Mind the Model, Not the Agent: The Primacy Bias in Model-Based RL

Zhongjian Qiao^a, Jiafei Lyu^a and Xiu Li^{a,*}

^aTsinghua Shenzhen International Graduate School, Tsinghua University

Abstract. The primacy bias in model-free reinforcement learning (MFRL), which refers to the agent's tendency to overfit early data and lose the ability to learn from new data, can significantly decrease the performance of MFRL algorithms. Previous studies have shown that employing simple techniques, such as resetting the agent's parameters, can substantially alleviate the primacy bias in MFRL. However, the primacy bias in model-based reinforcement learning (MBRL) remains unexplored. In this work, we focus on investigating the primacy bias in MBRL. We begin by observing that resetting the agent's parameters harms its performance in the context of MBRL. We further find that the primacy bias in MBRL is more closely related to the primacy bias of the world model instead of the primacy bias of the agent. Based on this finding, we propose world *model resetting*, a simple yet effective technique to alleviate the primacy bias in MBRL. We apply our method to two different MBRL algorithms, MBPO and DreamerV2. We validate the effectiveness of our method on multiple continuous control tasks on MuJoCo and DeepMind Control Suite, as well as discrete control tasks on Atari 100k benchmark. The experimental results show that world model resetting can significantly alleviate the primacy bias in the modelbased setting and improve the algorithm's performance. We also give a guide on how to perform world model resetting effectively.

1 Introduction

Deep Reinforcement Learning (DRL) has shown great potential in numerous fields such as robot control [3] and autonomous driving [22]. However, training an agent with good performance in DRL requires a large amount of data. Gathering samples from real environment can be costly, making it crucial to improve the algorithm's sample efficiency and train a better-performing agent using limited data. A direct approach to improving sample efficiency is to perform multiple training updates per environment step [5, 16], which means using a high UTD (Update-To-Data) ratio. This is analogous to updating multiple epochs for a dataset in supervised learning. However, unlike supervised learning, in DRL, the agent continuously interacts with the environment and learns from the replay buffer, which contains a dynamic and evolving data collection. As the policy is updated, the data distribution in the replay buffer also changes. In that case, using a high UTD ratio for updates can magnify the primacy bias [31], which refers to the agent's tendency to overfit early data distribution and lose the ability to learn from new data distribution effectively. Previous studies [5, 27, 31] have shown that such primacy



Figure 1. The basic flowchart of MBRL algorithms. For each data sampled from the environment, the world model is updated m times. The agent then is trained on data generated by the world model n times. In practice, n is a positive integer, and m is a fraction, which means that the world model is updated every multiple environment steps.

bias due to a high UTD ratio can seriously damage the agent's performance in model-free setting. However, we notice that in model-based reinforcement learning algorithms (MBRL), where a world model is utilized, it is common to use a high UTD ratio for updates to improve sample efficiency without causing significant harm to the agent's performance. For example, MBPO [17] employs a high UTD ratio of at least 20 to update the agent, which is enough for MFRL algorithms like SAC [14] and TD3 [10] to fail [27]. However, MBPO has demonstrated superior performance on many continuous control tasks. This phenomenon raises the following questions: *Does the primacy bias also exist in model-based setting? Does the primacy bias in MBRL differ from its form in model-free scenario if it does exist?*

In this paper, we focus on exploring the primacy bias in MBRL, which has been neglected in prior studies. We start by assuming that there are two forms of the primacy bias: the primacy bias of the agent and the primacy bias of the world model. Next, we apply the parameter resetting technique, a simple yet effective approach to reducing the primacy bias proposed by [31], in MBPO. Surprisingly, instead of improving the agent's performance, we observe a decline in performance compared to the original algorithm. This suggests that parameter resetting seems not to work in model-based setting. To further investigate this phenomenon, we design and conduct extensive experiments. Our findings reveal that the primacy bias still exists in model-based setting. However, it is more closely related to the primacy bias of the world model, which differs from the situation in model-free setting. Building upon this finding, we propose world model resetting, which is a modification to the parameter resetting method. This slight modification works in model-based setting, while original parameter resetting does not. To validate the effectiveness of world model resetting, we apply it to two MBRL algorithms:

^{*} Corresponding Author. Email: li.xiu@sz.tsinghua.edu.cn.



Figure 2. Left: Learning curves for Hopper-v2 task of MuJoCo for SAC and SAC with parameter resetting (abbreviated as SAC with reset). **Right:** Learning curves for Hopper-v2 task of MuJoCo for MBPO and MBPO with parameter resetting (abbreviated as MBPO with reset). The solid line represents the average return, the shaded area represents the 95% confidence interval. The results are evaluated by 5 runs.

