

No Prior Mask: Eliminate Redundant Action for Deep Reinforcement Learning

Dianyu Zhong*, Yiqin Yang*, Qianchuan Zhao†

Department of Automation, Tsinghua University
{zhongdy19, yangyiqi19}@mails.tsinghua.edu.cn, zhaoqc@tsinghua.edu.cn

Abstract

The large action space is one fundamental obstacle to deploying Reinforcement Learning methods in the real world. The numerous redundant actions will cause the agents to make repeated or invalid attempts, even leading to task failure. Although current algorithms conduct some initial explorations for this issue, they either suffer from rule-based systems or depend on expert demonstrations, which significantly limits their applicability in many real-world settings. In this work, we examine the theoretical analysis of what action can be eliminated in policy optimization and propose a novel redundant action filtering mechanism. Unlike other works, our method constructs the similarity factor by estimating the distance between the state distributions, which requires no prior knowledge. In addition, we combine the modified inverse model to avoid extensive computation in high-dimensional state space. We reveal the underlying structure of action spaces and propose a simple yet efficient redundant action filtering mechanism named No Prior Mask (NPM) based on the above techniques. We show the superior performance of our method by conducting extensive experiments on high-dimensional, pixel-input, and stochastic problems with various action redundancy tasks. Our code is public online at <https://github.com/zhongdy15/npm>.

Introduction

Although Reinforcement Learning (RL) methods have found tremendous success in challenging domains, including video games (Mnih et al. 2015) and classic games (Silver et al. 2016; Yang et al. 2021; Ma et al. 2021), deploying them in real-world applications is still limited. One of the main challenges towards that goal is dealing with large action spaces. For example, maximizing long-term portfolio value using various trading strategies in the finance system (Jiang, Xu, and Liang 2017), regulating the voltage level of all the units in the power system (Glavic, Fonteneau, and Ernst 2017), and generating recommendation sequences from an extensive collection of options (Sidney et al. 2005) pose significant challenges due to the numerous actions involved. The redundant actions in these tasks will cause the

agents to make repeated invalid attempts, even leading to task failure.

To solve this issue, current researchers attempt to utilize the underlying structure in the action space to filter the irrelevant actions. One paradigm for algorithm design incorporates the prior knowledge into the learning algorithm. There are several works in the existing empirical literature, especially in challenging domains such as Dota2 (Berner et al. 2019), StarCraftII (Vinyals et al. 2017) and Honor of Kings (Ye et al. 2020). These approaches build simple hand-crafted action masks to determine the availability of an action, such as checking if a valid target is nearby or if the ability is on cooldown. Ye et al. (2020) demonstrates that action mask can largely reduce the training time and Huang et al. (2021) conducts ablation studies on action mask, revealing its substantial potential in enhancing performance. However, despite the interpretability and efficiency of these methods, it remains exceptionally challenging to manually construct an accurate relationship between actions, particularly in high-dimensional and state-dependent scenarios with limited expert information.

On the other hand, recent advances were made in autonomously learning such underlying structures in the action space. For example, Dulac-Arnold et al. (2015) leverages prior information about the actions to embed them in a continuous space. Zahavy et al. (2018a) utilizes an external elimination signal provided by the environment to predict invalid actions. Tennenholtz and Mannor (2019) combines existing methods in natural language processing (NLP) to group similar actions and utilizing relations between different actions from expert data. Chandak et al. (2019) provides an algorithm to learn and use action representations on large-scale real-world problems. Baram, Tennenholtz, and Mannor (2021) proposes maximizing the next states' entropy instead of action entropy to minimize action redundancy.

Nevertheless, the aforementioned methods either rely heavily on prior information or are tightly coupling with the policy, which significantly limits their applicability in various real-world scenarios. Therefore, the shortcoming of the current redundant action filtering mechanism naturally leads to the following question:

How to learn the underlying structure accurately in the action space without prior knowledge?

*These authors contributed equally.

†Advising.

To answer this question, we start with the theoretical analysis of what action can be eliminated in policy optimization. Intuitively, actions with the same state transition should be classified into similar actions. Based on this analysis, we propose a redundant action classifying mechanism, which constructs the similarity factor by estimating the distance between the state distributions. In addition, we combine the modified inverse model to avoid extensive computation in the high-dimensional state space, which makes our algorithm simple yet efficient. We conduct extensive experiments on various domains, including synthetic action redundancy, combined action redundancy and state-dependent action redundancy. The experimental results show that our method performs better than other baselines.

Important Contributions:

1. We propose a mask-based reinforcement learning framework without any prior knowledge, which can scale up to high-dimensional pixel-based observations. To the best of our knowledge, our work is the first study of constructing no prior state-dependent action mask, leading our method to perform better than other methods.
2. We give a novel theoretical analysis revealing what kind of mask is reasonable and feasible. Our work stands out as the first attempt to bridge the gap between the practical technique of action mask and the theoretical foundations of Markov Decision Processes (MDPs).
3. The mask learning phase is reward-free, indicating our mask mechanism is unrelated to policy and easily transfers across multi-tasks. We validate our method’s effectiveness, simplicity, and transfer ability on both single-tasks and multi-tasks.

Related Work

RL in Large Action Space There are extensive works to apply RL algorithms to environments with large action spaces, which can be roughly divided into using prior knowledge and without prior knowledge. On the one hand, researchers find that better performances can be achieved when actions are factored into their primary categories (Sharma et al. 2017) according to the underlying compositional structure in Atari 2600 (Bellemare et al. 2013). Dulac-Arnold et al. (2015) attempt to embed actions with prior information in a continuous space and find the closest neighbors of promising ones. Zahavy et al. (2018b) directly learn about redundant or irrelevant actions with external elimination signals provided by the environment and eliminate them from being sampled in text-based games. Tennenholtz and Mannor (2019) adopts the Negative-Sampling procedure with expert demonstrations to gain a deeper understanding of actions. However, valuable prior information is scarce and expensive, which limits the application of these methods.

