number of times that the arm  $k$  has been selected within  $T$  rounds as follows:

$$N_T(x) = \sum_{t=0}^T \mathbb{1}_{k_t=x}. \quad (22)$$

Then we can obtain the empirical mean of the arm  $x$  within  $T$  rounds as follows:

$$V_T(x) = N_T(x) \sum_{t=0}^T R_t(x) \mathbb{1}_{k_t=x}. \quad (23)$$

The UCB methods encourage the decision maker (actor-wise) to maximize the UCB scores:

$$\text{Score}_x = V_T(x) + c \cdot \sqrt{\frac{\log(1 + \sum_j N_T(j))}{1 + N_T(x)}} \quad (24)$$

The optimal strategy for each actor is to pull the arm with the largest mean scores. At the beginning of each round, each actor will produce a sample mean from its mean reward model for each arm, and pulls the arm from which it obtained the largest sample. After observing the selected arm's scores, it updates its mean reward model.

**Remark.** In practical, Z-score Normalization are normally used to normalized  $V_T(x)$ , namely  $\frac{V_T(x) - \mathbb{E}[V_T]}{\mathcal{D}[V_T]}$ , which can be formulated as

$$\text{Score}_x = \frac{V_T(x) - \mathbb{E}[V_T(x)]}{\mathcal{D}[V_T(x)]} + c \cdot \sqrt{\frac{\log(1 + \sum_j N_T(j))}{1 + N_T(x)}} \quad (25)$$

#### E.4 POPULATION-BASED MAB

In the non-stationary scenario, the distributions of  $V_T(x)$  could be shifted in the course of the lifelong learning. The standard UCB-based MAB failed to adapt to the change of the reward distribution and thus we refer to a population-based MAB to handle this problem, which jointly train a population of MAB with different hyperparameters. The sampling and update procedure of MAB is slightly different from the origin MAB, which will be discussed in the following. The main implementation of our population-based MAB has been concluded in Algorithm 1.

##### E.4.1 MAB POPULATION FORMULATION

Assuming there are  $N$  bandits  $B_{\mathbf{h}_i}$  to from a population  $\mathcal{B} = \{B_{\mathbf{h}_1}, \dots, B_{\mathbf{h}_N}\}$ , wherein each bandit can be uniquely indexed by its hyper-parameter  $\mathbf{h}_i$  and keep other hyper-parameters remain the same such as the discretization. In this paper,  $\mathbf{h}_i = c$ , wherein  $c$  is the trade-off coefficient in equation 24, which is uniformly sampled from  $[0.5, 1.5]$ , i.e., randomly select a  $c \in [0.5, 1.5]$  while initializing each bandit.

##### E.4.2 POPULATION-BASED SAMPLE

During the sampling procedure, each bandit  $B_{\mathbf{h}_i}$  will sample  $D$  arm  $k_i \in \mathcal{K}$  with the Top-D ucb-scores. After all the bandits sample  $D$  arm, there are  $D \times N$  sampled arms. We summarize the number of times each arm is selected, and sorted in descending order by the number of times they are selected. Then, we can obtain an arm  $x_{j,t}$  that is selected the most times, which is the sample output of the population-based MAB. Finally, we uniformly sample a  $\psi_{j,t}$  from the region indexed by  $x_{j,t}$ .

**Example 2.** Assuming there are 7 bandits, and each bandit will sample  $D = 2$  arms from  $\mathcal{K} = \{1, \dots, 10\}$ . Assuming that the sample output is as follows:

$$1, 2, 1, 3, 2, 4, 5; 1, 1, 2, 2, 1, 1, 4.$$

Then, the arm  $k = 1$  is the arm being selected the most times, so we can get the sampled arm  $x_{j,t} = 1$ .

**Remark.** Noting that, if there are more than one arm that is selected the most times, we can uniformly sample one from these arms.

##### E.4.3 POPULATION-BASED UPDATE

With  $x_{j,t}$ , according to the behavior space equation 10, we can obtain a behavior  $\mu_{j,t} = \mathcal{F}_\psi(\Phi_{\mathbf{h}})$  or  $\mu_{j,t} = \mathcal{F}_{\psi_{j,t}}(\Phi_{\mathbf{h}_1}, \dots, \Phi_{\mathbf{h}_N})$ . Execute  $\mu_{j,t}$  in the environment and we can obtain the return  $G_{j,t}$ . With  $G_{j,t}$ , we can update each bandit in the population based on equation 22 - equation 24.

##### E.4.4 BANDIT REPLACEMENT

Similar to the sliding windows (Badia et al., 2020a), to tackle the non-stationary problem, we have to track the changes of optimization objectives in a timely manner. To achieve this, we update the replace in the population regularly so that it captures short-term information to improve its tracking performance.

Noting that there are many methods to solve this non-stationary problem at present, such as the sliding windows (Badia et al., 2020a). Since this is not the main proposition of this paper, we just choose a feasible implementation to handle this problem.

**Algorithm 1** Population-Based Multi-Arm Bandits (Actor-Wise)

---

```

// For Each Actor j
// Initialize Bandits Population
Initialize each bandit  $B_{\mathbf{h}_i}$  in the population with different hyper-parameters  $c$ .
Incorporate each bandit together to form a population of bandits  $\mathcal{B}$ .
for each episode  $t$  do
  for each  $B_{\mathbf{h}_i}$  in  $\mathcal{B}$  do
    Sample  $D$  arms with Top-D UCB Score via equation 24.
  end for
  Summarize  $N \times D$  arms and count the selected times of each arm.
  Uniformly sample an arm among arms that selected the most times to obtain arm  $x_{j,t}$ .
  Uniformly sample a  $\psi_{j,t}$  from the region indexed by  $x_{j,t}$ .
  Obtain a behavior  $\mu_{j,t} = \mathcal{F}_\psi(\Phi_{\mathbf{h}})$  or  $\mu_{j,t} = \mathcal{F}_{\psi_{j,t}}(\Phi_{\mathbf{h}_1}, \dots, \Phi_{\mathbf{h}_N})$ .
  Execute  $\mu_{j,t}$  and obtain the return  $G_{j,t}$ .
  for each  $B_{\mathbf{h}_i}$  in  $\mathcal{B}$  do
    Update  $B_{\mathbf{h}_i}$  via equation 22 and equation 23.
  end for
  Update each bandit in the population via equation 22 and equation 23.
  // Replace Bandit from The Population
  if  $t \bmod T_{replace} = 0$  then
    Remove one bandit from the bandit population Uniformly and recreate (reinitialize)
    one into it.
  end if
end for

```

---

## F ALGORITHM PSEUDOCODE

We concluded our algorithm in in the Algorithm. 2. Apart from that, we also concluded our model architecture in App. L.

---

### Algorithm 2 Learnable Behavior Control

---

```

Initialize the Data Buffer (DB), the Parameter Server (PS), the Learner Push Parameter Interval  $d_{push}$  and the Actor Pull Parameter Interval  $d_{pull}$ .
// LEARNER i
Initialize the network parameter  $\theta_i$  (for model structure, see App. L)
for Training Step t do
  Load data from DB.
  Estimate  $Q_{\theta_i}$  by  $Q^{\tilde{\pi}}(s_t, a_t) = \mathbb{E}_{\mu}[Q(s_t, a_t) + \sum_{k \geq 0} \gamma^k c_{[t+1:t+k]} \delta_{t+k}^Q Q]$ , wherein the target policy  $\pi_i = \text{softmax}(A_i)$ .
  Estimate  $V_{\theta_i}$  by  $V^{\tilde{\pi}}(s_t) = \mathbb{E}_{\mu}[V(s_t) + \sum_{k \geq 0} \gamma^k c_{[t:t+k-1]} \rho_{t+k} \delta_{t+k}^V V]$ , wherein the target policy  $\pi_i = \text{softmax}(A_i)$ .
  Update  $\theta_i$  via  $\nabla_{\theta} \mathcal{J}(\theta) = \mathbb{E}_{\pi} [\sum_{t=0}^{\infty} \Phi_t \nabla_{\theta} \log \pi_{\theta_i}(a_t | s_t)]$ .
  if t mod  $d_{push} = 0$  then
    Push  $\theta_i$  into PS.
  end if
end for
// ACTOR j
for kth episode at training step t do
  Sample a  $\psi_{j,k} = (\tau_1, \dots, \tau_N, \omega_1, \dots, \omega_N)$  via the MAB-based meta-controller (see App. E).
  Generalized Policy Selection. Adjusting the contribution proportion of the each learned policies for the behavior via a importance weight  $\omega = (\omega_1, \dots, \omega_N)$ .
  Policy-Wise Entropy Control. Adjusting the entropy of each policy via a  $\tau_i$ , (e.g.,  $\pi_{\tau_i} = \text{softmax}_{\tau_i}(\Phi_i)$ ).
  Behavior Distillation from Multiple Policies. Distilling the entropy-controlled policies into a behavior policy  $\mu_{j,k}$  via a mixture model  $\mu = \sum_i^N \pi_{\tau_i}$ .
  Obtaining episode  $\tau_{j,k}$  and reward  $G_{j,k}$  via executing  $\mu_{j,k}$ , and push  $\tau_{j,k}$  into DB.
  Update the meta-controller with  $(\mu_{j,k}, G_{j,k})$ .
  if t mod  $d_{pull} = 0$  then
    Pull  $\{\theta_1, \dots, \theta_N\}$  from PS.
  end if
end for

```

---

## G AN LBC-BASED VERSION OF RL

The behavior space is vital for RL methods, which can be used to categorize RL algorithms. Given the model structure  $\Phi$ , and the form of  $\mathcal{F}$ , the behavior space can be fully determined by  $\mathbf{H}$  and  $\Psi$ , which can be used to categorize RL methods. We say one algorithm belongs to  $\text{LBC-H}_N^C\text{-}\Psi_L^K$  when **1)** the hyper-parameters  $\mathbf{h}$  is a C-D vector and  $\mathbf{h}$  has N possible values corresponding to N different policy models, and **2)**  $\psi$  is a K-D vector and  $\psi$  has L possible values corresponding to L realizable behavior mappings at each training step. Based on that, we can offer a general view to understand prior methods from the perspective of behavior control, which is illustrated in Tab. 2.

Table 2: An LBC-based Version of RL Methods.

Algorithm	PBT	Agent57	DvD	LBC-BM (Ours)
$\mathcal{F}$	$\epsilon - greedy$	$\epsilon - greedy$	identical mapping	$\sum_{i=1}^N \omega_i \text{Softmax}_{\tau_i}(\Phi_{\mathbf{h}_i})$
$\mathbf{M}_{\mathbf{H}, \Psi}$	$\{\mathcal{F}_{\psi}(\Phi_{\mathbf{h}_j})   \mathbf{h}_j \in \mathbf{H}, \psi \in \Psi\}$	$\{\mathcal{F}_{\psi}(\Phi_{\mathbf{h}_j})   \mathbf{h}_j \in \mathbf{H}, \psi \in \Psi\}$	$\{\mathcal{F}_{\psi}(\Phi_{\mathbf{h}_j})   \mathbf{h}_j \in \mathbf{H}, \psi \in \Psi\}$	$\{\mathcal{F}_{\psi}(\Phi_{\mathbf{h}_1}, \dots, \Phi_{\mathbf{h}_N})   \mathbf{h}_j \in \mathbf{H}, \psi \in \Psi\}$
$\mathbf{H}$	$\{\mathbf{h}_i = (h_1, \dots, h_C)   i = 1, \dots, N\}$	$\{\gamma_i, \beta_i   i = 1, \dots, N\}$	$\{(\lambda_i)   i = 1, \dots, N\}$	$\{\mathbf{h}_i = (h_1, \dots, h_C)   i = 1, \dots, N\}$
$\Psi$	$\{(\epsilon_l)   l = 1, \dots, L\}$	$\{(\epsilon_l)   l = 1, \dots, L\}$	$\{1\}$	$\{(\omega_1, \tau_1, \dots, \omega_N, \tau_N)\}$
Category	$\text{LBC-H}_N^C\text{-}\Psi_L^1$	$\text{LBC-H}_N^2\text{-}\Psi_L^1$	$\text{LBC-H}_N^1\text{-}\Psi_1^1$	$\text{LBC-H}_N^C\text{-}\Psi_{\infty}^{K=2N}$
$ \mathbf{M}_{\mathbf{H}, \Psi} $	$N \times L$	$N \times L$	N	$\infty$
Meta-Controller ( $\mathbf{H}$ )	ES	MAB	MAB	Ensemble equation 10
Meta-Controller ( $\Psi$ )	Rule-Based	Rule-Based	Rule-Based	MAB

## H EXPERIMENT DETAILS

### H.1 IMPLEMENTATION DETAILS

The overall training architecture is on the top of the Learner-Actor framework (Espeholt et al., 2018), which supports large-scale training. Additionally, the recurrent encoder with LSTM (Schmidhuber, 1997) is used to handle the partially observable MDP problem (Bellemare et al., 2013). The burn-in technique is adopted to deal with the representational drift (Kapturowski et al., 2019), and we train each sample twice. A complete description of the hyper-parameters can see App. I.