MBPO and DreamerV2 [15]. We conduct extensive experiments on multiple continuous control tasks on MuJoCo [40] and DeepMind Control Suite (DMC) [38], as well as discrete control tasks on Atari 100k benchmark [20]. The results show that *world model resetting* can reduce the primacy bias and improve the performance of MBRL algorithms in both continuous and discrete domains. We also compare *world model resetting* with AutoMBPO [25], a parameter tuning method designed specifically for MBPO, to demonstrate the superiority of our method under high UTD ratio conditions. In the final part of our experiments, we examine the key factors that influence the effectiveness of *world model resetting*. This provides a guideline on effectively utilizing our method for better performance.

2 The Primacy Bias in MBRL

The main goal of our work is to explore the primacy bias in MBRL. In model-free RL algorithms, the primacy bias typically refers to the agent's tendency to overfit early training data. However, in model-based RL algorithms, we point out that there are two forms of the primacy bias:

The Primacy Bias of the agent: The Q-network and the policy network overfit the early data in the replay buffer and cannot effectively use the data in the later stage of training, resulting in learning a suboptimal policy.

The Primacy Bias of the world model: The world model overfits the data distribution under the early policy and cannot effectively represent the environment dynamics under the evolving policy, failing the algorithm.

Unlike model-free RL algorithms, which only suffer from the first form of the primacy bias, the above two forms of primacy bias can appear in MBRL algorithms at the same time, and we assume that either of them may cause the failure of the entire algorithm. In Figure 1, we present the sources of these two forms of primacy bias. We refer to m as model UTD ratio and n as agent UTD ratio. A high model UTD ratio may lead to overfitting of the world model, which magnifies the primacy bias of the world model. Similarly, a high agent UTD ratio can potentially lead to overfitting of the agent, which stretches the primacy bias of the agent.

In the following part of this section, we first apply the parameter resetting technique to MBRL but observe that it can't improve the



Figure 3. Left: Learning curves for Hopper-v2 of MuJoCo for MBPO with different model UTD ratio. **Right:** Learning curves for Hopper-v2 of MuJoCo for MBPO with different agent UTD ratio. The solid line represents the average return, the shaded area represents the 95% confidence interval. The results are evaluated by 5 runs.

algorithm's performance. Instead, it results in performance degradation. Based on this observation, we empirically investigate the impact of the two forms of the primacy bias on the performance of MBRL algorithms. As a conclusion, we identify the main form of the primacy bias in model-based setting.

2.1 Parameter resetting doesn't improve the performance

Resetting the agent's parameters periodically has been shown a simple yet powerful technique to eliminate the primacy bias in modelfree scenario, thus improving the algorithm's performance, especially in high agent UTD ratio scenario. But in MBRL setting, can parameter resetting still be effective?

We apply parameter resetting to model-free RL algorithm SAC and model-based RL algorithm MBPO to answer that question. We experiment on Hopper-v2 task on MuJoCo. To make the effect of parameter resetting as significant as possible, we refer to the experiment settings suggested in [31]. For SAC, we set the agent UTD ratio to 32, set the reset interval to 1×10^5 environment steps, and reset the entire Q-network (including target Q-network) and policy network. For MBPO, the optimization algorithm used is also SAC; we set the reset interval to 2×10^4 environment steps, set the agent UTD ratio to 32 (the default agent UTD ratio is 20), and the other hyper-parameters are consistently with the default setting. We train SAC for a total of 400k environment steps and MBPO for 130k environment steps. Figure 2 shows our experimental results. We can observe that SAC with a high agent UTD ratio has poor performance, but the algorithm performance significantly improves with parameter resetting, consistent with previous studies. However, for MBPO, a high agent UTD ratio does not lead to worse performance. Applying parameter resetting actually harms the performance instead of improving it. To investigate this phenomenon, we conduct further experiments as follows.

2.2 High model UTD ratio hurts the performance

The phenomenon observed in Figure 2 seems quite perplexing. Why is parameter resetting effective in model-free scenarios but not in model-based setting? Does this imply the absence of the primacy bias in MBRL?

To answer these questions, we conduct further experiments. To investigate the existence of the two forms of primacy bias and their impact on the performance of MBRL algorithms, we conduct experiments still using MBPO algorithm on Hopper-v2 task of MuJoCo.



Figure 4. Left: Performance comparison between MBPO and MBPO with a pre-initialized buffer on Hopper-v2 task. Middle: Model mean squared error comparison between MBPO and MBPO with a pre-initialized buffer. Right: Value-aware model error comparison between MBPO and MBPO with a pre-initialized buffer. All the results are evaluated by 5 runs.