Unlike prior-based methods, learning in large action spaces without prior knowledge is hard. Chandak et al. (2019) showed how to learn and use representations without prior knowledge in which an embedding is used as part

of the policy’s structure to train an agent. Baram, Tennenholtz, and Mannor (2021) explored to incorporate maximum entropy between distributions over the next states and penalized it in reward as action redundancy term. However, these methods have a tight coupling with the policy, limiting the optimized policy’s generalization performance across multi-tasks.

Meanwhile, there are several works on discretizing continuous actions and learning the structure within the action space. Metz et al. (2017) introduces a creative approach to discretizing high-dimensional continuous actions through the sequential combination of one-dimensional discrete actions. Tavakoli, Pardo, and Kormushev (2018) explores the intriguing concept of utilizing hypergraphs to combine actions within the action space. Tavakoli, Fatemi, and Kormushev (2021) presents an innovative neural architecture that capitalizes on a shared decision module and distinct network branches for each action dimension. Seyde et al. (2022) offers a new lens on Q-learning by amalgamating action discretization with value decomposition. By framing single-agent control as a cooperative multi-agent RL problem, this paper contributes a fresh perspective to addressing continuous control challenges. Dadashi et al. (2021) introduces an elegant approach by discretizing continuous actions based on prior knowledge and subsequently learning exclusively within the corresponding discrete action space during online training.

Action Mask Action mask has achieved wide attention in dealing with numerous invalid actions tasks, such as Dota 2 (Berner et al. 2019), Starcraft II (Vinyals et al. 2017) and Honor of Kings (Ye et al. 2020). These methods build simple action filters manually, determining whether the action is available. At each time step, the algorithm restricts available actions from the full action space based on the constructed filters and optimizes policy in a sub-action space. Ye et al. (2020) demonstrates that action mask can largely reduce the training time from 80 hours to 65 hours. Huang et al. (2021) reveals that action mask considerably improves performance on real-time strategy (RTS) games. However, learning a valid action mask without prior knowledge and promoting algorithm performance is not explored well.

MDP Homomorphisms and Bisimulation There are several works in MDP homomorphisms and action abstractions. Taylor, Precup, and Panagaden (2008); Ravindran (2004) define a behavior similarity metric for measuring state-action pair similarity and obtains error bounds. Panagaden et al. (2023) extend the definition of MDP homomorphisms to the setting of continuous state and action spaces and derive a policy gradient theorem on the abstract MDP. Thomas et al. (2017) learn controllable factors along with policies. Sallans and Hinton (2004) use an undirected graphical model to approximate the value function with factored states and actions. However, the methods utilizing the Hausdorff metric (Taylor, Precup, and Panagaden 2008) or the Kantorovich metric (Panangaden et al. 2023) lead to unaffordable computational complexity in high-dimensional states, which limits the applicability to broad scenarios.

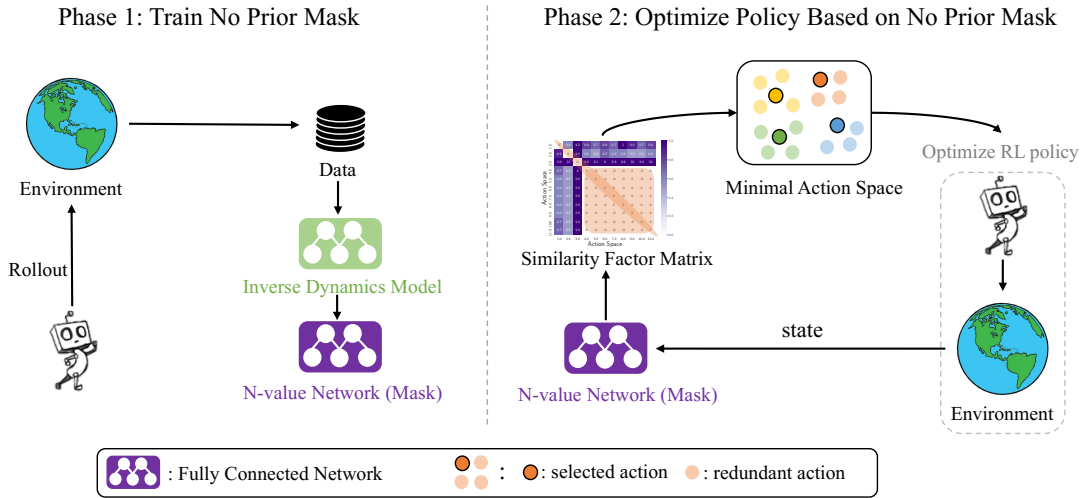


Figure 1: The framework of NPM. Our method consists of two phases: The first is to train the no prior mask. The second phase uses the trained mask to calculate the similarity factor matrix and eliminate the redundant actions.

Action Priors There are several works in accelerating learning by reducing the number of actions in RL. Sherstov and Stone (2005); Fernández and Veloso (2006); Rosman and Ramamoorthy (2012); Biza et al. (2021) utilize a set of different optimal policies from similar tasks to construct a small fraction of the domain’s action set or bias exploration away from sub-optimal actions. Rosman and Ramamoorthy (2015) define action priors over observation features rather than states, and Abel et al. (2015) model the optimal actions through experience to avoid searching for sub-optimal actions. Yan, Schwing, and Wang (2022) combine multiple implicit priors with task-specific demonstrations to outperform in challenging environments. However, the research on Action Priors focus on reducing sub-optimal actions instead of redundant actions discussed in our work.