### H.2 EXPERIMENTAL SETUP

We employ additional environments to evaluate the scores during model training, and we record the undiscounted episode returns averaged over 5 seeds using a windowed mean over 32 episodes. We evaluated all agents on 57 Atari 2600 games from the arcade learning environment (Bellemare et al., 2013, ALE) by recording the average score of the population of agents during model training. Noting that episodes will be truncated at 100K frames as other baseline algorithms (Hessel et al., 2018; Badia et al., 2020a; Schmitt et al., 2020; Badia et al., 2020b; Kapturowski et al., 2019).

### H.3 RESOURCES USED

All the experiment is accomplished using 10 workers with 72 cores CPU and 3 learners with 3 Tesla-V100-SXM2-32GB GPU.



## I HYPER-PARAMETERS

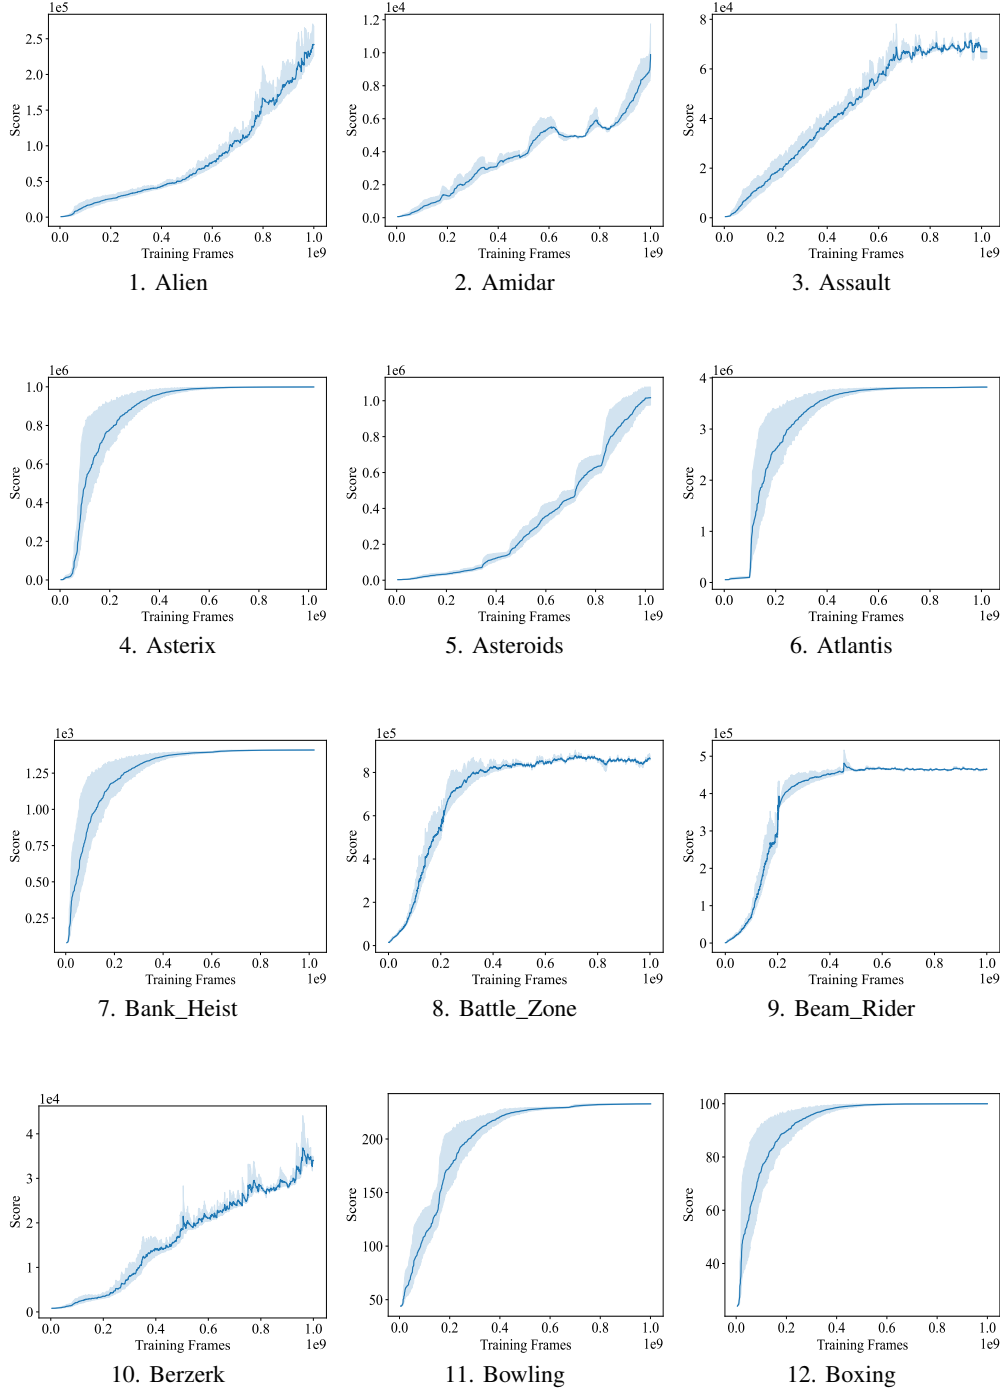
The hyper-parameters that we used in all experiments are like those of NGU Badia et al. (2020b) and Agent57 (Badia et al., 2020a). However, for completeness and readability, we detail them below in Tab. 3. We also include the hyper-parameters we used in the population-based MAB. For more details on the parameters in ALE, we refer the readers to see (Machado et al., 2018).

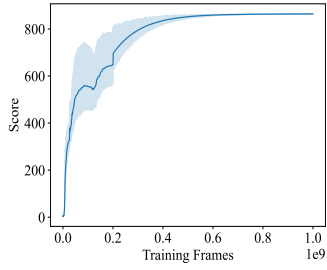
Table 3: Hyper-Parameters for Atari Experiments.

Parameter	Value	Parameter	Value
Burn-in	40	Replay	2
Seq-length	80	Burn-in Stored Recurrent State	Yes
Bootstrap	Yes	Batch size	64
$V$ -loss Scaling ( $\xi$ )	1.0	$Q$ -loss Scaling ( $\alpha$ )	5.0
$\pi$ -loss Scaling ( $\beta$ )	5.0	Importance sampling clip $\bar{c}$	1.05
Importance Sampling Clip $\bar{\rho}$	1.05	LSTM Units	256
Weight Decay Rate	0.01	Optimizer	Adam weight decay
Learning Rate	5.3e-4	Weight Decay Schedule	Anneal linearly to 0
Warmup Steps	4000	Learning Rate Schedule	Anneal linearly to 0
AdamW $\beta_1$	0.9	Auxiliary Forward Dynamic Task	Yes
AdamW $\epsilon$	1e-6	Learner Push Model Every $d_{push}$ Steps	25
AdamW $\beta_2$	0.98	Auxiliary Inverse Dynamic Task	Yes
AdamW Clip Norm	50.0	Actor Pull Model Every $d_{pull}$ Steps	64
$\gamma_1$	0.997	$\mathcal{RS}_1$	$\text{sign}(x) \cdot (\sqrt{ x  + 1} - 1) + 0.001 \cdot x$
$\gamma_2$	0.999	$\mathcal{RS}_2$ (log scaling)	$\log( x  + 1) \cdot (2 \cdot \mathbb{1}_{r \geq 0} - \mathbb{1}_{r \leq 0})$
$\gamma_3$	0.99	$\mathcal{RS}_3$	$0.3 \cdot \min(\tanh x, 0) + 5 \cdot \max(\tanh x, 0)$
Population Num.	7	UCB $c$	Uniformly sampled from $[0.5, 1.5]$
D of Top-D	4	Replacement Interval $T_{replace}$	50
Range of $\tau_i$	$[0, \exp 4]$	Range of $\omega_i$	$[0, 1]$
Discrete Accuracy of $\tau_i$	0.2	Discrete Accuracy of $\omega_i$	0.1
Max episode length	30 <i>min</i>	Image Size	(84, 84)
Grayscaled/RGB	Grayscaled	Life information	Not allowed
Action Space	Full	Sticky action probability	0.0
Num. Action Repeats	4	Random noops range	30
Num. Frame Stacks	4	Num. Atari Games	57 (Full)