First, to investigate the effect of the primacy bias of the world model on algorithm's performance, we keep the other hyperparameters at their default settings and only modify the model UTD ratio; we observe how this change affects the algorithm's performance. The default model UTD ratio in the original MBPO paper is $\frac{1}{250}$, which means the world model is updated every 250 environment steps. We denote the default model UTD ratio as dUTD for convenience. In addition to dUTD, we set the model UTD ratio to $2 \times dUTD$, $5 \times dUTD$, and $10 \times \text{dUTD}$ for four experimental groups. We train MBPO and plot the learning curves in Figure 3 Left. We can observe that when the model UTD ratio doubles, algorithm's performance barely differs. However, when model UTD ratio becomes 5 times and 10 times, there is a significant decrease in the performance of MBPO. This indicates the existence of the primacy bias of the world model, and it does not require a high model UTD ratio (even at 10 times the default model UTD ratio, which is only $\frac{1}{25}$) for the primacy bias of the world model to affect algorithm's performance significantly. Similarly, to investigate whether the primacy bias of the agent affects algorithm's performance, we change the agent UTD ratio while keeping other hyperparameters default and conduct sets of experiments. We vary the agent UTD ratio to 1, 5, 10, 20, 40, 80, conducting a total of six groups of experiments. We plot the learning curves in Figure 3 Right. We observe that the performance of MBPO improves with the increase of the agent UTD ratio, until the ratio reaches 80, after which the performance begins to decline. Compared with SAC, where the algorithm's performance is impaired when the agent UTD ratio is 20, and increasing the model UTD ratio by only 5 times affects the performance of MBPO, this experimental result suggests that in the setting of MBRL, the primacy bias of the agent is not as significant as MFRL, while the primacy bias of the world model is more pronounced.

The question is, why are the types of the primacy bias different in model-based and model-free setting? We can explain it as follows: Due to the agent in MBRL being trained with data generated by the world model, a couple of factors are at play. Firstly, the number of samples the model generates is typically much larger than the number of samples gathered from the real environment. This abundance of data reduces the likelihood of overfitting. Secondly, the samples generated by the model contain noise, which can be seen as a form of data augmentation. This means that even if the agent UTD ratio is high, it is unlikely to suffer from overfitting, as is often observed in model-free setting. In fact, if the agent UTD ratio is not high enough, underfitting can occur, leading to a decline in performance. On the other hand, when the world model is trained with samples from the real environment, overfitting becomes more likely.

In summary, we draw an important conclusion that: in the setting of MBRL, the primacy bias is more towards the primacy bias of the world model, rather than the primacy bias of the agent.

2.3 The data in the replay buffer is enough for learning the world model

Once the world model overfits the early data, even if there is highquality data in the environment buffer (we call it D_{env}), the world model cannot learn from that data efficiently. Actually, even though the agent trained with a high UTD ratio updated model performs poorly, the data in D_{env} is usable. Still, the model cannot effectively learn from it. We design experiments to verify this. We save the MBPO agent trained with model UTD ratio of 10. Due to the algorithm's poor performance, this agent's policy is far from the optimal policy. As a result, the collected data in D_{env} is also poor, but does this poor data hinder the learning process of the world model? To answer that question, we train another MBPO algorithm with the same model UTD ratio of 10. The difference is that we use the saved policy to initialize the replay buffer, rather than a random policy. We plot the learning curves in Figure 4 Left. It shows that the learning speed of the agent with the pre-initialized buffer is much faster, and the final return is higher, which means the data in the pre-initialized buffer is suitable for policy learning. To further show the quality of the world model, we also record the change process of Model Mean Squared Error (MMSE) and Value-aware Model Error (VME). MMSE measures the difference between model MDP and true MDP, and VME measures the error between model Bellman operator and true Bellman operator given a Q function. The specific definition and calculation of MMSE and VME can be found in Appendix¹ A. We plot the two kinds of model error in Figure 4 Middle, Right. We can see that both model errors are lower with the pre-initialized buffer, which means the data in the pre-initialized buffer is enough for learning a world model.

¹ The full paper is available at https://arxiv.org/abs/2310.15017

Algorithm 1 MBPO with model reset

- 1: **Require:** total environment steps N, agent UTD ratio U_a , model UTD ratio U_m , reset interval I_r , initial policy π_{θ} , predictive model p_{ϕ} , environment dataset \mathcal{D}_{env} , model dataset \mathcal{D}_{model}
- 2: for step from 1 to N do
- 3: Take action in environment according to π_{θ} ; add to \mathcal{D}_{env}
- 4: Sample s_t uniformly from \mathcal{D}_{env}
- 5: Perform k-step model rollout starting from s_t using policy π_{θ} ; add to \mathcal{D}_{model}
- 6: **for** agent step in 1 to U_a **do**
- 7: Update policy π_{θ} using data from $\mathcal{D}_{env} \cup \mathcal{D}_{model}$: $\theta \leftarrow \theta \lambda_{\pi} \nabla_{\theta} J_{\pi}(\theta, \mathcal{D}_{env} \cup \mathcal{D}_{model})$
- 8: end for
- 9: **for** model step in 1 to U_m **do**
- 10: Train model p_{ϕ} on \mathcal{D}_{env} via maximum likelihood
- 11: end for
- 12: **if** step $\% I_r == 0$ **then**
- 13: Reset the world model
- 14: **end if**
- 15: end for

3 Reset the world model instead of the agent

In the previous section, we verify the existence of the primacy bias in MBRL and show that, the primacy bias of the world model is more pronounced than the primacy bias of the agent.