Background

We consider the Markov Decision Processes (MDPs), defined by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma)$, where \mathcal{S} is a state space, \mathcal{A} is an action space, $\gamma \in [0, 1)$ is the discount factor and $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S}), r : \mathcal{S} \rightarrow [0, r_{\max}]$ are the transition function and reward function, respectively. In this work, we restrict our focus to MDPs with finite action sets, and $|\mathcal{A}|$ denotes the size of the action set. We also assume a fixed distribution $\mu_0 \in \Delta(\mathcal{S})$ as the initial state distribution. The goal of an RL agent is to learn a policy $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ under dataset \mathcal{D} , which maximizes the expectation of a discounted cumulative reward: $J(\pi) = \mathbb{E}_{\mu_0, \pi} [\sum_{t=0}^{\infty} \gamma^t r(s_t)]$.

For any policy π , the corresponding state-action value function is $Q^\pi(s, a) = \mathbb{E}[\sum_{k=0}^{\infty} \gamma^k r_{t+k} | S_t = s, A_t = a, \pi]$.

Action Mask Mechanism

The action mask mechanism significantly reduces the complexity of exploration for some complicated cases, such as real-time strategy games or job scheduling. In standard value iteration methods, the agent will make decisions by maximizing $Q(s, a)$ among the full action space \mathcal{A} :

$$a = \arg \max_{a \in \mathcal{A}} Q(s, a).$$

However, if $|\mathcal{A}|$ is too large and contains numerous redundant actions, the difficulty of the exploration is greatly increased. For this reason, we define the invalid action space $\mathcal{A}^{\text{invalid}} \subseteq \mathcal{A}$ to reduce the difficulty of the exploration, which can be state-dependent or time-varying. Given $\mathcal{A}^{\text{invalid}}$, value-based RL methods will replace the full action space by the valid action space $\mathcal{A} \setminus \mathcal{A}^{\text{invalid}}$:

$$a = \arg \max_{a \in \mathcal{A} \setminus \mathcal{A}^{\text{invalid}}} Q(s, a). \tag{1}$$

As for policy gradient methods, the action mask mechanism will forcibly set the sampling probability of the redundant actions to zero and re-normalize:

$$\pi_{\text{mask}}(a|s) = \begin{cases} \frac{\pi(a|s)}{1 - \sum_{a \in \mathcal{A}^{\text{invalid}}} \pi(a|s)} & , a \notin \mathcal{A}^{\text{invalid}} \\ 0 & , a \in \mathcal{A}^{\text{invalid}}. \end{cases} \tag{2}$$

Method

We start with the theoretical analysis of what action can be eliminated in policy optimization. Based on this analysis, we propose a redundant action classifying mechanism, which constructs the similarity factor by estimating the distance between the state distributions. Further, we introduce the modified inverse model to avoid extensive computation. Finally, we show the overall framework of our algorithm in Figure 1 and Algorithm 1.

Identify Redundant Actions Based on the Similarity Factor

What actions can be eliminated in policy optimization? At first sight, actions transferring from the same initial state s_t to the same next state s_{t+1} should be classified into similar actions since they have the same effect on the state transition. We formulate this intuition by introducing the similarity factor as follows:

Definition 1 (Similarity Factor). We define the similarity factor $M(s_t, a_i, a_j)$ by KL-divergence following state s_t and any two actions a_i, a_j :

$$M(s_t, a_i, a_j) = D_{\text{KL}}(P(s_{t+1}|s_t, a_i) \| P(s_{t+1}|s_t, a_j)). \quad (3)$$

We can use the similarity factor to identify whether the actions are redundant or irrelevant based on the definition. Specifically, if the value of $M(s_t, a_i, a_j)$ is small, the corresponding two actions have similar effects on the state s_t and should be aligned into the same cluster. Therefore, we have the following formulation:

Definition 2 (State-dependent Action Clusters). We define the state-dependent action cluster $A_{s,\epsilon}^m$ over state s based on the similarity factor:

$$\mathcal{A} = \bigcup_{m \in \{0,1,\dots,l-1\}} A_{s,\epsilon}^m, \quad (4)$$

$$\text{s.t. } \forall a_i, a_j \in A_{s,\epsilon}^m, \quad M(s, a_i, a_j) < \epsilon \quad (5)$$

where $A_{s,\epsilon}^m$ is the sub-action space and $A_{s,\epsilon}^i \cap A_{s,\epsilon}^j = \emptyset$ if $i \neq j$. In practice, we choose arbitrarily one action $a_{A_m} \in A_{s,\epsilon}^m$ (e.g. the first action) to represent the sub-action space $A_{s,\epsilon}^m$ and mask out all the redundant actions. Then, based on the selected action a_{A_m} , we can reduce the original action space \mathcal{A} in the state s into the minimal action space $\mathcal{A}_s^\epsilon = \{a_{A_0}, a_{A_1}, \dots, a_{A_{l-1}}\}$. Next, similar to Equation 1 and Equation 2, the policy gradient methods and value-based methods can be optimized by simply replacing the full action space with the minimal action space \mathcal{A}_s^ϵ .