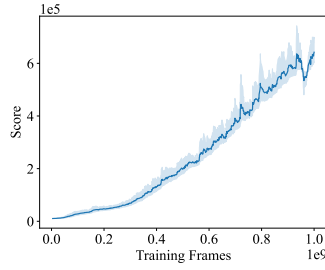
## J EXPERIMENTAL RESULTS

### J.1 ATARI GAMES LEARNING CURVES

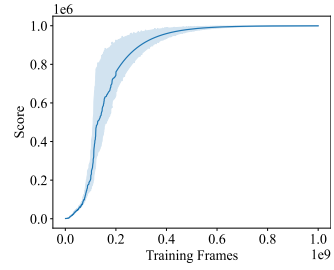




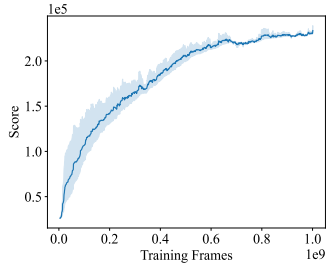
13. Breakout



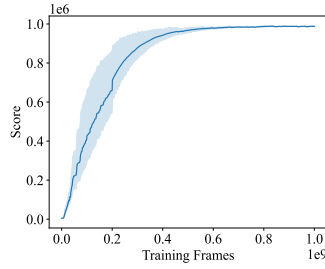
14. Centipede



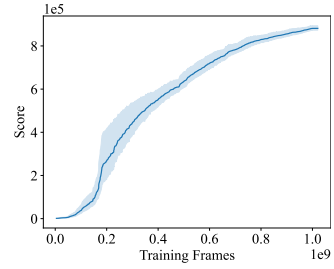
15. Chopper\_Command



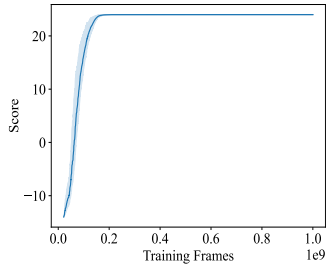
16. Crazy\_Climber



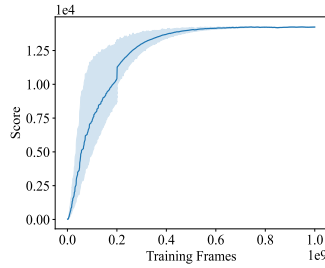
17. Defender



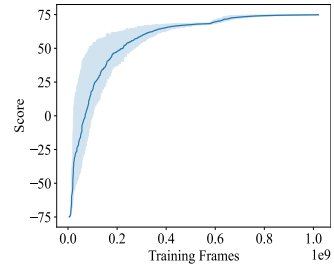
18. Demon\_Attack



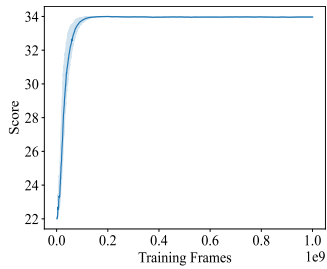
19. Double\_Dunk



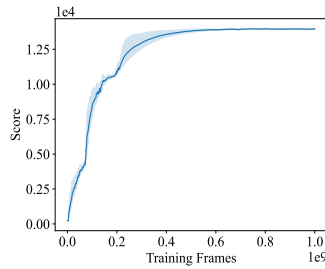
20. Enduro



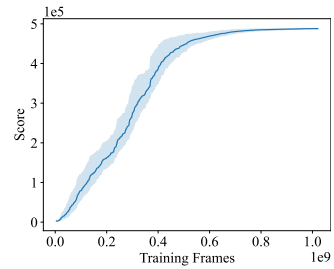
21. Fishing\_Derby



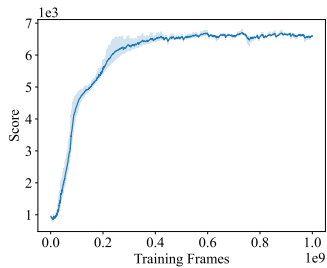
22. Freeway



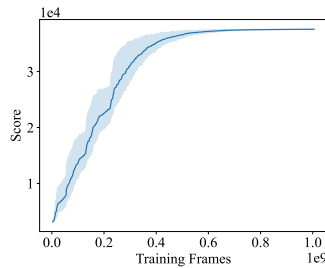
23. Frostbite



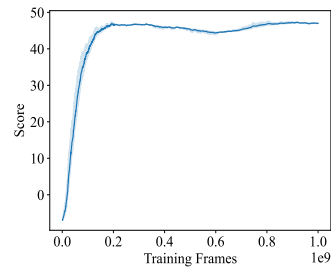
24. Gopher



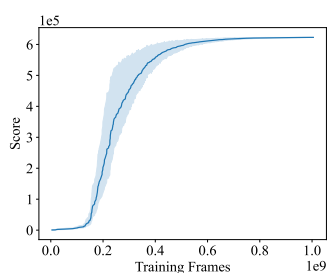
25. Gravitar



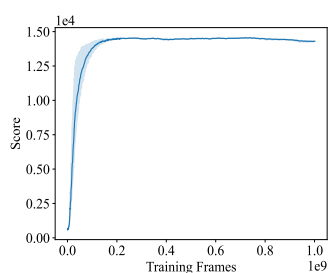
26. Hero



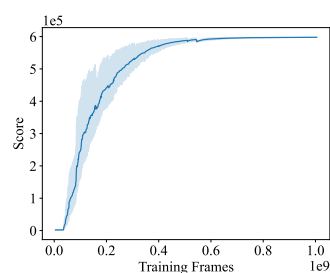
27. Ice\_Hockey



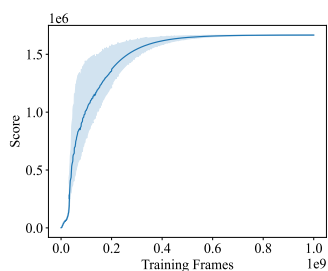
28. Jamesbond



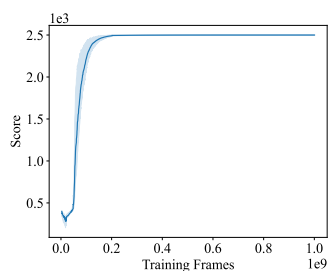
29. Kangaroo



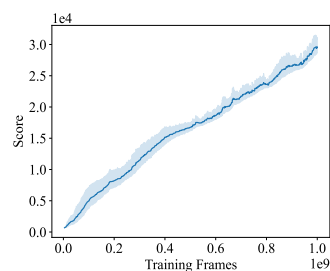
30. Krull



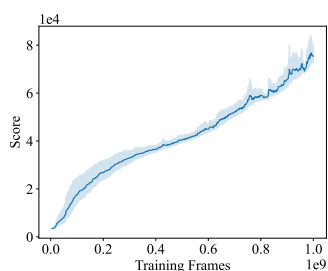
31. Kung\_Fu\_Master



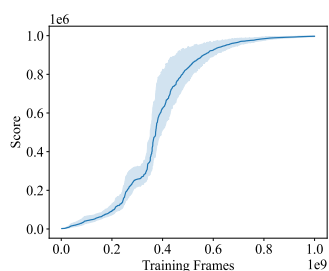
32. Montezuma\_Revenge



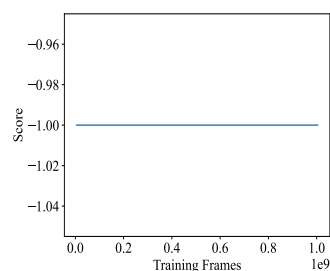
33. Ms\_Pacman



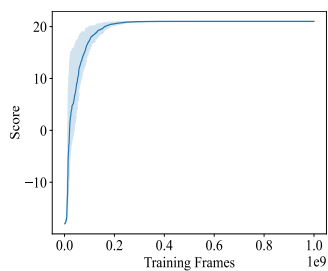
34. Name\_This\_Game



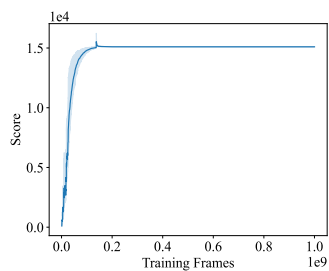
35. Phoenix



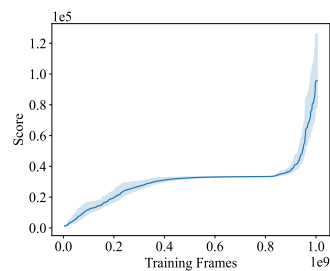
36. Pitfall



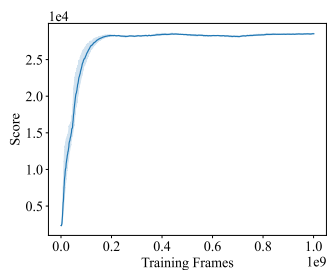
37. Pong



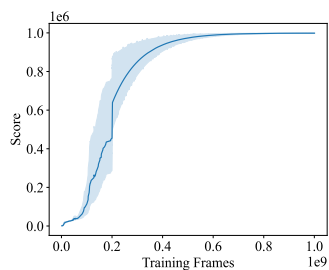
38. Private\_Eye



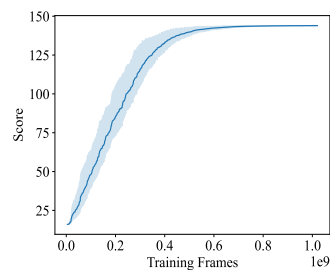
39. Qbert



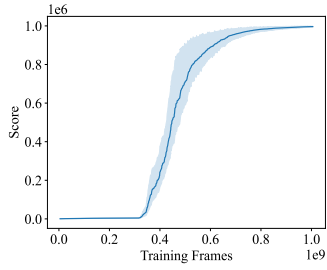
40. Riverraid



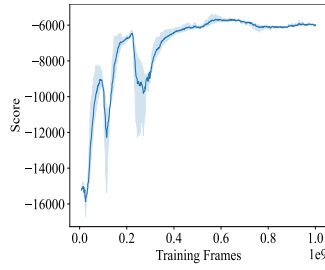
41. Road Runner



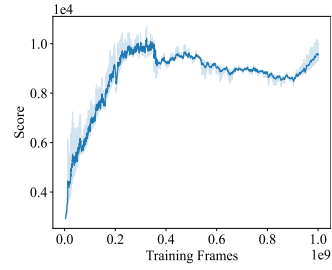
42. Robotank



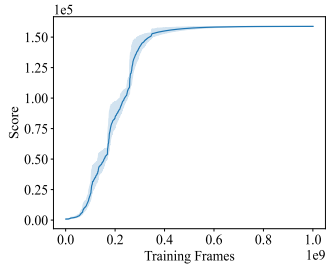
43. Seaquest



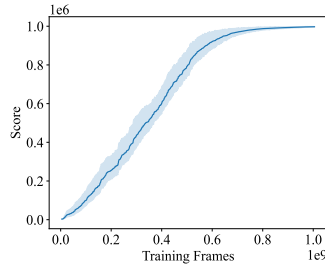
44. Skiing



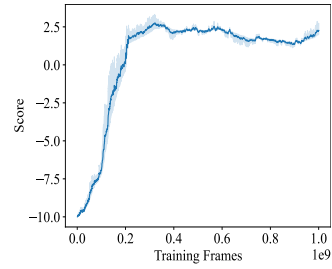
45. Solaris



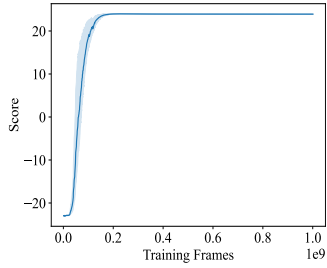
46. Space\_Invaders



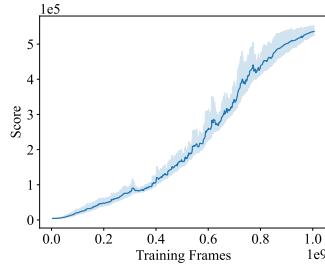
47. Star\_Gunner



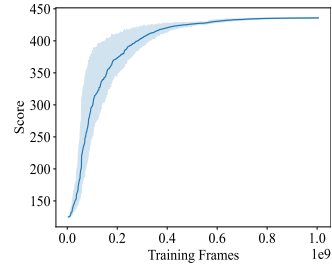
48. Surround



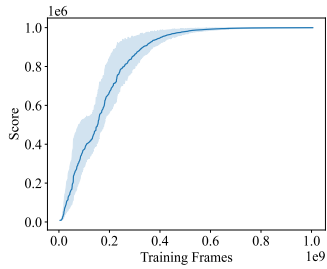
49. Tennis



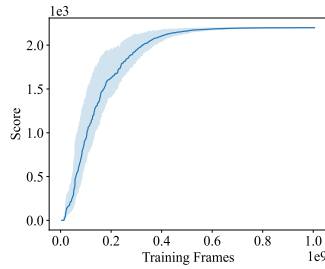
50. Time\_Pilot



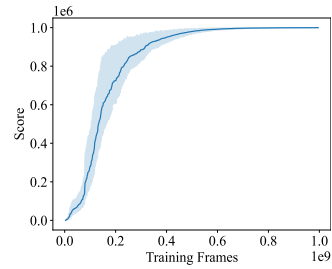
51. Tutankham



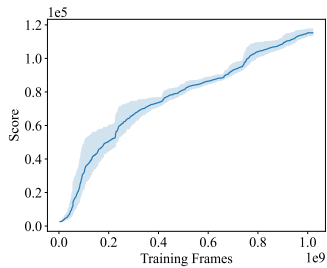
52. Up\_N\_Down



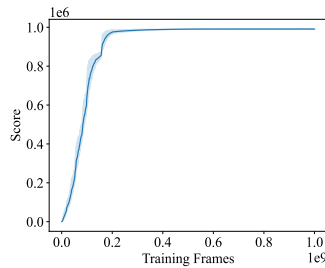
53. Venture



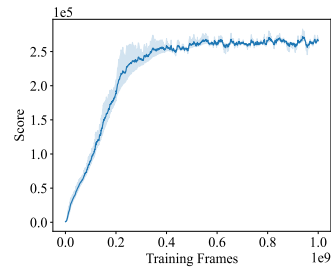
54. Video\_Pinball



55. Wizard\_of\_Wor



56. Yars\_Revenge



57. Zaxxon

## J.2 ATARI GAMES TABLE OF SCORES BASED ON HUMAN AVERAGE SCORES

In this part, we detail the raw score of several representative SOTA algorithms , including the SOTA model-free algorithms, SOTA model-based algorithms, and other SOTA algorithms. Apart from the raw score, we also report the Human Average Scores for each game and Human Normalized Score (HNS) for each algorithm in Atari 57 games. More details on these algorithms can see Machado et al. (2018); Toromanoff et al. (2019).

Table 4: Score table of SOTA model-free algorithms on HNS(%).