Based on these conclusions, we can infer that the parameter resetting technique is not ineffective but rather used in the wrong way. We advocate that parameter resetting can still improve the performance of MBRL algorithms; only a minor adjustment is needed: reset the parameters of the world model, instead of the agent. We call this method *world model resetting*, summarized as follows:

world model resetting

During the MBRL algorithm training process, periodically resetting the parameters of the agent world model, and the rest of the training process remains unchanged.

The complete pseudo-code of applying *world model resetting* to MBPO is presented in Algorithm 1.

4 **Experiments**

This section aims to validate the effectiveness of *world model resetting* (abbreviated as model reset). To achieve this, we design and conduct numerous experiments on both continuous and discrete tasks. We aim to answer the following questions: (1) Can model reset improve the performance of MBRL algorithms? (2) Is the world model with model reset more accurate? (3) Can model reset show superiority over parameter tuning methods? (4) What are the key factors that influence the success of model reset?

4.1 Overall performance of model reset

The first question is: can model reset improve the performance of MBRL algorithms? To answer that question, we apply model reset to two MBRL algorithms: MBPO and DreamerV2. We choose them because they utilize different types of world model: probabilistic edynamics model and recurrent state-space model (RSSM), respectively. For MBPO, we conduct experiments on four continuous control tasks





 Table 1. Point estimates and 95% confidence interval for the performance of DreamerV2 and DreamerV2 with model reset on 20 tasks on DMC. The results are evaluated over 5 seeds.

Method	IQM	Median	Mean
DreamerV2	762(716,796)	772(732,814)	765(722,801)
DreamerV2 + reset	783 (732,805)	797 (740,819)	786 (734,810)

on MuJoCo. For DreamerV2, we conduct experiments on multiple continuous control tasks on DMC and discrete control tasks on Atari 100k benchmark to verify the effectiveness of model reset across different domains.

For the world model setting of MBPO and DreamerV2, we use the networks and default hyperparameters as specified in the original paper. For MBPO, the world model is an ensemble probabilistic model consisting of 7 models, each with 4 hidden layers. We reset all ensemble models and only reset the last 2 hidden layers every 2×10^4 environment steps. For DreamerV2, the world model consists of an image encoder and a RSSM, and the RSSM consists of a recurrent model, a representation model, and a transition predictor. Among these components, the transition predictor is directly used for behavior learning. Therefore, we only focus on resetting the transition predictor. Due to the complexity of RSSM, we opt for a more conservative Exponential Moving Average (EMA) approach to reset the parameters, which means $\phi_{t+1} = (1 - \alpha)\phi_t + \alpha\phi_{random}$. represents the parameters of the transition predictor, and we set α to 0.8. The transition predictor consists of one hidden layer, and we reset this hidden layer every 2×10^5 environment steps.

4.2 Model reset enhances model accuracy

We plot our results for MBPO with model UTD ratio of 10 on Mu-JoCo tasks in Figure 5. Except for HalfCheetah-v2 task, which appears to be unaffected by the primacy bias, the other three tasks significantly improve algorithm's performance with model reset. As for DreamerV2, we show the aggregated results on Atari 100k benchmark and DMC in Figure 6 and Table 1. We report **the best results**



Figure 6. Point estimates and 95% confidence interval for the performance of DreamerV2 and DreamerV2 with model reset on Atari 100k. The metric is suggested by [1]. The results are evaluated over 5 seeds.



Figure 7. Model mean squared error and Value-aware model error comparison between MBPO and MBPO with model reset. The experiments are conducted on four tasks on MuJoCo. The results are evaluated by 5 runs.

over different model UTD ratios for methods with and without model reset. The results indicate that the algorithm's performance has improved in continuous and discrete control tasks with model reset.² We include more detailed experimental results in Appendix C.1.

The next question is: is the world model with model reset more accurate at collected data points? This question is crucial because it determines whether model reset improves algorithm's performance by enhancing the model's accuracy, rather than other potential factors. To answer that question, we conduct experiments to evaluate the world model in MBPO during the training process. We do not choose DreamerV2 because its world model component is too complex and relies on temporal states, making it difficult to evaluate the model.

We first observe the changes in model error (including model mean square error and value-aware model error) during the training process of MBPO and MBPO with reset. The results are presented in Figure

7. The results indicate that in the later stage of training, the model with model reset can better fit the data in the current buffer. Periodic reset allows the model to learn from new data distribution effectively. In addition to the ability to fit the current data, we also want to know the model's ability to predict the environment dynamics under the near-optimal policy. This can reflect whether the policy has the potential to reach the optimal policy. Therefore, we additionally train a SAC agent on four tasks of MuJoCo for 1 million steps. We keep the replay buffer at the end of the training. When training MBPO, we use the reserved replay buffer to calculate model mean square error. We present the curves in Figure 8. The results show that the model's ability to predict the dynamics under near-optimal policy has improved with model reset. This suggests that the algorithm has greater potential to reach the optimal policy. In fact, a good model does not need to be accurate on every data point but rather should be accurate on the data points that the policy frequently visits. Model reset can be seen as a forgetting mechanism that allows the model to focus more on the data distribution under the current policy, thereby obtaining a better model.