Actions in the same clusters have similar effects on the state. Intuitively, eliminating the redundant actions will have little effect on the policy’s performance. Based on the above motivation, we have the following theoretical analysis:

Lemma 1. For any policy π and $A_{s,\epsilon}^m$, we can choose arbitrarily one action $a_{A_m} \in A_{s,\epsilon}^m$ to represent the sub-action space. We denote

$$\pi_{A_{s,\epsilon}^m}(a | s) = \begin{cases} \sum_{a' \in A_{s,\epsilon}^m} \pi(a' | s) & , a = a_{A_m} \\ 0 & , a \in A_{s,\epsilon}^m \setminus \{a_{A_m}\} \\ \pi(a | s) & , \text{o.w} \end{cases} \quad (6)$$

then

$$\|V^\pi - V^{\pi_{A_{s,\epsilon}^m}}\|_\infty \leq \frac{\gamma r_{\max}}{(1-\gamma)^2} \sqrt{2\epsilon} \quad (7)$$

Lemma 1 shows that eliminating the redundant actions based on the similarity factor has little effect on the policy’s performance. Please refer to Appendix for the proof.

Estimate Similarity Factor Using Modified Inverse Dynamic Model

Although we can identify the redundant actions based on the similarity factor, accurately calculating $M(s_t, a_i, a_j)$ is extremely hard since it requires the distance on the state distribution. Therefore, the key issue is reducing the computational complexity of $M(s_t, a_i, a_j)$. We derive a form equivalent to Definition 1, shown in Lemma 2 to solve this issue. Compared with Definition 1, Lemma 2 transfers the computation on the state space to the action space, making calculating $M(s_t, a_i, a_j)$ feasible.

Lemma 2. The similarity factor M can be divided by two terms with the same form N , which we refer to as N-value network:

$$M(s_t, a_i, a_j) = N(s_t, a_i, a_i) - N(s_t, a_i, a_j)$$

where

$$N(s_t, a_i, a_j) = \mathbb{E}_{s_{t+1} \sim P(\cdot | s_t, a_i)} \log \left[\frac{P^\pi(a_j | s_t, s_{t+1})}{\pi(a_j | s_t)} \right].$$

For $i = j$, then we have the former term

$$N(s_t, a_i, a_i) = \mathbb{E}_{s_{t+1} \sim P(\cdot | s_t, a_i)} \log \left[\frac{P^\pi(a_i | s_t, s_{t+1})}{\pi(a_i | s_t)} \right].$$

In Lemma 2, we introduce the inverse dynamic model $P^\pi(a_t | s_t, s_{t+1})$ to calculate $M(s_t, a_i, a_j)$, which greatly reduce the computational complexity of our method. Please refer to Appendix for the proof.

We note that $P^\pi(a_t | s_t, s_{t+1})$ is different from $P(a_t | s_t, s_{t+1})$, which ignores the π and do not take the policy distribution over s_t into account. The conflict induced by the $P(a_t | s_t, s_{t+1})$ will be severe when training on the data collected by varying policies.

When the policy π is fixed, $P^\pi(a_t | s_t, s_{t+1})$ is equivalent to $P(a_t | s_t, s_{t+1})$. However, in more complex environments, where the agent needs to explore a vast state space, a fixed policy may not lead to effective exploration. To solve this issue, we have to adopt multiple exploration policies (e.g., curiosity reward (Pathak et al. 2017)). With the collected data coming from gradually varying policies, the modified inverse dynamics model lead to better performance in Experiment Section.

Therefore, we construct the modified inverse dynamic model $P_\psi^{\text{inv}}(a_t | s_t, s_{t+1}, \pi(\cdot | s_t))$ with parameter ψ to estimate $P^\pi(a_t | s_t, s_{t+1})$. In Experiment Section, we show that the modified inverse dynamic model can converge faster and achieve higher accuracy compared with $P(a_t | s_t, s_{t+1})$ with the collected data coming from gradually varying policies.

Practical Implementation

In this subsection, we show how to train the modified inverse dynamic model P_ψ^{inv} and N-value network N_θ parameterized with θ . First, we train the inverse dynamic model with the following loss:

$$L(\psi) = H_{(s_t, a_t, s_{t+1}, \pi(\cdot | s_t)) \sim \mathcal{D}} [P_\psi^{\text{inv}}(\cdot | s_t, s_{t+1}, \pi(\cdot | s_t)), P_{a_t}^{\text{one-hot}}(\cdot)], \quad (8)$$

where $H(\cdot, \cdot)$ is the cross entropy loss and $|P_{a_t}^{\text{one-hot}}(\cdot)| = |\mathcal{A}|$. $P_{a_t}^{\text{one-hot}}(\cdot)$ denotes one-hot distribution, where the index value corresponding to a_t is 1, otherwise is 0.

As for the N-value network N_θ , we adopt the state-action pair (s, a_t) as inputs and output a $|\mathcal{A}|$ -dimensional vector. Actually, there are some alternative plans to construct an N-value network. For example, we can input only state s and output a similarity factor matrix. However, with the increase of $|\mathcal{A}|$, the output dimension is positively correlated with