Games Scale	RND	Average Human	AGENT57 100B	HNS(%)	Ours 1B	HNS(%)	MEME 1B	HNS(%)
Alien	227.8	7127.8	<b>297638.17</b>	<b>4310.30%</b>	279703.5	4050.37%	83683.43	1209.50%
Amidar	5.8	1719.5	<b>29660.08</b>	<b>1730.42%</b>	12996.3	758.04%	14368.9	838.13%
Assault	222.4	742	<b>67212.67</b>	<b>12892.66%</b>	62025.7	11894.40%	46635.86	8932.54%
Asterix	210	8503.3	991384.42	11951.51%	<b>999999</b>	<b>12055.38%</b>	769803.92	9279.71%
Asteroids	719	47388.7	150854.61	321.70%	<b>1106603.5</b>	<b>2369.60%</b>	364492.07	779.46%
Atlantis	12850	29028.1	1528841.76	9370.64%	<b>3824506.3</b>	<b>23560.59%</b>	1669226.33	10238.39%
Bank Heist	14.2	753.1	23071.5	3120.49%	1410	188.90%	<b>87792.55</b>	<b>11879.60%</b>
Battle Zone	236	37187.5	<b>934134.88</b>	<b>2527.36%</b>	857369	2319.62%	776770	2101.50%
Beam Rider	363.9	16926.5	300509.8	1812.19%	<b>457321</b>	<b>2758.97%</b>	51870.2	310.98%
Berzerk	123.7	2630.4	<b>61507.83</b>	<b>2448.80%</b>	35340	1404.89%	38838.35	1544.45%
Bowling	23.1	160.7	251.18	165.76%	233.1	152.62%	<b>261.74</b>	<b>173.43%</b>
Boxing	0.1	12.1	<b>100</b>	<b>832.50%</b>	<b>100</b>	<b>832.50%</b>	99.85	831.25%
Breakout	1.7	30.5	790.4	2738.54%	<b>864</b>	<b>2994.10%</b>	831.08	2879.79%
Centipede	2090.9	12017	412847.86	4138.15%	<b>728080</b>	<b>7313.94%</b>	245892.18	2456.16%
Chopper Command	811	7387.8	999900	15191.11%	<b>999999</b>	<b>15192.62%</b>	912225	13858.02%
Crazy Climber	10780.5	36829.4	<b>565909.85</b>	<b>2131.10%</b>	233090	853.43%	339274.67	1261.07%
Defender	2874.5	18688.9	677642.78	4266.80%	<b>995950</b>	<b>6279.56%</b>	543979.5	3421.60%
Demon Attack	152.1	1971	143161.44	7862.41%	<b>900170</b>	<b>49481.44%</b>	142176.58	7808.26%
Double Dunk	-18.6	-16.4	23.93	1933.18%	<b>24</b>	<b>1936.36%</b>	23.7	1922.73%
Enduro	0	860.5	2367.71	275.16%	<b>14332.5</b>	<b>1665.60%</b>	2360.64	274.33%
Fishing Derby	-91.7	-38.8	<b>86.97</b>	<b>337.75%</b>	75	315.12%	77.05	319.00%
Freeway	0	29.6	32.59	110.10%	<b>34</b>	<b>114.86%</b>	33.97	114.76%
Frostbite	65.2	4334.7	<b>541280.88</b>	<b>12676.32%</b>	13792.4	321.52%	526239.5	12324.03%
Gopher	257.6	2412.5	117777.08	5453.59%	<b>488900</b>	<b>22675.87%</b>	119457.53	5531.58%
Gravitar	173	3351.4	19213.96	599.07%	6372.5	195.05%	<b>20875</b>	<b>651.33%</b>
Hero	1027	30826.4	114736.26	381.58%	37545.6	122.55%	<b>199880.6</b>	<b>667.31%</b>
Ice Hockey	-11.2	0.9	<b>63.64</b>	<b>618.51%</b>	47.53	485.37%	47.22	482.81%
Jamesbond	29	302.8	135784.96	49582.16%	<b>623300.5</b>	<b>227637.51%</b>	117009.92	42724.95%
Kangaroo	52	3035	<b>24034.16</b>	<b>803.96%</b>	14372.6	480.07%	17311.17	578.58%
Krull	1598	2665.5	251997.31	23456.61%	<b>593679.5</b>	<b>55464.31%</b>	155915.32	14455.96%
Kung Fu Master	258.5	22736.3	206845.82	919.07%	<b>1666665</b>	<b>7413.57%</b>	476539.53	2118.90%
Montezuma Revenge	0	4753.3	9352.01	196.75%	2500	52.60%	<b>12437</b>	<b>261.65%</b>
Ms Pacman	307.3	6951.6	<b>63994.44</b>	<b>958.52%</b>	31403	468.01%	29747.91	443.10%
Name This Game	2292.3	8049	54386.77	904.94%	<b>81473</b>	<b>1375.45%</b>	40077.73	656.37%
Phoenix	761.5	7242.6	908264.15	14002.29%	<b>999999</b>	<b>15417.71%</b>	849969.25	13102.83%
Pitfall	-229.4	6463.7	18756.01	283.66%	-1	3.41%	<b>46734.79</b>	<b>701.68%</b>
Pong	-20.7	14.6	20.67	117.20%	<b>21</b>	<b>118.13%</b>	19.31	113.34%
Private Eye	24.9	69571.3	79716.46	114.59%	15100	21.68%	<b>100798.9</b>	<b>144.90%</b>
Qbert	163.9	13455	<b>580328.14</b>	<b>4365.06%</b>	151730	1140.36%	238453.5	1792.85%
Riverraid	1338.5	17118	63318.67	392.79%	27964.3	168.74%	<b>90333.12</b>	<b>563.99%</b>
Road Runner	11.5	7845	243025.8	3102.24%	<b>999999</b>	<b>12765.53%</b>	399511.83	5099.90%
Robotank	2.2	11.9	127.32	1289.90%	<b>144</b>	<b>1461.86%</b>	114.46	1157.32%
Seaquest	68.4	42054.7	999997.63	2381.56%	<b>1000000</b>	<b>2381.57%</b>	960181.39	2286.73%
Skiing	-17098	-4336.9	-4202.6	101.05%	-5903.34	87.72%	<b>-3273.43</b>	<b>108.33%</b>
Solaris	1236.3	12326.7	<b>44199.93</b>	<b>387.39%</b>	10732.5	85.63% 28175.53	242.91%	
Space Invaders	148	1668.7	48680.86	3191.48%	<b>159999.6</b>	<b>10511.71%</b>	57828.45	3793.02%
Star Gunner	664	10250	839573.53	8751.40%	<b>999999</b>	<b>10424.94%</b>	264286.33	2750.08%
Surround	-10	6.5	9.5	118.18%	2.726	77.13%	<b>9.82</b>	<b>120.12%</b>
Tennis	-23.8	-8.3	23.84	307.35%	<b>24</b>	<b>308.39%</b>	22.79	300.58%
Time Pilot	3568	5229.2	405425.31	24190.78%	<b>531614</b>	<b>31787.02%</b>	404751.67	24150.23%
Tutankham	11.4	167.6	<b>2354.91</b>	<b>1500.33%</b>	436.2	271.96%	1030.27	652.29%
Up N Down	533.4	11693.2	623805.73	5584.98%	<b>999999</b>	<b>8955.95%</b>	524631	4696.30%
Venture	0	1187.5	2623.71	220.94%	2200	185.26%	<b>2859.83</b>	<b>240.83%</b>
Video Pinball	0	17667.9	992340.74	5616.63%	<b>999999</b>	<b>5659.98%</b>	617640.95	3495.84%
Wizard of Wor	563.5	4756.5	<b>157306.41</b>	<b>3738.20%</b>	118900	2822.24%	71942	1702.33%
Yars Revenge	3092.9	54576.9	998532.37	1933.49%	<b>998970</b>	<b>1934.34%</b>	633867.66	1225.19%
Zaxxon	32.5	9173.3	<b>249808.9</b>	<b>2732.54%</b>	241570.6	2642.42%	77942.17	852.33%
Mean HNS				4762.17%		<b>10077.52%</b>		4081.14
Median HNS				<b>1933.49%</b>		1665.60%		1225.19

Table 5: Score table of SOTA model-based algorithms on HNS(%).

Games Scale	MuZero 20B	HNS(%)	EfficientZero 100K	HNS(%)	Ours 1B	HNS(%)
Alien	<b>741812.63</b>	<b>10747.61%</b>	808.5	8.42%	279703.5	4050.37%
Amidar	<b>28634.39</b>	<b>1670.57%</b>	148.6	8.33%	12996.3	758.04%
Assault	<b>143972.03</b>	<b>27665.44%</b>	1263.1	200.29%	62025.7	11894.40%
Asterix	998425	12036.40%	25557.8	305.64%	<b>999999</b>	<b>12055.38%</b>
Asteroids	678558.64	1452.42%	N/A	N/A	<b>1106603.5</b>	<b>2369.60%</b>
Atlantis	1674767.2	10272.64%	N/A	N/A	<b>3824506.3</b>	<b>23560.59%</b>
Bank Heist	1278.98	171.17%	351	45.58%	<b>1410</b>	<b>188.90%</b>
Battle Zone	848623	2295.95%	13871.2	36.90%	<b>857369</b>	<b>2319.62%</b>
Beam Rider	454993.53	2744.92%	N/A	N/A	<b>457321</b>	<b>2758.97%</b>
Berzerk	<b>85932.6</b>	<b>3423.18%</b>	N/A	N/A	35340	1404.89%
Bowling	<b>260.13</b>	<b>172.26%</b>	N/A	N/A	233.1	152.62%
Boxing	<b>100</b>	<b>832.50%</b>	52.7	438.33%	<b>100</b>	<b>832.50%</b>
Breakout	<b>864</b>	<b>2994.10%</b>	414.1	1431.94%	<b>864</b>	<b>2994.10%</b>
Centipede	<b>1159049.27</b>	<b>11655.72%</b>	N/A	N/A	728080	7313.94%
Chopper Command	991039.7	15056.39%	1117.3	4.66%	<b>999999</b>	<b>15192.62%</b>
Crazy Climber	<b>458315.4</b>	<b>1718.06%</b>	83940.2	280.86%	233090	853.43%
Defender	839642.95	5291.18%	N/A	N/A	<b>995950</b>	<b>6279.56%</b>
Demon Attack	143964.26	7906.55%	13003.9	706.57%	<b>900170</b>	<b>49481.44%</b>
Double Dunk	23.94	1933.64%	N/A	N/A	<b>24</b>	<b>1936.36%</b>
Enduro	2382.44	276.87%	N/A	N/A	<b>14332.5</b>	<b>1665.60%</b>
Fishing Derby	<b>91.16</b>	<b>345.67%</b>	N/A	N/A	75	315.12%
Freeway	33.03	111.59%	21.8	73.65%	<b>34</b>	<b>114.86%</b>
Frostbite	<b>631378.53</b>	<b>14786.59%</b>	296.3	5.41%	13792.4	321.52%
Gopher	130345.58	6036.85%	3260.3	139.34%	<b>488900</b>	<b>22675.87%</b>
Gravitar	<b>6682.7</b>	<b>204.81%</b>	N/A	N/A	6372.5	195.05%
Hero	<b>49244.11</b>	<b>161.81%</b>	3915.9	9.69%	37545.6	122.55%
Ice Hockey	<b>67.04</b>	<b>646.61%</b>	N/A	N/A	47.53	485.37%
Jamesbond	41063.25	14986.94%	517	178.23%	<b>623300.5</b>	<b>227637.51%</b>
Kangaroo	<b>16763.6</b>	<b>560.23%</b>	724.1	22.53%	14372.6	480.07%
Krull	269358.27	25082.93%	5663.3	380.82%	<b>593679.5</b>	<b>55464.31%</b>
Kung Fu Master	204824	910.08%	30944.8	136.52%	<b>1666665</b>	<b>7413.57%</b>
Montezuma Revenge	0	0.00%	N/A	N/A	<b>2500</b>	<b>52.60%</b>
Ms Pacman	<b>243401.1</b>	<b>3658.68%</b>	1281.2	14.66%	31403	468.01%
Name This Game	<b>157177.85</b>	<b>2690.53%</b>	N/A	N/A	81473	1375.45%
Phoenix	955137.84	14725.53%	N/A	N/A	<b>999999</b>	<b>15417.71%</b>
Pitfall	<b>0</b>	<b>3.43%</b>	N/A	N/A	-1	3.41%
Pong	<b>21</b>	<b>118.13%</b>	20.1	115.58%	<b>21</b>	<b>118.13%</b>
Private Eye	<b>15299.98</b>	<b>21.96%</b>	96.7	0.10%	15100	21.68%
Qbert	72276	542.56%	14448.5	107.47%	<b>151730</b>	<b>1140.36%</b>
Riverraid	<b>323417.18</b>	<b>2041.12%</b>	N/A	N/A	27964.3	168.74%
Road Runner	613411.8	7830.48%	17751.3	226.46%	<b>999999</b>	<b>12765.53%</b>
Robotank	131.13	1329.18%	N/A	N/A	<b>144</b>	<b>1461.86%</b>
Seaquest	999976.52	2381.51%	1100.2	2.46%	<b>1000000</b>	<b>2381.57%</b>
Skiing	-29968.36	-100.86%	N/A	N/A	<b>-5903.34</b>	<b>87.72%</b>
Solaris	56.62	-10.64%	N/A	N/A	<b>10732.5</b>	<b>85.63%</b>
Space Invaders	74335.3	4878.50%	N/A	N/A	<b>159999.6</b>	<b>10511.71%</b>
Star Gunner	549271.7	5723.01%	N/A	N/A	<b>999999</b>	<b>10424.94%</b>
Surround	<b>9.99</b>	<b>121.15%</b>	N/A	N/A	2.726	77.13%
Tennis	0	153.55%	N/A	N/A	<b>24</b>	<b>308.39%</b>
Time Pilot	<b>476763.9</b>	<b>28485.19%</b>	N/A	N/A	531614	31787.02%
Tutankham	<b>491.48</b>	<b>307.35%</b>	N/A	N/A	436.2	271.96%
Up N Down	715545.61	6407.03%	17264.2	149.92%	<b>999999</b>	<b>8955.95%</b>
Venture	0.4	0.03%	N/A	N/A	<b>2200</b>	<b>185.26%</b>
Video Pinball	981791.88	5556.92%	N/A	N/A	<b>999999</b>	<b>5659.98%</b>
Wizard of Wor	<b>197126</b>	<b>4687.87%</b>	N/A	N/A	118900	2822.24%
Yars Revenge	553311.46	1068.72%	N/A	N/A	<b>998970</b>	<b>1934.34%</b>
Zaxxon	<b>725853.9</b>	<b>7940.46%</b>	N/A	N/A	241570.6	2642.42%
Mean HNS		4994.97%		194.3%		<b>10077.52%</b>
Median HNS		<b>2041.12%</b>		109%		1665.60%