² One might argue that the results in Table 1 show similar performance with and without reset. We point out that the results here are the best results over different model UTD ratios, and the effect of model reset is more obvious for a fixed high model UTD ratio, which we show in Appendix C.1.



Figure 8. Model mean squared error over the data in reserved SAC replay buffer. The experiments are conducted on four tasks of MuJoCo. The results are evaluated by 5 runs.

4.3 Comparison with hyperparameter tuning methods

In the previous sections, we empirically show that model reset can mitigate the primacy bias for MBRL algorithms and improve the algorithms' performance. In this part, we aim to answer the question: **under high model UTD ratio conditions where the primacy bias is pronounced**, whether model reset can show superiority over hyperparameter tuning methods, i.e., tuning real data ratio or rollout length?

As we discuss in Section 2.2, the primacy bias of the world model can be pronounced under a high model UTD ratio, thus harming the performance of MBRL algorithms. A natural idea is to automatically tune model UTD ratio to mitigate the primacy bias. Previous studies follow this idea and propose hyperparameter tuning methods, i.e., [7] utilizes validation sets for dynamically adjusting the model UTD ratio, [18] introduces an event-triggered mechanism to determine when to update the model. Although adjusting model UTD ratio is an effective way to mitigate the primacy bias, it comes at a risk of sacrificing sample efficiency. Therefore, we set model UTD ratio to a large fixed value, intending to compare the effects of model reset and tuning hyperparameters other than model UTD ratio on the performance of MBRL algorithms.

We then introduce our experimental setup. We choose MBPO as our base MBRL algorithm and set the model UTD ratio to $10 \times dUTD$, which is high enough to show primacy bias. As for the hyperparameter tuning method, we choose AutoMBPO [25], a hyperparameter adjusting method which introduces a parametric hypercontroller to sequentially select the value of hyperparameters to maximize the performance of MBPO. The hyper-controller dynamically adjusts hyperparameters including real data ratio, model UTD ratio, agent UTD ratio and rollout length during the training process of MBPO. Since we fix the model UTD ratio, we only adjust real data ratio, agent UTD ratio and rollout length. For model reset, we follow the setup in Section 4.1, resetting all ensemble models and resetting only the last 2 hidden layers every 2×10^4 steps.

We conduct experiments on four tasks of MuJoCo, comparing the performance of AutoMBPO and MBPO with model reset. The experimental results are shown in Figure 9. We can observe that model



Figure 9. Learning curves for AutoMBPO and MBPO with model reset on four tasks of MuJoCo. Model UTD ratio is 10. The results are evaluated by 5 runs.

reset can incur higher overall performance improvement than AutoMBPO under the model UTD ratio of 10. Moreover, the implementation of AutoMBPO is complex since it contains an additional hyper-controller while model reset only needs a few lines of code to change. These demonstrate the advantage of model reset under high model UTD ratio conditions.

4.4 The key factors influencing model reset

In this section, we aim to investigate the key factors that influence the effectiveness of model reset and maximize the use of model reset to improve the performance of MBRL algorithms. Due to space limit, we only present our main points and conclusions here. The detailed results can be found in Appendix C.

Model UTD ratio Model UTD ratio determines the extent to which the model utilizes collected data. A high model UTD ratio implies that the model is prone to fit the current data, potentially increasing the risk of overfitting. In our experiments, we observe the model UTD ratio significantly impacts the effect of model reset. We vary the model UTD ratio of MBPO and DreamerV2 and observe how model reset affects the degree of performance improvement. For MBPO, we conduct experiments on four tasks on MuJoCo. We vary the model UTD ratio to be $2 \times \text{ dUTD}$, $5 \times \text{ dUTD}$, and $10 \times$ dUTD. For DreamerV2, we conduct experiments on 20 tasks on DMC and 26 tasks on Atari 100k. For DMC, We change the default model UTD ratio $(\frac{1}{5})$ to 5 times and 10 times. For Atari 100k, We change the default model UTD ratio $(\frac{1}{16})$ to 4 times and 8 times. Experimental results indicate that, as the model UTD ratio increases, the level of overfitting of the model also increases. Therefore, the degree of performance improvement from model reset is also higher.

Model ensemble size In MBRL algorithms, such as MBPO, the world model can be an ensemble model. The ensemble model utilizes ensemble learning to reduce the variance of prediction and thus mitigate the risk of overfitting to some extent. Therefore, model ensemble size is also one of the potential influencing factors. We conduct experiments using MBPO. We vary the ensemble size to 1,

3, 5, and 7 and apply model reset. The results indicate that as model ensemble size increases, the effectiveness of model reset diminishes.