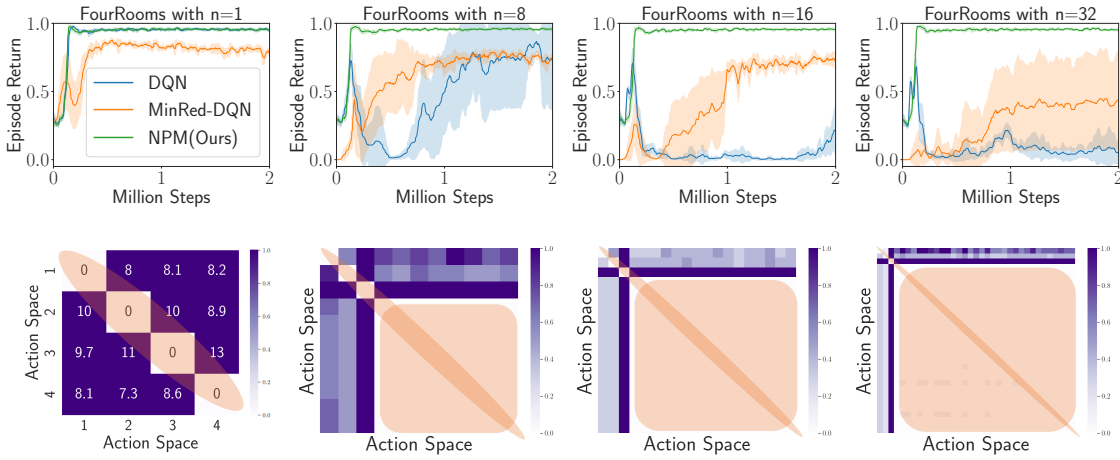


Figure 2: Upper: The experimental results on FourRooms with synthetic action redundancy $n \in \{1, 8, 16, 32\}$ with five random seeds. Below: Visualization of similarity factor matrix corresponding to each task.

$|\mathcal{A}|^2$, leading the method difficult to guarantee convergence. Therefore, based on the Lemma 2, we train the N-value network with the following loss:

$$L(\theta) = \mathbb{E}_{(s_t, a_t, s_{t+1}, \pi(\cdot|s_t)) \sim \mathcal{D}} [N_\theta(s_t, a_t, \cdot) - \log \frac{P_\psi^{\text{inv}}(\cdot|s_t, s_{t+1}, \pi(\cdot|s_t))}{\pi(\cdot|s_t)}]_2 \quad (9)$$

where $P_\psi^{\text{inv}}(\cdot|s_t, s_{t+1}, \pi(\cdot|s_t))$ is the trained inverse dynamic model. In practice, the modified inverse dynamic model P_ψ^{inv} and N-value network N_θ are trained iteratively. P_ψ^{inv} updates faster than N-value network to guarantee the convergence. Based on the trained N-value network, we calculate the similarity factor between any two actions a_i, a_j :

$$M(s_t, a_i, a_j) = N_\theta(s_t, a_i, a_j) - N_\theta(s_t, a_i, a_j). \quad (10)$$

In the practical implementation, we iterate over each action a_i in the discrete action set to obtain the similarity factor matrix $M(s_t, \cdot, \cdot)$ ($|\mathcal{A}| \times |\mathcal{A}|$), where every element represents the KL-divergence between any two actions. Then, we can utilize $M(s_t, \cdot, \cdot)$ to construct state-dependent action clusters and eliminate the redundant actions as illustrated in Lemma 1.

We show the overall framework in Algorithm 1 in Appendix and Figure 1. Our algorithm consists of two parts. The former is to train the similarity factor model. The latter is eliminating redundant actions during policy optimization. We note that in complex tasks we instead adopt the curiosity reward (e.g., $r(s, a) = N(s, a)^{-\frac{1}{2}}$) (Pathak et al. 2017) to update the initial random policy while training the similarity factor model.

Experiments

In this section, we aim to evaluate the effectiveness of the redundant action filtering mechanism and answer the following questions: (1) How does NPM perform compared with other methods in various action redundancy tasks? (2) Why




Figure 3: Left: Four room task. Right: Maze task.

and When is NPM effective? (3) Whether the learned mask can be transferred across multi-tasks?

Experimental Setting

We evaluate our method on three action redundancy tasks: synthetic action redundancy, combined action redundancy, and state-dependent action redundancy including Minigrid environments (Chevalier-Boisvert, Willems, and Pal 2018; Chevalier-Boisvert et al. 2018) and Atari benchmark. In all experiments, we utilize image as the sole input and provide evidence for NPM’s scalability to high-dimensional pixel-based observations.

Synthetic Action Redundancy. We consider a controlled synthetic four-room task consistent with (Baram, Tennenholtz, and Mannor 2021), shown in Figure 3. In this task, agent (red circle) must achieve the goal (star) to earn the reward within 100 steps. Meanwhile, the wind in the vertical direction will interfere agent with probability p . We construct a synthetic redundant action space $\mathcal{A} =$

$\{\text{Top, Bottom, Left, Right, \dots, Right}\}$, for which the action Right was repeated n times. To reach the goal, the agent must counteract redundancy in action space.

Combined Action Redundancy. We consider the continuous-state maze task consistent with (Chandak et al. 2019), shown in Figure 3. Unlike the Synthetic Action Redundancy, the agent has m equally spaced actuators (each actuator moves the agent in the direction the actuator is pointing towards), and it can choose whether each actuator




Figure 4: The experimental results on Maze with combined action redundancy with actions = {16, 64, 256}. The experiments are conducted under five random seeds.