Table 6: Score table of SOTA exploration-based algorithms on HNS(%).

Games Scale	Go-Explore 10B	HNS	Ours 1B	HNS
Alien	<b>959312</b>	<b>13899.77%</b>	279703.5	4050.37%
Amidar	<b>19083</b>	<b>1113.22%</b>	12996.3	758.04%
Assault	30773	5879.64%	<b>62025.7</b>	<b>11894.40%</b>
Asterix	999500	12049.37%	<b>999999</b>	<b>12055.38%</b>
Asteroids	112952	240.48%	<b>1106603.5</b>	<b>2369.60%</b>
Atlantis	286460	1691.24%	<b>3824506.3</b>	<b>23560.59%</b>
Bank Heist	<b>3668</b>	<b>494.49%</b>	1410	188.90%
Battle Zone	<b>998800</b>	<b>2702.36%</b>	857369	2319.62%
Beam Rider	371723	2242.15%	<b>457321</b>	<b>2758.97%</b>
Berzerk	<b>131417</b>	<b>5237.69%</b>	35340	1404.89%
Bowling	<b>247</b>	<b>162.72%</b>	233.1	152.62%
Boxing	91	757.50%	<b>100</b>	<b>832.50%</b>
Breakout	774	2681.60%	<b>864</b>	<b>2994.10%</b>
Centipede	613815	6162.78%	<b>728080</b>	<b>7313.94%</b>
Chopper Command	996220	15135.16%	<b>999999</b>	<b>15192.62%</b>
Crazy Climber	<b>235600</b>	<b>863.07%</b>	233090	853.43%
Defender	N/A	N/A	<b>995950</b>	<b>6279.56%</b>
Demon Attack	239895	13180.65%	<b>900170</b>	<b>49481.44%</b>
Double Dunk	<b>24</b>	<b>1936.36%</b>	<b>24</b>	<b>1936.36%</b>
Enduro	1031	119.81%	<b>14332.5</b>	<b>1665.60%</b>
Fishing Derby	67	300.00%	<b>75</b>	<b>315.12%</b>
Freeway	<b>34</b>	<b>114.86%</b>	<b>34</b>	<b>114.86%</b>
Frostbite	<b>999990</b>	<b>23420.19%</b>	13792.4	321.52%
Gopher	134244	6217.75%	<b>488900</b>	<b>22675.87%</b>
Gravitar	<b>13385</b>	<b>415.68%</b>	6372.5	195.05%
Hero	<b>37783</b>	<b>123.34%</b>	37545.6	122.55%
Ice Hockey	33	365.29%	<b>47.53</b>	<b>485.37%</b>
Jamesbond	200810	73331.26%	<b>623300.5</b>	<b>227637.51%</b>
Kangaroo	<b>24300</b>	<b>812.87%</b>	14372.6	480.07%
Krull	63149	5765.90%	<b>593679.5</b>	<b>55464.31%</b>
Kung Fu Master	24320	107.05%	<b>1666665</b>	<b>7413.57%</b>
Montezuma Revenge	<b>24758</b>	<b>520.86%</b>	2500	52.60%
Ms Pacman	<b>456123</b>	<b>6860.25%</b>	31403	468.01%
Name This Game	<b>212824</b>	<b>3657.16%</b>	81473	1375.45%
Phoenix	19200	284.50%	<b>999999</b>	<b>15417.71%</b>
Pitfall	<b>7875</b>	<b>121.09%</b>	-1	3.41%
Pong	<b>21</b>	<b>118.13%</b>	<b>21</b>	<b>118.13%</b>
Private Eye	<b>69976</b>	<b>100.58%</b>	15100	21.68%
Qbert	<b>999975</b>	<b>7522.41%</b>	151730	1140.36%
Riverraid	<b>35588</b>	<b>217.05%</b>	27964.3	168.74%
Road Runner	999900	12764.26%	<b>999999</b>	<b>12765.53%</b>
Robotank	143	1451.55%	<b>144</b>	<b>1461.86%</b>
Seaquest	539456	1284.68%	<b>1000000</b>	<b>2381.57%</b>
Skiing	<b>-4185</b>	<b>101.19%</b>	-5903.34	87.72%
Solaris	<b>20306</b>	<b>171.95%</b>	10732.5	85.63%
Space Invaders	93147	6115.54%	<b>159999.6</b>	<b>10511.71%</b>
Star Gunner	609580	6352.14%	<b>999999</b>	<b>10424.94%</b>
Surround	N/A	N/A	<b>2.726</b>	<b>77.13%</b>
Tennis	<b>24</b>	<b>308.39%</b>	<b>24</b>	<b>308.39%</b>
Time Pilot	183620	10838.67%	<b>531614</b>	<b>31787.02%</b>
Tutankham	<b>528</b>	<b>330.73%</b>	436.2	271.96%
Up N Down	553718	4956.94%	<b>999999</b>	<b>8955.95%</b>
Venture	<b>3074</b>	<b>258.86%</b>	2200	185.26%
Video Pinball	<b>999999</b>	<b>5659.98%</b>	<b>999999</b>	<b>5659.98%</b>
Wizard of Wor	<b>199900</b>	<b>4754.03%</b>	118900	2822.24%
Yars Revenge	<b>999998</b>	<b>1936.34%</b>	998970	1934.34%
Zaxxon	18340	200.28%	<b>241570.6</b>	<b>2642.42%</b>
Mean HNS		4989.31%		<b>10077.52%</b>
Median HNS		1451.55%		<b>1665.60%</b>

Table 7: Score Table of GDI-H<sup>3</sup> and LBC-BM (Ours) on HNS(%).

Games Scale	GDI-H <sup>3</sup> 200M	HNS	Ours 1B	HNS
Alien	48735	703.00%	<b>279703.5</b>	<b>4050.37%</b>
Amidar	1065	61.81%	<b>12996.3</b>	<b>758.04%</b>
Assault	<b>97155</b>	<b>18655.23%</b>	62025.7	11894.40%
Asterix	<b>999999</b>	<b>12055.38%</b>	<b>999999</b>	<b>12055.38%</b>
Asteroids	760005	1626.94%	<b>1106603.5</b>	<b>2369.60%</b>
Atlantis	<b>3837300</b>	<b>23639.67%</b>	3824506.3	23560.59%
Bank Heist	1380	184.84%	<b>1410</b>	<b>188.90%</b>
Battle Zone	824360	2230.29%	<b>857369</b>	<b>2319.62%</b>
Beam Rider	422390	2548.07%	<b>457321</b>	<b>2758.97%</b>
Berzerk	14649	579.46%	<b>35340</b>	<b>1404.89%</b>
Bowling	205.2	132.34%	<b>233.1</b>	<b>152.62%</b>
Boxing	<b>100</b>	<b>832.50%</b>	<b>100</b>	<b>832.50%</b>
Breakout	<b>864</b>	<b>2994.10%</b>	<b>864</b>	<b>2994.10%</b>
Centipede	195630	1949.80%	<b>728080</b>	<b>7313.94%</b>
Chopper Command	<b>999999</b>	<b>15192.62%</b>	<b>999999</b>	<b>15192.62%</b>
Crazy Climber	<b>241170</b>	<b>919.76%</b>	233090	853.43%
Defender	970540	6118.89%	<b>995950</b>	<b>6279.56%</b>
Demon Attack	787985	43313.70%	<b>900170</b>	<b>49481.44%</b>
Double Dunk	<b>24</b>	<b>1936.36%</b>	<b>24</b>	<b>1936.36%</b>
Enduro	14300	1661.82%	<b>14332.5</b>	<b>1665.60%</b>
Fishing Derby	65	296.22%	<b>75</b>	<b>315.12%</b>
Freeway	<b>34</b>	<b>114.86%</b>	<b>34</b>	<b>114.86%</b>
Frostbite	11330	263.84%	<b>13792.4</b>	<b>321.52%</b>
Gopher	473560	21964.01%	<b>488900</b>	<b>22675.87%</b>
Gravitar	5915	180.66%	<b>6372.5</b>	<b>195.05%</b>
Hero	<b>38225</b>	<b>124.83%</b>	37545.6	122.55%
Ice Hockey	47.11	481.90%	<b>47.53</b>	<b>485.37%</b>
Jamesbond	620780	226716.95%	<b>623300.5</b>	<b>227637.51%</b>
Kangaroo	14636	488.90%	14372.6	480.07%
Krull	594540	55544.92%	593679.5	55464.31%
Kung Fu Master	<b>1666665</b>	<b>7413.57%</b>	<b>1666665</b>	<b>7413.57%</b>
Montezuma Revenge	<b>2500</b>	<b>52.60%</b>	<b>2500</b>	<b>52.60%</b>
Ms Pacman	<b>11573</b>	<b>169.55%</b>	<b>31403</b>	<b>468.01%</b>
Name This Game	36296	590.68%	<b>81473</b>	<b>1375.45%</b>
Phoenix	959580	14794.07%	<b>999999</b>	<b>15417.71%</b>
Pitfall	-4.3	3.36	<b>-1</b>	<b>3.41%</b>
Pong	21	118.13%	<b>21</b>	<b>118.13%</b>
Private Eye	15100	21.68%	15100	21.68%
Qbert	28657	214.38%	<b>151730</b>	<b>1140.36%</b>
Riverraid	<b>28349</b>	<b>171.17%</b>	27964.3	168.74%
Road Runner	999999	12765.53%	<b>999999</b>	<b>12765.53%</b>
Robotank	113.4	1146.39%	<b>144</b>	<b>1461.86%</b>
Seaquest	<b>1000000</b>	<b>2381.57%</b>	<b>1000000</b>	<b>2381.57%</b>
Skiing	<b>-6025</b>	<b>86.77%</b>	-5903.34	87.72%
Solaris	9105	70.95%	<b>10732.5</b>	<b>85.63%</b>
Space Invaders	154380	10142.17%	<b>159999.6</b>	<b>10511.71%</b>
Star Gunner	677590	7061.61%	<b>999999</b>	<b>10424.94%</b>
Surround	2.606	76.40%	<b>2.726</b>	<b>77.13%</b>
Tennis	<b>24</b>	<b>308.39%</b>	<b>24</b>	<b>308.39%</b>
Time Pilot	450810	26924.45%	<b>531614</b>	<b>31787.02%</b>
Tutankham	418.2	260.44%	<b>436.2</b>	<b>271.96%</b>
Up N Down	966590	8656.58%	<b>999999</b>	<b>8955.95%</b>
Venture	2000	168.42%	<b>2200</b>	<b>185.26%</b>
Video Pinball	978190	5536.54%	<b>999999</b>	<b>5659.98%</b>
Wizard of Wor	63735	1506.59%	<b>118900</b>	<b>2822.24%</b>
Yars Revenge	968090	1874.36%	<b>998970</b>	<b>1934.34%</b>
Zaxxon	216020	2362.89%	<b>241570.6</b>	<b>2642.42%</b>
Mean HNS		4989.31%		<b>10077.52%</b>
Median HNS		1451.55%		<b>1665.60%</b>