When and how to reset the model The last factor we examine is how to conduct model reset, such as how often to perform a reset and which part of the network to reset. These factors may have a significant impact on the effectiveness of model reset. Through this investigation, we hope to provide insights and a guideline on performing model reset effectively.

In our experiments, if the model has multiple hidden layers, resetting only the last few layers yields the best results compared to resetting the first few hidden layers or randomly resetting hidden layers. We conjecture that the first few layers can extract the general features of the task [33], allowing the world model to adapt to the new data distribution after a reset quickly. Therefore, retaining the first few hidden layers when performing reset is preferable. We also find that, for ensemble model, resetting the majority or all ensemble members has a better effect than resetting only one or few of them. Regarding when to perform reset, we find that it depends on the specific task and model UTD ratio. Generally, if the model UTD ratio is higher, the reset frequency should also be higher.

5 Related Work

The primacy bias in reinforcement learning The primacy bias is originally a concept in cognitive science [29], which means first impressions influence human consciousness and people tend to pay less attention to subsequent impressions. [31] first propose that in deep reinforcement learning, the primacy bias also exists, and is related to the agent's overfitting of early data. In fact, the phenomenon of overfitting [45, 44] in deep reinforcement learning has been studied a lot. [37] point out that when using the Q function for generalization, there might be overfitting; [28] show the approximator will lose the ability to generalize during training in off-policy RL; [4] show the "self-overfitting" issue exists in image-based RL with an image decoder. It is worth mentioning that overfitting and primacy bias in RL cannot be equated. For example, overfitting can also occur in offline setting [11, 23], but since the dataset is static, it cannot be called the primacy bias. The primacy bias often occurs in high UTD ratio scheme [41, 9], causing the failure of the algorithm. Previous works attempt to handle this in various ways, such as randomized ensemble [5], network normalization [16, 35], replay ratio scaling [8], and so on. [31] propose a simple method called parameter resetting. By periodically resetting the parameters of the agent, it can effectively deal with the scheme with a high UTD ratio. These methods are mainly about model-free setting, while we further focus on model-based setting.

Overfitting in model-based RL In model-based setting, most previous work focuses on handling the issue that the agent overfits an inaccurate world model [19, 2]. Commonly used methods are uncertainty estimation [6], and policy regularization [46]. We need to mention that the overfitting here is different from the primacy bias we discuss above, because it is caused by an inaccurate model, and we focus on the overfitting caused by a high UTD ratio. There are also studies focusing on the overfitting of the world model itself. Standard solutions include using different model structures, such as Bayesian neural networks [32], and supervised learning methods, such as spectral regularization [12] and dropout [36]. Some recent studies aim to mitigate model overfitting by adjusting model UTD ratio by

utilizing validation sets. However, different network structures may bring computational burdens, and adjusting UTD ratio may sacrifice sample efficiency. *World model resetting* proposed by us neither changes the network structure nor adjusts UTD ratio, but can be seen as a way to make the world model pay more attention to current transitions, thereby alleviating overfitting. However, previous works do not explore the primary and secondary relationship between agent overfitting and model overfitting, which is the focus of our work.

Model adaptation Our work is also related to model adaptation. In model-based RL, the purpose of model adaptation is to enable the world model to adapt when encountering new data or new tasks quickly. Meta learning [39, 34] is often used for model adaptation [26, 13]. These methods use a prior model to help the agent quickly adapt to new tasks. Other methods include system identification [24], curious replay [21], drifting Gaussian process [30]. Our work differs from these methods because we focus on model learning in a fixed environment, while they are mainly about a changing environment. PDML [43] learns a model which can adapt to the evolving policy in the same environment. However, PDML guides the model to pay more attention to the current transitions mainly by re-weighting the transition distribution in the replay buffer, while we consider it from the perspective of the primacy bias, allowing the world model to quickly adapt to the current policy by world model resetting, which is simpler to apply.

6 Limitations

Our work has several limitations. One is that model reset is only effective when the primacy bias is significant (such as, in high model UTD ratio scenario). Another limitation is that applying resetting technique to complex network structures (such as Transformers [42].) is challenging. However, addressing these issues is beyond the scope of this work. For future work, we can focus on exploring other more effective methods to reduce the primacy bias, which we believe is an interesting topic.

7 Conclusion

In this paper, we focus on the primacy bias in MBRL. We first assume that the primacy bias is composed of the primacy bias of the agent and the primacy bias of the world model. However, our experiments find parameter resetting is ineffective for MBRL algorithms. In further investigations, we discover that a high agent UTD ratio benefits the algorithm, while a high model UTD ratio hurts the performance. Therefore, we conclude that the primacy bias in MBRL is primarily composed of the primacy bias of the world model. Additionally, we propose *world model resetting*, which works in MBRL setting.