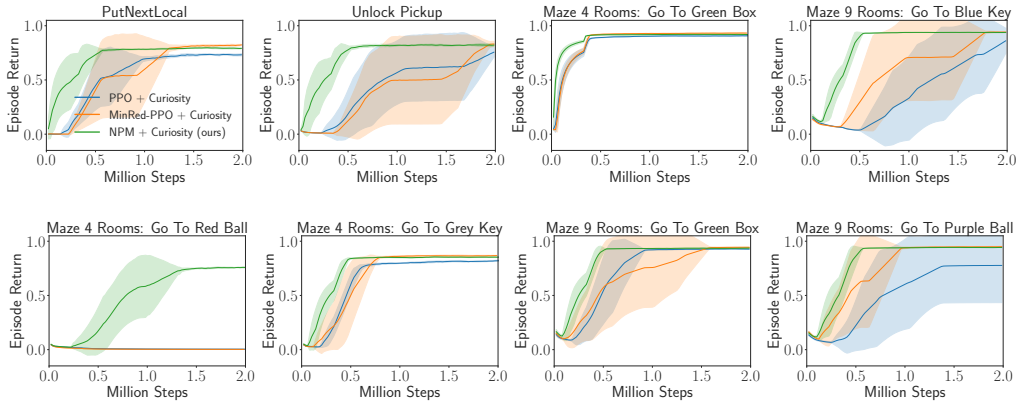


Figure 5: Above: The experimental results in the state-dependent action redundancy tasks under five random seeds. Below: Transfer results of the learned mask in the Maze 4 Rooms and Maze 9 Rooms.

should be on or off. Therefore, the size of the action set is exponential in the number of actuators, that is, $|\mathcal{A}| = 2^m$. The net outcome of an action is the vectorial summation of the displacements associated with the selected actuators. To earn the reward, the agent must achieve the goal (star) within 150 steps.

State-dependent Action Redundancy. We consider the Minigrid environments (Chevalier-Boisvert, Willems, and Pal 2018; Chevalier-Boisvert et al. 2018) as the state-dependent action redundancy tasks. As shown in Figure 12 in Appendix, there are common identifiers among these tasks: agent, box, key, ball, and door. We need to control the agent to complete different tasks. For example, PutNextLocal is putting an object next to another object inside a single room with no doors. UnlockPickup is picking up a box that is placed in another room behind a locked door. The Maze 4 and 9 Rooms control the agent to go to an object, and the object may be in another room. The action space is $\mathcal{A} = \{\text{Turn Left, Turn Right, Move Forward, Pick Up, Drop, Toggle, Noop}\}$. Note that there are state-dependent redundancies in these tasks. For instance, the agent can pick up an object only when this object is in front of the agent. Otherwise, the pick-up action will not work, and the agent will stay still. Therefore, we need to eliminate the action redundancy to complete these tasks efficiently.

Baselines: We compare NPM against the state-of-the-art baseline, MinRed (Baram, Tennenholtz, and Mannor 2021), which maximizes the next states’ entropy to minimize ac-

tion redundancy. In this paper, we compare our method with two modified algorithms based on MinRed, named MinRed-DQN and MinRed-PPO. In addition, we compare our method with DQN (Mnih et al. 2015), PPO (Schulman et al. 2017), and PPO with curiosity reward (Pathak et al. 2017).

Experimental Results

Results of Synthetic Action Redundancy. We compare NPM with MinRed-DQN and DQN in the Four-Rooms task with synthetic action redundancy $n = \{1, 8, 16, 32\}$, which is shown in Figure 2. The experimental results show that with the increase of n , DQN converges more slowly or collapses. The performance of MinRed-DQN will drop if n is large. Differently, NPM converges quickly and achieves better performance.

Further, we visualize the similarity factor matrix $M(s_t, \cdot, \cdot)$, where every element represents KL-divergence between two actions defined in Definition 1. The experimental results in Figure 2 show that the KL value between redundant actions is 0. This demonstrates that the learned similarity factor model successfully captures the action structure, and we can use it to mask the redundant actions.

Results of Combined Action Redundancy. In the maze, we compare NPM with MinRed-DQN and DQN with $|\mathcal{A}| = \{2^4, 2^6, 2^8\}$. The experimental results in Figure 4 show that NPM performs better than MinRed-DQN and DQN. We further visualize the similarity factor. Specifically, we first visu-




Figure 6: Left: 2-D representations for the displacements in the Cartesian coordinate caused by each action. Right: t-SNE projection based on similarity factor. We find that the actions with similar effects gather together based on the similarity factor. The relative position of points of different colors in the Cartesian coordinate is consistent with the t-SNE projection result.

alize the displacement in the Cartesian coordinate caused by each action. We mark actions with a similar effect as similar colors (There are some overlapping points). Then, we calculate the t-SNE projection based on the similarity factor shown in Figure 6. The experimental results show two interesting phenomena: The actions with similar effects gather together based on the similarity factor. Second, the relative position of points of different colors in the Cartesian coordinate is consistent with the t-SNE projection result. This demonstrates that the relationship between different actions is also revealed in the learned representations.

Results of State-dependent Action Redundancy. We compare NPM with PPO and Minred+PPO with the curiosity reward. The experimental result in Figure 5 shows NPM achieves superior performance than baselines. In addition, we test the transfer ability of the learned mask. Specifically, we train the action mask only in the “Maze 4 Rooms: Go To Green Box” and “Maze 9 Rooms: Go to Blue Key” in the first phase. Next, we respectively train the RL policy with the corresponding action masks in the second phase in the “Maze 4 Rooms: Go To Red Ball, Grey Key” and “Maze 9 Rooms: Go to Blue Key, Purple Ball”. The experimental results in Figure 5 show that the learned mask of NPM can transfer across multi-tasks.

Further, we visualize the learned similarity factor. The results in Figure 13 and Figure 14 in Appendix show that the agent has different action clusters in various state positions, which can be successfully captured based on the learned similarity factor.