### J.3 ATARI GAMES TABLE OF SCORES BASED ON HUMAN WORLD RECORDS

In this part, we detail the raw score of several representative SOTA algorithms, including the SOTA model-free algorithms, SOTA model-based algorithms, and other SOTA algorithms. Apart from the raw score, we also report the Human World Records for each game and **Human World Records Breakthrough (HWRB)** for each algorithm in Atari 57 games. More details on these algorithms can see Machado et al. (2018); Toromanoff et al. (2019).

Table 8: Score table of SOTA model-free algorithms on HWRB.

Games Scale	RND	Human World Records	AGENT57 100B	HWRB	Ours 1B	HWRB	MEME 1B	HWRB
Alien	227.8	251916	297638.17	1	279703.5	1	83683.43	0
Amidar	5.8	104159	29660.08	0	12996.3	0	14368.9	0
Assault	222.4	8647	67212.67	1	62025.7	1	46635.86	1
Asterix	210	1000000	991384.42	0	999999	0	769803.92	0
Asteroids	719	10506650	150854.61	0	1106603.5	0	364492.07	0
Atlantis	12850	10604840	1528841.76	0	3824506.3	0	1669226.33	0
Bank Heist	14.2	82058	23071.5	0	1410	0	87792.55	1
Battle Zone	236	801000	934134.88	1	857369	1	776770	0
Beam Rider	363.9	999999	300509.8	0	457321	0	51870.2	0
Berzerk	123.7	1057940	61507.83	0	35340	0	38838.35	0
Bowling	23.1	300	251.18	0	233.1	0	261.74	0
Boxing	0.1	100	100	1	100	1	99.85	0
Breakout	1.7	864	790.4	0	864	1	831.08	0
Centipede	2090.9	1301709	412847.86	0	728080	0	245892.18	0
Chopper Command	811	999999	999900	0	999999	1	912225	0
Crazy Climber	10780.5	219900	565909.85	1	233090	1	339274.67	1
Defender	2874.5	6010500	677642.78	0	995950	0	543979.5	0
Demon Attack	152.1	1556345	143161.44	0	900170	0	142176.58	0
Double Dunk	-18.6	21	23.93	1	24	1	23.7	1
Enduro	0	9500	2367.71	0	14332.5	1	2360.64	0
Fishing Derby	-91.7	71	86.97	1	75	1	77.05	1
Freeway	0	38	32.59	0	34	0	33.97	0
Frostbite	65.2	454830	541280.88	1	13792.4	0	526239.5	1
Gopher	257.6	355040	117777.08	0	488900	1	119457.53	0
Gravitar	173	162850	19213.96	0	6372.5	0	20875	0
Hero	1027	1000000	114736.26	0	37545.6	0	199880.6	0
Ice Hockey	-11.2	36	63.64	1	47.53	1	47.22	1
Jamesbond	29	45550	135784.96	1	623300.5	1	117009.92	1
Kangaroo	52	1424600	24034.16	0	14372.6	0	17311.17	0
Krull	1598	104100	251997.31	1	593679.5	1	155915.32	1
Kung Fu Master	258.5	1000000	206845.82	0	1666665	1	476539.53	0
Montezuma Revenge	0	1219200	9352.01	0	2500	0	12437	0
Ms Pacman	307.3	290090	63994.44	0	31403	0	29747.91	0
Name This Game	2292.3	25220	54386.77	1	81473	1	40077.73	1
Phoenix	761.5	4014440	908264.15	0	999999	0	849969.25	0
Pitfall	-229.4	114000	18756.01	0	-1	0	46734.79	0
Pong	-20.7	21	20.67	0	21	1	19.31	0
Private Eye	24.9	101800	79716.46	0	15100	0	100798.9	0
Qbert	163.9	2400000	580328.14	0	151730	0	238453.5	0
Riverraid	1338.5	1000000	63318.67	0	27964.3	0	90333.12	0
Road Runner	11.5	2038100	243025.8	0	999999	0	399511.83	0
Robotank	2.2	76	127.32	1	144	1	114.46	1
Seaquest	68.4	999999	999997.63	0	1000000	1	960181.39	0
Skiing	-17098	-3272	-4202.6	0	-5903.34	0	-3273.43	0
Solaris	1236.3	111420	44199.93	0	10732.5	0	28175.53	0
Space Invaders	148	621535	48680.86	0	159999.6	0	57828.45	0
Star Gunner	664	77400	839573.53	1	999999	1	264286.33	1
Surround	-10	9.6	9.5	0	2.726	0	9.82	1
Tennis	-23.8	21	23.84	1	24	1	22.79	1
Time Pilot	3568	65300	405425.31	1	531614	1	404751.67	1
Tutankham	11.4	5384	2354.91	0	436.2	0	1030.27	0
Up n Down	533.4	82840	623805.73	1	999999	1	524631	1
Venture	0	38900	2623.71	0	2200	0	2859.83	0
Video Pinball	0	89218328	992340.74	0	999999	0	617640.95	0
Wizard of Wor	563.5	395300	157306.41	0	118900	0	71942	0
Yars Revenge	3092.9	15000105	998532.37	0	998970	0	633867.66	0
Zaxxon	32.5	83700	249808.9	1	241570.6	1	77942.17	0
$\Sigma$ HWRB				18		24		16

Table 9: Score table of SOTA model-based algorithms on HWRB.

Games Scale	MuZero 20B	HWRB	EfficientZero 100K	HWRB	Ours 1B	HWRB
Alien	741812.63	1	808.5	0	279703.5	1
Amidar	28634.39	0	148.6	0	12996.3	0
Assault	143972.03	1	1263.1	0	62025.7	1
Asterix	998425	0	25557.8	0	999999	0
Asteroids	678558.64	0	N/A	N/A	1106603.5	0
Atlantis	1674767.2	0	N/A	N/A	3824506.3	0
Bank Heist	1278.98	0	351	0	1410	0
Battle Zone	848623	1	13871.2	0	857369	1
Beam Rider	454993.53	0	N/A	N/A	457321	0
Berzerk	85932.6	0	N/A	N/A	35340	0
Bowling	260.13	0	N/A	N/A	233.1	0
Boxing	100	1	52.7	0	100	1
Breakout	864	1	414.1	0	864	1
Centipede	1159049.27	0	N/A	N/A	728080	0
Chopper Command	991039.7	0	1117.3	0	999999	1
Crazy Climber	458315.4	1	83940.2	0	233090	1
Defender	839642.95	0	N/A	N/A	995950	0
Demon Attack	143964.26	0	13003.9	0	900170	0
Double Dunk	23.94	1	N/A	N/A	24	1
Enduro	2382.44	0	N/A	N/A	14332.5	1
Fishing Derby	91.16	1	N/A	N/A	75	1
Freeway	33.03	0	21.8	0	34	0
Frostbite	631378.53	1	296.3	0	13792.4	0
Gopher	130345.58	0	3260.3	0	488900	1
Gravitar	6682.7	0	N/A	N/A	6372.5	0
Hero	49244.11	0	3915.9	0	37545.6	0
Ice Hockey	67.04	1	N/A	N/A	47.53	1
Jamesbond	41063.25	0	517	0	623300.5	1
Kangaroo	16763.6	0	724.1	0	14372.6	0
Krull	269358.27	1	5663.3	0	593679.5	1
Kung Fu Master	204824	0	30944.8	0	1666665	1
Montezuma Revenge	0	0	N/A	N/A	2500	0
Ms Pacman	243401.1	0	1281.2	0	31403	0
Name This Game	157177.85	1	N/A	N/A	81473	1
Phoenix	955137.84	0	N/A	N/A	999999	0
Pitfall	0	0	N/A	N/A	-1	0
Pong	21	1	20.1	0	21	1
Private Eye	15299.98	0	96.7	0	15100	0
Qbert	72276	0	14448.5	0	151730	0
Riverraid	323417.18	0	N/A	N/A	27964.3	0
Road Runner	613411.8	0	17751.3	0	999999	0
Robotank	131.13	1	N/A	N/A	144	1
Seaquest	999976.52	0	1100.2	0	1000000	1
Skiing	-29968.36	0	N/A	N/A	-5903.34	0
Solaris	56.62	0	N/A	N/A	10732.5	0
Space Invaders	74335.3	0	N/A	N/A	159999.6	0
Star Gunner	549271.7	1	N/A	N/A	999999	1
Surround	9.99	1	N/A	N/A	2.726	0
Tennis	0	0	N/A	N/A	24	1
Time Pilot	476763.9	1	N/A	N/A	531614	1
Tutankham	491.48	0	N/A	N/A	436.2	0
Up N Down	715545.61	1	17264.2	0	999999	1
Venture	0.4	0	N/A	N/A	2200	0
Video Pinball	981791.88	0	N/A	N/A	999999	0
Wizard of Wor	197126	0	N/A	N/A	118900	0
Yars Revenge	553311.46	0	N/A	N/A	998970	0
Zaxxon	725853.9	1	N/A	N/A	241570.6	1
$\Sigma$ HWRB		19		0		24

Table 10: Score table of SOTA exploration-based algorithms on HWRB.

Games Scale	Go-Explore 10B	HWRB	Ours 1B	HWRB
Alien	959312	1	279703.5	1
Amidar	19083	0	12996.3	0
Assault	30773	1	62025.7	1
Asterix	999500	0	999999	0
Asteroids	112952	0	1106603.5	0
Atlantis	286460	0	3824506.3	0
Bank Heist	3668	0	1410	0
Battle Zone	998800	1	857369	1
Beam Rider	371723	0	457321	0
Berzerk	131417	0	35340	0
Bowling	247	0	233.1	0
Boxing	91	0	100	1
Breakout	774	0	864	1
Centipede	613815	0	728080	0
Chopper Command	996220	0	999999	1
Crazy Climber	235600	1	233090	1
Defender	N/A	N/A	995950	0
Demon Attack	239895	0	900170	0
Double Dunk	24	1	24	1
Enduro	1031	0	14332.5	1
Fishing Derby	67	0	75	1
Freeway	34	0	34	0
Frostbite	999990	1	13792.4	0
Gopher	134244	0	488900	1
Gravitar	13385	0	6372.5	0
Hero	37783	0	37545.6	0
Ice Hockey	33	0	47.53	1
Jamesbond	200810	1	623300.5	1
Kangaroo	24300	0	14372.6	0
Krull	63149	0	593679.5	1
Kung Fu Master	24320	0	1666665	1
Montezuma Revenge	24758	0	2500	0
Ms Pacman	456123	1	31403	0
Name This Game	212824	1	81473	1
Phoenix	19200	0	999999	0
Pitfall	7875	0	-1	0
Pong	21	1	21	1
Private Eye	69976	0	15100	0
Qbert	999975	0	151730	0
Riverraid	35588	0	27964.3	0
Road Runner	999900	0	999999	0
Robotank	143	1	144	1
Sequest	539456	0	1000000	1
Skiing	-4185	0	-5903.34	0
Solaris	20306	0	10732.5	0
Space Invaders	93147	0	159999.6	0
Star Gunner	609580	1	999999	1
Surround	N/A	N/A	2.726	0
Tennis	24	1	24	1
Time Pilot	183620	1	531614	1
Tutankham	528	0	436.2	0
Up N Down	553718	1	999999	1
Venture	3074	0	2200	0
Video Pinball	999999	0	999999	0
Wizard of Wor	199900	0	118900	0
Yars Revenge	999998	0	998970	0
Zaxxon	18340	0	241570.6	1
$\sum$ HWRB		15		<b>24</b>

Table 11: Score Table of GDI-H<sup>3</sup> and LBC-BM (Ours) on HWRB.