In the validation experiments of our method, we demonstrate that *world model resetting* can improve the performance of the MBRL algorithms across multiple domains, and the world model with model reset is more accurate. We also compare *world model resetting* with AutoMBPO to illustrate the advantage of our method under high model UTD ratio conditions. Finally, we investigate the key factors that influence the effectiveness of model reset. Our intention is not to propose a totally new and state-of-the-art method for improving MBRL algorithms' performance. Instead, *world model resetting* is simple yet effective, and we hope our work can shed light on addressing the issues of model overfitting and the primacy bias in MBRL.

Acknowledgements

This work was supported by the STI 2030-Major Projects under Grant 2021ZD0201404. The authors would also like to thank the anonymous reviewers for valuable comments.

References

- R. Agarwal, M. Schwarzer, P. S. Castro, A. C. Courville, and M. Bellemare. Deep reinforcement learning at the edge of the statistical precipice. *Advances in neural information processing systems*, 34: 29304–29320, 2021.
- [2] D. Arumugam, D. Abel, K. Asadi, N. Gopalan, C. Grimm, J. K. Lee, L. Lehnert, and M. L. Littman. Mitigating planner overfitting in modelbased reinforcement learning. arXiv preprint arXiv:1812.01129, 2018.
- [3] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig. Safe learning in robotics: From learning-based control to safe reinforcement learning. *Annual Review of Control, Robotics,* and Autonomous Systems, 5:411–444, 2022.
- [4] E. Cetin, P. J. Ball, S. Roberts, and O. Celiktutan. Stabilizing off-policy deep reinforcement learning from pixels. In *International Conference* on Machine Learning, pages 2784–2810. PMLR, 2022.
- [5] X. Chen, C. Wang, Z. Zhou, and K. W. Ross. Randomized ensembled double q-learning: Learning fast without a model. In *International Conference on Learning Representations*, 2020.
- [6] K. Chua, R. Calandra, R. McAllister, and S. Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. Advances in neural information processing systems, 31, 2018.
- [7] N. Dorka, T. Welschehold, and W. Burgard. Dynamic updateto-data ratio: Minimizing world model overfitting. arXiv preprint arXiv:2303.10144, 2023.
- [8] P. D'Oro, M. Schwarzer, E. Nikishin, P.-L. Bacon, M. G. Bellemare, and A. Courville. Sample-efficient reinforcement learning by breaking the replay ratio barrier. In *Deep Reinforcement Learning Workshop NeurIPS* 2022, 2022.
- [9] W. Fedus, P. Ramachandran, R. Agarwal, Y. Bengio, H. Larochelle, M. Rowland, and W. Dabney. Revisiting fundamentals of experience replay. In *International Conference on Machine Learning*, pages 3061– 3071. PMLR, 2020.
- [10] S. Fujimoto, H. Hoof, and D. Meger. Addressing function approximation error in actor-critic methods. In *International conference on machine learning*, pages 1587–1596. PMLR, 2018.
- [11] S. Fujimoto, D. Meger, and D. Precup. Off-policy deep reinforcement learning without exploration. In *International conference on machine learning*, pages 2052–2062. PMLR, 2019.
- [12] F. Gogianu, T. Berariu, M. C. Rosca, C. Clopath, L. Busoniu, and R. Pascanu. Spectral normalisation for deep reinforcement learning: an optimisation perspective. In *International Conference on Machine Learning*, pages 3734–3744. PMLR, 2021.
- [13] J. Guo, M. Gong, and D. Tao. A relational intervention approach for unsupervised dynamics generalization in model-based reinforcement learning. In *International Conference on Learning Representations*, 2021.
- [14] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Offpolicy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, pages 1861– 1870. PMLR, 2018.
- [15] D. Hafner, T. P. Lillicrap, M. Norouzi, and J. Ba. Mastering atari with discrete world models. In *International Conference on Learning Rep*resentations, 2020.
- [16] T. Hiraoka, T. Imagawa, T. Hashimoto, T. Onishi, and Y. Tsuruoka. Dropout q-functions for doubly efficient reinforcement learning. In *International Conference on Learning Representations*, 2021.
- [17] M. Janner, J. Fu, M. Zhang, and S. Levine. When to trust your model: Model-based policy optimization. Advances in neural information processing systems, 32, 2019.
- [18] T. Ji, Y. Luo, F. Sun, M. Jing, F. He, and W. Huang. When to update your model: Constrained model-based reinforcement learning. *Advances in Neural Information Processing Systems*, 35:23150–23163, 2022.
- [19] N. Jiang, A. Kulesza, S. Singh, and R. Lewis. The dependence of effective planning horizon on model accuracy. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1181–1189, 2015.
- [20] Ł. Kaiser, M. Babaeizadeh, P. Miłos, B. Osiński, R. H. Campbell, K. Czechowski, D. Erhan, C. Finn, P. Kozakowski, S. Levine, et al. Model based reinforcement learning for atari. In *International Conference on Learning Representations*, 2019.