Ablation Study

How to select ϵ is an important issue in our algorithm. As shown by Lemma 1, to limit the performance error to an acceptable level, we need to adjust the value of ϵ as the reward range $[0, r_{\max}]$ in the environment increases. In typical scenarios where the range of reward is not particularly large, it is appropriate to choose the range of hyperparameters ϵ between 0.05 and 0.5 in practical implementation. We conduct the ablation studies for $\epsilon \in \{0.01, 0.05, 0.1, 0.5, 1, 5\}$. The

results in Figure 7 in Appendix show that NPM performs robust when $\epsilon \in (0.05, 0.5)$.

Modified inverse dynamic model: We compare the modified inverse dynamic model with the original model in the combined action redundancy domain (maze task). We select various action numbers n in the maze task. The experimental results in Figure 8 in Appendix show that the modified inverse dynamic model (named modified) learns faster and achieves higher accuracy than the original dynamic model (named original). The result is consistent with the statement in section .

Curiosity Reward: We conduct ablation studies for the curiosity reward used in the state-dependent action redundancy tasks. The results in Figure 9 in Appendix show that the curiosity reward is essential in complex tasks. Further, we visualize the visitation locations in the Unlock-Pickup task. The result in Figure 10 in Appendix shows that agent driven by the curiosity reward succeeds in adequate exploration.

Extend to Atari Tasks: We conduct experiments on another state-dependent action redundancy tasks (e.g., Atari 2600). The experimental results in Figure 15 and Figure 16 in Appendix show our method can effectively remove the redundant actions and improve the current baselines.

Conclusion

In this paper, we construct a framework to reveal the underlying structure of action space without prior knowledge. Our theoretical analysis has yielded a crucial finding: the policy derived from our framework has minimal impact on the performance of the original policy. Building upon these insights, we propose a simple yet efficient algorithm, known as NPM. In practice, our approach consistently outperforms existing methods across a wide range of tasks. In future, we will explore the extension of NPM into continuous action spaces and its potential applications in robotics.

Acknowledgments

This work is supported by National Natural Science Foundation of China under Grant No. 62192751, in part by Key R&D Project of China under Grant No. 2017YFC0704100, the 111 International Collaboration Program of China under Grant No. BP2018006, and in part by the BNRist Program under Grant No. BNR2019TD01009, the National Innovation Center of High Speed Train R&D project (CX/KJ-2020-0006), in part by the InnoHK Initiative, The Government of HKSAR; and in part by the Laboratory for AI-Powered Financial Technologies.

References

- Abel, D.; Hershkowitz, D.; Barth-Maron, G.; Brawner, S.; O’Farrell, K.; MacGlashan, J.; and Tellex, S. 2015. Goal-based action priors. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 25, 306–314.
- Baram, N.; Tennenholtz, G.; and Mannor, S. 2021. Action redundancy in reinforcement learning. In *Uncertainty in Artificial Intelligence*, 376–385. PMLR.
- Bellemare, M. G.; Naddaf, Y.; Veness, J.; and Bowling, M. 2013. The arcade learning environment: An evaluation platform for general agents. *Journal of Artificial Intelligence Research*, 47: 253–279.
- Berner, C.; Brockman, G.; Chan, B.; Cheung, V.; Debiak, P.; Dennison, C.; Farhi, D.; Fischer, Q.; Hashme, S.; Hesse, C.; et al. 2019. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*.
- Biza, O.; Wang, D.; Platt, R.; van de Meent, J.-W.; and Wong, L. L. 2021. Action priors for large action spaces in robotics. *arXiv preprint arXiv:2101.04178*.
- Chandak, Y.; Theodorou, G.; Kostas, J.; Jordan, S.; and Thomas, P. 2019. Learning action representations for reinforcement learning. In *International conference on machine learning*, 941–950. PMLR.
- Chevalier-Boisvert, M.; Bahdanau, D.; Lahlou, S.; Willems, L.; Saharia, C.; Nguyen, T. H.; and Bengio, Y. 2018. Babyai: A platform to study the sample efficiency of grounded language learning. *arXiv preprint arXiv:1810.08272*.
- Chevalier-Boisvert, M.; Willems, L.; and Pal, S. 2018. Minimalistic Gridworld Environment for Gymnasium.
- Dadashi, R.; Hussenot, L.; Vincent, D.; Girgin, S.; Raichuk, A.; Geist, M.; and Pietquin, O. 2021. Continuous Control with Action Quantization from Demonstrations. In *Deep RL Workshop NeurIPS 2021*.
- Dulac-Arnold, G.; Evans, R.; van Hasselt, H.; Sunehag, P.; Lillicrap, T.; Hunt, J.; Mann, T.; Weber, T.; Degris, T.; and Coppin, B. 2015. Deep reinforcement learning in large discrete action spaces. *arXiv preprint arXiv:1512.07679*.
- Fernández, F.; and Veloso, M. 2006. Probabilistic policy reuse in a reinforcement learning agent. In *Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, 720–727.
- Glavic, M.; Fonteneau, R.; and Ernst, D. 2017. Reinforcement learning for electric power system decision and control: Past considerations and perspectives. *IFAC-PapersOnLine*, 50(1): 6918–6927.
- Huang, S.; Ontañón, S.; Bamford, C.; and Grela, L. 2021. Gym- μ RTS: Toward Affordable Full Game Real-time Strategy Games Research with Deep Reinforcement Learning. In *2021 IEEE Conference on Games (CoG)*, 1–8. IEEE.
- Jiang, Z.; Xu, D.; and Liang, J. 2017. A deep reinforcement learning framework for the financial portfolio management problem. *arXiv preprint arXiv:1706.10059*.
- Ma, X.; Yang, Y.; Hu, H.; Liu, Q.; Yang, J.; Zhang, C.; Zhao, Q.; and Liang, B. 2021. Offline reinforcement learning with value-based episodic memory. *arXiv preprint arXiv:2110.09796*.
- Metz, L.; Ibarz, J.; Jaitly, N.; and Davidson, J. 2017. Discrete sequential prediction of continuous actions for deep rl. *arXiv preprint arXiv:1705.05035*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Riedmiller, M.; Fidjeland, A. K.; Ostrovski, G.; et al. 2015. Human-level control through deep reinforcement learning. *nature*, 518(7540): 529–533.
- Panangaden, P.; Rezaei-Shoshtari, S.; Zhao, R.; Meger, D.; and Precup, D. 2023. Policy Gradient Methods in the Presence of Symmetries and State Abstractions. *arXiv preprint arXiv:2305.05666*.
- Pathak, D.; Agrawal, P.; Efros, A. A.; and Darrell, T. 2017. Curiosity-driven exploration by self-supervised prediction. In *International conference on machine learning*, 2778–2787. PMLR.
- Ravindran, B. 2004. *An algebraic approach to abstraction in reinforcement learning*. University of Massachusetts Amherst.
- Rosman, B.; and Ramamoorthy, S. 2012. What good are actions? Accelerating learning using learned action priors. In *2012 IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL)*, 1–6. IEEE.
- Rosman, B.; and Ramamoorthy, S. 2015. Action priors for learning domain invariances. *IEEE Transactions on Autonomous Mental Development*, 7(2): 107–118.
- Sallans, B.; and Hinton, G. E. 2004. Reinforcement learning with factored states and actions. *The Journal of Machine Learning Research*, 5: 1063–1088.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Seyde, T.; Werner, P.; Schwarting, W.; Gilitschenski, I.; Riedmiller, M.; Rus, D.; and Wulfmeier, M. 2022. Solving Continuous Control via Q-learning. In *The Eleventh International Conference on Learning Representations*.
- Sharma, S.; Suresh, A.; Ramesh, R.; and Ravindran, B. 2017. Learning to factor policies and action-value functions: Factored action space representations for deep reinforcement learning. *arXiv preprint arXiv:1705.07269*.
- Sherstov, A. A.; and Stone, P. 2005. Improving action selection in MDP’s via knowledge transfer. In *AAAI*, volume 5, 1024–1029.