Games Scale	GDI-H <sup>3</sup> 200M	HWRB	Ours 1B	HWRB
Alien	48735	0	<b>279703.5</b>	1
Amidar	1065	0	<b>12996.3</b>	0
Assault	<b>97155</b>	1	62025.7	1
Asterix	<b>999999</b>	0	<b>999999</b>	0
Asteroids	760005	0	<b>1106603.5</b>	0
Atlantis	<b>3837300</b>	0	3824506.3	0
Bank Heist	1380	0	<b>1410</b>	0
Battle Zone	824360	1	<b>857369</b>	1
Beam Rider	422390	0	<b>457321</b>	0
Berzerk	14649	0	<b>35340</b>	0
Bowling	205.2	0	<b>233.1</b>	0
Boxing	<b>100</b>	1	<b>100</b>	1
Breakout	<b>864</b>	1	<b>864</b>	1
Centipede	195630	0	<b>728080</b>	0
Chopper Command	<b>999999</b>	1	<b>999999</b>	1
Crazy Climber	<b>241170</b>	1	233090	1
Defender	970540	0	<b>995950</b>	0
Demon Attack	787985	0	<b>900170</b>	0
Double Dunk	<b>24</b>	1	<b>24</b>	1
Enduro	14300	1	<b>14332.5</b>	1
Fishing Derby	65	0	<b>75</b>	1
Freeway	<b>34</b>	0	<b>34</b>	0
Frostbite	11330	0	<b>13792.4</b>	0
Gopher	473560	1	<b>488900</b>	1
Gravitar	5915	0	<b>6372.5</b>	0
Hero	<b>38225</b>	0	37545.6	0
Ice Hockey	47.11	0	<b>47.53</b>	1
Jamesbond	620780	1	<b>623300.5</b>	1
Kangaroo	14636	0	14372.6	0
Krull	594540	1	593679.5	1
Kung Fu Master	<b>1666665</b>	1	<b>1666665</b>	1
Montezuma Revenge	<b>2500</b>	0	<b>2500</b>	0
Ms Pacman	<b>11573</b>	0	<b>31403</b>	0
Name This Game	36296	1	<b>81473</b>	1
Phoenix	959580	0	<b>999999</b>	0
Pitfall	-4.3	0	<b>-1</b>	0
Pong	21	1	<b>21</b>	1
Private Eye	15100	0	15100	0
Qbert	28657	0	<b>151730</b>	0
Riverraid	<b>28349</b>	0	27964.3	0
Road Runner	999999	0	<b>999999</b>	0
Robotank	113.4	0	<b>144</b>	0
Seaquest	<b>1000000</b>	1	<b>1000000</b>	1
Skiing	<b>-6025</b>	0	-5903.34	0
Solaris	9105	0	<b>10732.5</b>	0
Space Invaders	154380	0	<b>159999.6</b>	0
Star Gunner	677590	1	<b>999999</b>	1
Surround	2.606	0	<b>2.726</b>	0
Tennis	<b>24</b>	1	<b>24</b>	1
Time Pilot	450810	1	<b>531614</b>	1
Tutankham	418.2	0	<b>436.2</b>	0
Up N Down	966590	1	<b>999999</b>	1
Venture	2000	0	<b>2200</b>	0
Video Pinball	978190	0	<b>999999</b>	0
Wizard of Wor	63735	0	<b>118900</b>	0
Yars Revenge	968090	0	<b>998970</b>	0
Zaxxon	216020	1	<b>241570.6</b>	1
$\sum$ HWRB		22		<b>24</b>

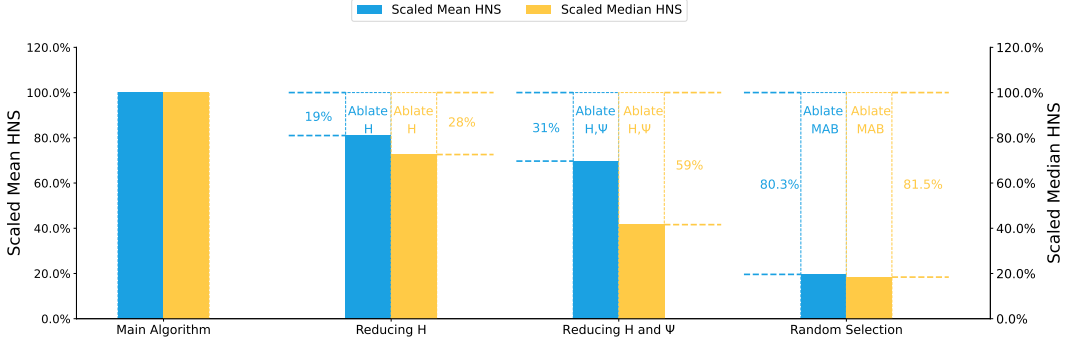


Figure 6: Ablation Results on Atari Benchmark (Machado et al., 2018). All the results are scaled by that of our main algorithm to improve readability. In these figures, we sequentially demonstrate how much performance (%) will degrade after ablating each component of LBC.

## K ABLATION STUDY

In this section, we will demonstrate the settings of our ablation studies first. Then we will introduce the algorithms of the ablation study, which has been concluded in Tab. 12. After that, we will introduce the ablation study results and case studies of t-SNE, including the results on the Atari benchmark in Fig. 6 and t-SNE analysis in App. K.3.

### K.1 ABLATION STUDY SETUP

We summarized all the algorithms of the ablation study in Tab. 12. All algorithms are tested in the same experimental setup. More details on these experimental setups can see App. H. The hyper-parameters can see App. I.

Table 12: Algorithms of Ablation Study.

Algorithm	Main Algorithm	Reducing $\mathbf{H}$	Reducing $\mathbf{H}$ and $\Psi$	Random Selection
Ablation Variables	Baseline	$\mathbf{H}$	$\Psi$ and $\mathbf{H}$	Meta-Controller ( $\Psi$ )
$\mathcal{F}_\Psi$	$\sum_{i=1}^3 \omega_i \text{Softmax}_{\tau_i}(\Phi_{\mathbf{h}_i})$	$\sum_{i=1}^3 \omega_i \text{Softmax}_{\tau_i}(\Phi_{\mathbf{h}_i})$	$\text{Softmax}_\tau(\Phi_{\mathbf{h}})$	$\sum_{i=1}^3 \omega_i \text{Softmax}_{\tau_i}(\Phi_{\mathbf{h}_i})$
$\Phi_{\mathbf{H}}$	$(\Phi_{\mathbf{h}_1}, \dots, \Phi_{\mathbf{h}_3})$	$\Phi_{\mathbf{h}_1}$	$\Phi_{\mathbf{h}_1}$	$(\Phi_{\mathbf{h}_1}, \dots, \Phi_{\mathbf{h}_3})$
$\mathbf{h}_i$	$(\gamma_i, \mathcal{RS}_i)$	$(\gamma_i, \mathcal{RS}_i)$	$(\gamma_i, \mathcal{RS}_i)$	$(\gamma_i, \mathcal{RS}_i)$
$\mathbf{H}$	$\{\mathbf{h}_i   i = 1, 2, 3\}$	$\{\mathbf{h}_i   i = 1\}$	$\{\mathbf{h}_i   i = 1\}$	$\{\mathbf{h}_i   i = 1, 2, 3\}$
$\psi_i$	$(\omega_1, \tau_1, \omega_2, \tau_2, \omega_3, \tau_3)$	$(\omega_1, \tau_1, \omega_2, \tau_2, \omega_3, \tau_3)$	$(\tau)$	$(\omega_1, \tau_1, \omega_2, \tau_2, \omega_3, \tau_3)$
$\Psi$	$\{\psi_i   i = 1, \dots, \infty\}$	$\{\psi_i   i = 1, \dots, \infty\}$	$\{\psi_i   i = 1, \dots, \infty\}$	$\{\psi_i   i = 1, \dots, \infty\}$
$ \mathbf{H} $	3	1	1	3
$ \Psi $	$\infty$	$\infty$	$\infty$	$\infty$
Category	LBC- $\mathbf{H}_3^2$ - $\Psi_\infty^6$	LBC- $\mathbf{H}_1^2$ - $\Psi_\infty^6$	LBC- $\mathbf{H}_1^2$ - $\Psi_\infty^1$	LBC- $\mathbf{H}_3^2$ - $\Psi_\infty^6$
$ \mathbf{M}_{\mathbf{H}, \Psi} $	$\infty$	$\infty$	$\infty$	$\infty$
Meta-Controller ( $\Psi$ )	MAB	MAB	MAB	Random Selection

### K.2 ABLATION STUDY RESULTS

The ablation study results can be found in Fig. 6. From left to right, the behavior space of the first three algorithms decreases in turn, and the final performance of these three algorithms decreases in turn. We can draw the following corollary:

**Corollary 5** (Smaller Behavior Space, Lower Final Performance). *Given any RL methods, assuming each behavior can be visited infinitely, decreasing the behavior space and keeping other conditions unchanged will degrade the final performance of the algorithm, and vice versa.*

The behavior space of Random Selection is the same as our main algorithm. Obviously, the appropriate behaviors fail to be selected with a random selection, resulting in a great decrease of the performance in limited training frames.



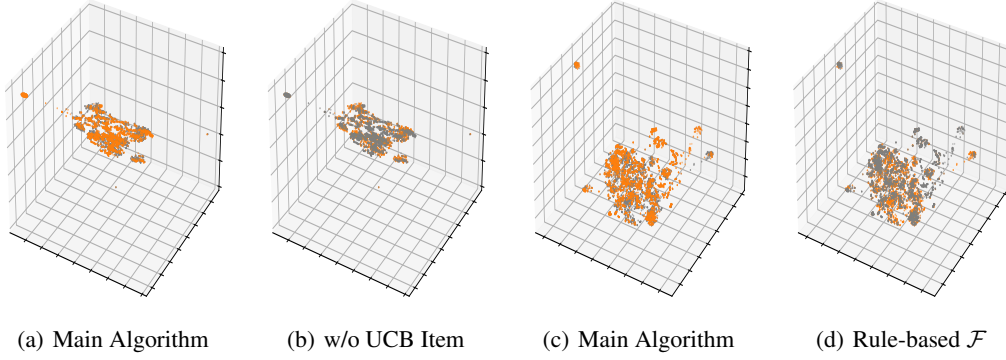


Figure 7: Visualizing Behavior Diversity via t-SNE. (a) and (b) are drawn from the t-SNE analysis of visited states (points highlighted with  $\bullet$ ) in Chopper Command, and (c) and (d) are drawn the t-SNE analysis of visited states in Atlantis.

### K.3 T-SNE ANALYSIS

In this paper, we adopt the ucb-score to encourage the actors to try more different behavior. To demonstrate the effectiveness of the ucb-score, we conduct the t-SNE analysis of the methods removing the ucb-score in 1 and 2 of Fig. 7.

To demonstrate that the behavior diversity can be boosted by our algorithm, we conducted the t-SNE analysis of the methods with rule-based  $\mathcal{F}$  in 2 and 3 of Fig. Fig. 7.