- [21] I. Kauvar, C. Doyle, L. Zhou, and N. Haber. Curious replay for modelbased adaptation. arXiv preprint arXiv:2306.15934, 2023.
- [22] B. R. Kiran, I. Sobh, V. Talpaert, P. Mannion, A. A. Al Sallab, S. Yogamani, and P. Pérez. Deep reinforcement learning for autonomous driving: A survey. *IEEE Transactions on Intelligent Transportation Systems*, 23(6):4909–4926, 2021.
- [23] A. Kumar, A. Zhou, G. Tucker, and S. Levine. Conservative q-learning for offline reinforcement learning. Advances in Neural Information Processing Systems, 33:1179–1191, 2020.
- [24] A. Kumar, Z. Fu, D. Pathak, and J. Malik. Rma: Rapid motor adaptation for legged robots. arXiv preprint arXiv:2107.04034, 2021.
- [25] H. Lai, J. Shen, W. Zhang, Y. Huang, X. Zhang, R. Tang, Y. Yu, and Z. Li. On effective scheduling of model-based reinforcement learning. *Advances in Neural Information Processing Systems*, 34:3694–3705, 2021.
- [26] K. Lee, Y. Seo, S. Lee, H. Lee, and J. Shin. Context-aware dynamics model for generalization in model-based reinforcement learning. In *International Conference on Machine Learning*, pages 5757–5766. PMLR, 2020.
- [27] Q. Li, A. Kumar, I. Kostrikov, and S. Levine. Efficient deep reinforcement learning requires regulating overfitting. In *The Eleventh International Conference on Learning Representations*, 2022.
- [28] C. Lyle, M. Rowland, and W. Dabney. Understanding and preventing capacity loss in reinforcement learning. In *International Conference on Learning Representations*, 2021.
- [29] P. H. Marshall and P. R. Werder. The effects of the elimination of rehearsal on primacy and recency. *Journal of Verbal Learning and Verbal Behavior*, 11(5):649–653, 1972.
- [30] F. Meier and S. Schaal. Drifting gaussian processes with varying neighborhood sizes for online model learning. In *International Conference* on Robotics and Automation (ICRA), pages 264–269. IEEE, 2016.
- [31] E. Nikishin, M. Schwarzer, P. D'Oro, P.-L. Bacon, and A. Courville. The primacy bias in deep reinforcement learning. In *International Conference on Machine Learning*, pages 16828–16847. PMLR, 2022.
- [32] M. Okada and T. Taniguchi. Variational inference mpc for bayesian model-based reinforcement learning. In *Conference on robot learning*, pages 258–272. PMLR, 2020.
- [33] K. Ota, T. Oiki, D. Jha, T. Mariyama, and D. Nikovski. Can increasing input dimensionality improve deep reinforcement learning? In *International conference on machine learning*, pages 7424–7433. PMLR, 2020.
- [34] N. Schweighofer and K. Doya. Meta-learning in reinforcement learning. *Neural Networks*, 16(1):5–9, 2003.
- [35] L. Smith, I. Kostrikov, and S. Levine. A walk in the park: Learning to walk in 20 minutes with model-free reinforcement learning. arXiv preprint arXiv:2208.07860, 2022.
- [36] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [37] R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [38] Y. Tassa, Y. Doron, A. Muldal, T. Erez, Y. Li, D. d. L. Casas, D. Budden, A. Abdolmaleki, J. Merel, A. Lefrancq, et al. Deepmind control suite. arXiv preprint arXiv:1801.00690, 2018.
- [39] S. Thrun. Lifelong learning algorithms. In *Learning to learn*, pages 181–209. Springer, 1998.
- [40] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for modelbased control. In 2012 IEEE/RSJ international conference on intelligent robots and systems, pages 5026–5033. IEEE, 2012.
- [41] H. P. Van Hasselt, M. Hessel, and J. Aslanides. When to use parametric models in reinforcement learning? Advances in Neural Information Processing Systems, 32, 2019.
- [42] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. Advances in neural information processing systems, 30, 2017.
- [43] X. Wang, W. Wongkamjan, R. Jia, and F. Huang. Live in the moment: Learning dynamics model adapted to evolving policy. In *International Conference on Machine Learning*, pages 36470–36493. PMLR, 2023.
- [44] A. Zhang, N. Ballas, and J. Pineau. A dissection of overfitting and generalization in continuous reinforcement learning. arXiv preprint arXiv:1806.07937, 2018.
- [45] C. Zhang, O. Vinyals, R. Munos, and S. Bengio. A study on overfitting in deep reinforcement learning. arXiv preprint arXiv:1804.06893, 2018.
- [46] R. Zheng, X. Wang, H. Xu, and F. Huang. Is model ensemble necessary? model-based rl via a single model with lipschitz regularized value function. In *The Eleventh International Conference on Learning Representations*, 2022.