- Sidney, K. D.; Craig, S. D.; Gholson, B.; Franklin, S.; Picard, R.; and Graesser, A. C. 2005. Integrating affect sensors in an intelligent tutoring system. In *Affective Interactions: The Computer in the Affective Loop Workshop at*, 7–13.
- Silver, D.; Huang, A.; Maddison, C. J.; Guez, A.; Sifre, L.; Van Den Driessche, G.; Schrittwieser, J.; Antonoglou, I.; Panneershelvam, V.; Lanctot, M.; et al. 2016. Mastering the game of Go with deep neural networks and tree search. *nature*, 529(7587): 484–489.
- Tavakoli, A.; Fatemi, M.; and Kormushev, P. 2021. Learning to Represent Action Values as a Hypergraph on the Action Vertices. In *International Conference on Learning Representations (ICLR 2021)*.
- Tavakoli, A.; Pardo, F.; and Kormushev, P. 2018. Action branching architectures for deep reinforcement learning. In *Proceedings of the aai conference on artificial intelligence*, volume 32.
- Taylor, J.; Precup, D.; and Panagaden, P. 2008. Bounding performance loss in approximate MDP homomorphisms. *Advances in Neural Information Processing Systems*, 21.
- Tennenholtz, G.; and Mannor, S. 2019. The Natural Language of Actions. *ArXiv*, abs/1902.01119.
- Thomas, V.; Pondard, J.; Bengio, E.; Sarfati, M.; Beaudoin, P.; Meurs, M.-J.; Pineau, J.; Precup, D.; and Bengio, Y. 2017. Independently controllable features. *arXiv preprint arXiv:1708.01289*.
- Vinyals, O.; Ewalds, T.; Bartunov, S.; Georgiev, P.; Vezhnevets, A. S.; Yeo, M.; Makhzani, A.; Küttler, H.; Agapiou, J.; Schrittwieser, J.; et al. 2017. Starcraft ii: A new challenge for reinforcement learning. *arXiv preprint arXiv:1708.04782*.
- Yan, K.; Schwing, A.; and Wang, Y.-X. 2022. CEIP: Combining Explicit and Implicit Priors for Reinforcement Learning with Demonstrations. *Advances in Neural Information Processing Systems*, 35: 7614–7627.
- Yang, Y.; Ma, X.; Li, C.; Zheng, Z.; Zhang, Q.; Huang, G.; Yang, J.; and Zhao, Q. 2021. Believe what you see: Implicit constraint approach for offline multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 10299–10312.
- Ye, D.; Liu, Z.; Sun, M.; Shi, B.; Zhao, P.; Wu, H.; Yu, H.; Yang, S.; Wu, X.; Guo, Q.; et al. 2020. Mastering complex control in moba games with deep reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 6672–6679.
- Zahavy, T.; Haroush, M.; Merlis, N.; Mankowitz, D. J.; and Mannor, S. 2018a. Learn What Not to Learn: Action Elimination with Deep Reinforcement Learning. In Bengio, S.; Wallach, H.; Larochelle, H.; Grauman, K.; Cesa-Bianchi, N.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc.
- Zahavy, T.; Haroush, M.; Merlis, N.; Mankowitz, D. J.; and Mannor, S. 2018b. Learn what not to learn: Action elimination with deep reinforcement learning. *Advances in neural information processing systems*, 31.