From (a) and (b) of Fig. 7, we find that removing the UCB item (i.e.,  $\sqrt{\frac{\log(1+\sum_{j \neq k}^K N_{\Psi_j})}{1+N_{\Psi_k}}}$ ) from the optimization target of behavior selection, the behavior diversity fade away. It can prove the effectiveness of the diversity control of our methods.

From (c) and (d) of Fig. 7, we find that compared with the rule-based  $\mathcal{F}$ , our method can acquire a diverse set of behaviors though we do not contain a diversity-based multi-objective model training which confirms the Corollary 2.

## L MODEL ARCHITECTURE

Since the network structure is not the focus of our work, we keep most of the components of the network, e.g., the LSTM Core, RL Head and Convolutional Layers the same as that of Agent57 (Badia et al., 2020a). To improve the reproducibility of our work, we still summarize our model architecture of our main algorithm in detail in Fig. 8. Wherein,  $\mathbf{h}_j^t$  is the hyper-parameters (e.g., discounted factor  $\gamma$ ) of policy model  $j$  and  $\psi_t$  is the parameters of the constructed behavior space (e.g.,  $\epsilon$  in  $\epsilon$ -greedy). More details on the hyper-parameters of each policy model can see App. I.  $\psi_t$  will be adaptively selected to control the behaviors across learning process via MAB. More implementation details on the MAB can see App. E. It is worth noting that our framework is not limited to an implementation of report in our body. In Fig. 8, we show a general way of integrating multiple policy models and automatically adjusting the proportion of any multiple policy models in the ensembled behavior policy.

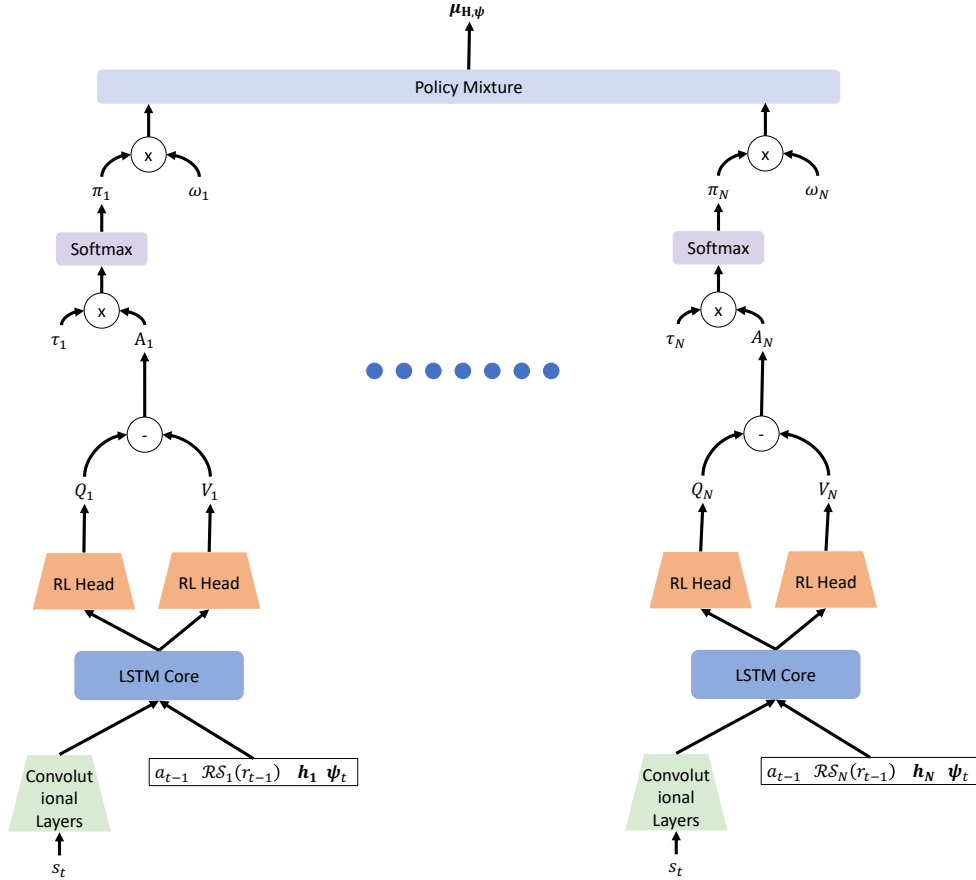


Figure 8: Model Architecture of our main algorithm.

## M SUPPLEMENTARY MATERIAL

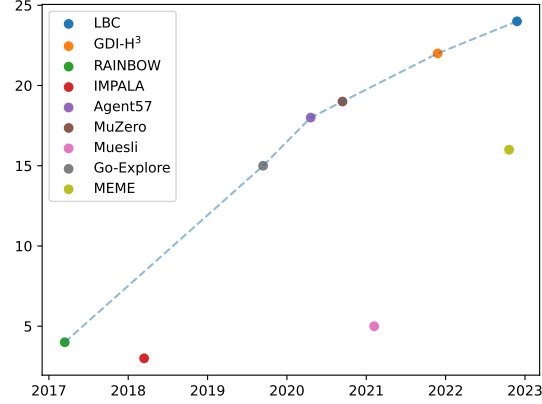


Figure 9: Human World Records Breakthrough of Atari RL Benchmarks.

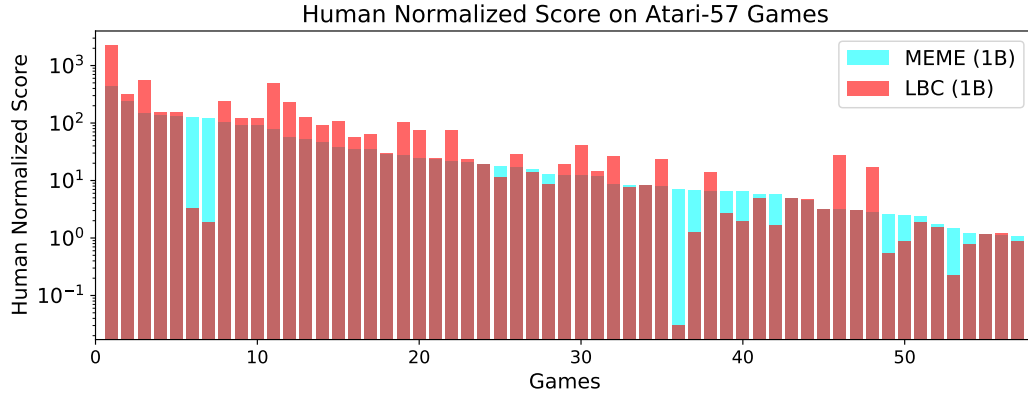
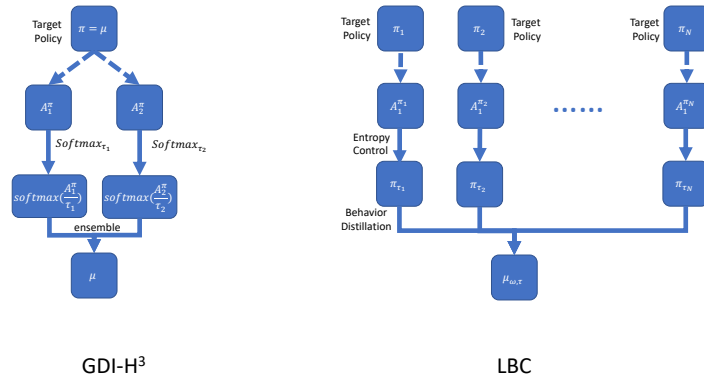


Figure 10: Performance Comparison between MEME and LBC based on HNS.(log scale)

Figure 11: Different Learning Framework of GDI-H<sup>3</sup> and LBC

## N CASE STUDY: KL DIVERGENCE

In this section, to further investigate the cause of the degradation phenomenon of the behavior space in GDI-H<sup>3</sup> (i.e., due to a same learned policy  $\pi$  for  $A_1^\pi$  and  $A_2^\pi$  under different reward shaping) and demonstrate that in our behavior space (i.e., different learned policies for  $A_1^{\pi_1}$  and  $A_2^{\pi_2}$ ), the diversity can be maintained since different policy models are distinguishable across learning. For fairness, we designed two implementations to explore the degradation phenomenon in the behavior space of GDI-H<sup>3</sup> including: i) an implementation with two different learned policies under different reward shaping (yellow in Fig. 12) ii) an implementation that learns two advantage functions of a same target policy (i.e., the behavior policy) under different reward shaping as GDI-H<sup>3</sup> (blue in Fig. 12). The learning framework of these two implementations can be found in 11.

For a fair comparison, we keep the two reward shaping the same as used in GDI-H<sup>3</sup>, namely, i)  $\log(\text{abs}(r) + 1.0) \cdot (2 \cdot 1_{\{r \geq 0\}} - 1_{\{r < 0\}})$  for  $A_1$  and ii)  $\text{sign}(r) \cdot ((\text{abs}(r) + 1.0)^{0.25} - 1.0) + 0.001 \cdot r$  for  $A_2$ .

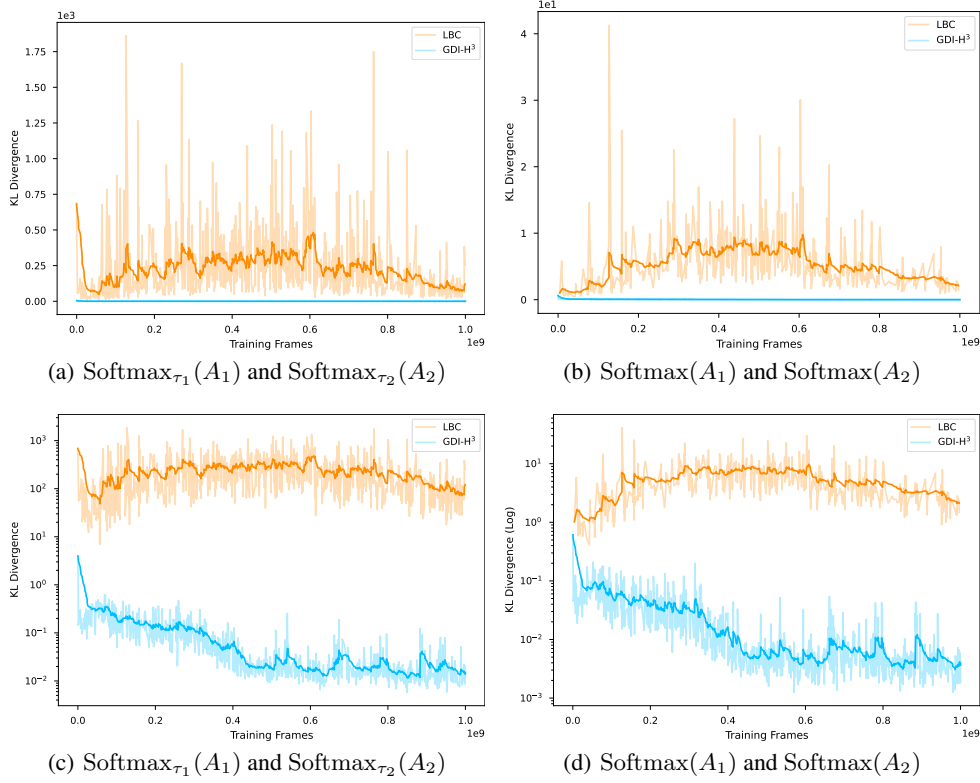


Figure 12: KL Divergence of GDI-H<sup>3</sup> and LBC in Chopper Command (Smoothed by 0.9 for the ease of reading).

From Fig. 12, we can find the distance between  $\text{Softmax}_{\tau_1}(A_1^\pi)$  and  $\text{Softmax}_{\tau_2}(A_2^\pi)$  of  $A_1^\pi$  and  $A_2^\pi$  decrease rapidly in GDI-H<sup>3</sup> while LBC can maintain a more diverse set of policies. The optional behaviors for each actor gradually diminish, and the behavior space of GDI-H<sup>3</sup> degenerates across the learning process. In contrast, LBC can maintain the capacity of the behavior space and avoid degradation since LBC maintains a population of different policy models (Corresponding to a population of different policies.)