

# What Model Does MuZero Learn?

Jinke He<sup>a,\*</sup>, Thomas M. Moerland<sup>b</sup>, Joery A. de Vries<sup>a</sup> and Frans A. Oliehoek<sup>a</sup>

<sup>a</sup>Delft University of Technology  
<sup>b</sup>Leiden University

**Abstract.** Model-based reinforcement learning (MBRL) has drawn considerable interest in recent years, given its promise to improve sample efficiency. Moreover, when using deep-learned models, it is possible to learn compact and generalizable models from data. In this work, we study MuZero, a state-of-the-art deep model-based reinforcement learning algorithm that distinguishes itself from existing algorithms by learning a value-equivalent model. Despite MuZero's success and impact in the field of MBRL, existing literature has not thoroughly addressed why MuZero performs so well in practice. Specifically, there is a lack of in-depth investigation into the value-equivalent model learned by MuZero and its effectiveness in model-based credit assignment and policy improvement, which is vital for achieving sample efficiency in MBRL. To fill this gap, we explore two fundamental questions through our empirical analysis: 1) to what extent does MuZero achieve its learning objective of a value-equivalent model, and 2) how useful are these models for policy improvement? Among various other insights, we conclude that MuZero's learned model cannot effectively generalize to evaluate unseen policies. This limitation constrains the extent to which we can additionally improve the current policy by planning with the model.

## 1 Introduction

In recent years, deep reinforcement learning (DRL) [9, 48] has achieved remarkable progress, finding applications in a variety of real-world problems such as video compression [33], chip design [35], inventory management [32] and plasma control for nuclear fusion [7]. However, despite these advancements, sample inefficiency remains a significant obstacle that limits the broader applicability of deep reinforcement learning in practical settings.

Model-based reinforcement learning (MBRL) [37] addresses sample inefficiency in DRL by learning predictive models of the environment. In the typical RL cycle, an agent interacts with the environment (acting) and uses the collected data to refine its policy (learning). Accordingly, MBRL methods fall into two non-mutually exclusive categories [51]: (1) those that use the learned model to improve acting and (2) those that use the learned model to improve learning.

One notable example of the first category is model-based exploration. In environments where rewards are sparse, shallow exploration techniques such as epsilon-greedy exploration often fail. Model-based exploration addresses this by not only allowing the agent to use model prediction errors or model uncertainty as intrinsic learning signals, but also enabling more effective exploration of "interesting" regions of the environment by "planning to explore"

[39, 22, 43, 40, 42, 3, 31]. In addition, decision-time planning methods such as Monte Carlo Tree Search (MCTS) [4] compute local policies online by planning with the model. A representative class of methods that use MCTS for decision-time planning is AlphaZero [47, 44, 46], which defeated a human Go world champion for the first time in human history.

The other class of MBRL methods aims to improve the agent's policy and value functions without consuming additional data through model-based credit assignment. At a high level, these methods generate synthetic data through the learned model to simulate potential outcomes of specific actions or policies. This synthetic data is then used to update the agent's value estimates and improve the policy with a policy improvement operator. Notably, the DYNA architecture [48] treats synthetic data as real data, integrating it into model-free learning algorithms, whereas the Dreamer methods [20, 18, 19] use synthetic data exclusively for policy computation. In addition to using synthetic data directly, AlphaZero employs MCTS to compute a refined local policy and improved value estimates, which then serve as better learning targets for policy and value functions in replacement of model-free targets. Furthermore, van Hasselt et al. [51] advocate for backward planning, which assigns credits to hypothetical states through a learned inverse model. They argue that planning backward for credit assignment can be more robust to model errors than planning forward, as updating fictional states can be less harmful than updating real states with fictional values.

Model-based exploration and credit assignment improve the data efficiency of RL methods by collecting more useful data and extracting more information from it. Beyond these two categories, learning a predictive model of the environment can also serve as an auxiliary task for representation learning [23, 25]. In addition, in model-free planning, differentiable computation graphs that resemble the structure of planning with a model (i.e., implicit models) have been found useful as architectural priors for value and policy functions, demonstrating improved performance in combinatorial planning domains while trained via model-free losses [38, 49, 8, 15, 16].

Unlike tabular methods that learn the dynamics and reward for each state-action pair, deep model-based RL (DMBRL) approaches typically learn a state representation, on top of which dynamics and reward functions are estimated. However, determining what relevant information to include in these state representations and how to effectively learn them remains challenging. Namely, not all features of the observations are relevant [30], and missing relevant features may lead to history dependence [34]. A prevalent approach within DMBRL for learning these state representations is to model the next observation [26, 17, 18, 19, 20]. However, accurately predicting high-dimensional observations, such as images, requires considerable ef-

\* Corresponding Author. Email: J.He-4@tudelft.nl.

fort in designing and training high-capacity neural networks with effective inductive biases [26]. This challenge has long been a barrier for MBRL methods, as the efficacy of model-based policy optimization strongly depends on the quality of these representations. Moreover, this approach often wastes significant representational power and training resources on encoding task-irrelevant information within the state representations, causing inefficiency in learning.

One approach to addressing this challenge is the development of value-equivalent models [45, 12, 13, 14, 8, 49, 38]. These models are specifically trained to predict the (multi-step) Bellman update, focusing solely on value-relevant aspects of the task dynamics without needing to reconstruct any observation. MuZero [41], a well-known MBRL algorithm, exemplifies this approach by achieving state-of-the-art performance in Atari games [2] and superhuman performance in the game of Go, Chess, and Shogi.

MuZero inherits much of its structure from its predecessor, AlphaZero, which uses MCTS guided by both a learned policy network and a learned value network to make decisions and generate learning targets. However, MuZero distinguishes itself by integrating a model learned jointly with the value and policy networks, contrasting AlphaZero’s use of a ground-truth model for simulation and search. Importantly, MuZero’s model is not trained to predict the next state or observation, but instead focuses on predicting task-relevant quantities such as future rewards, policies, and values. This approach of learning implicit models sets MuZero apart from traditional MBRL algorithms and promises a shift towards more efficient model learning. Despite MuZero’s empirical success and its considerable impact on MBRL [1, 24, 33, 52], a recent study by [6] shows that MuZero’s dynamics model can diverge significantly from real transitions, highlighting a gap in our understanding of how these models function and their efficacy in model-based credit assignment and planning. This discrepancy underscores the necessity for a detailed investigation into the capabilities and limitations of value-equivalent models within MuZero, which motivates this work.

The most relevant study in this direction is by [21], who studied the role of planning in MuZero’s learning. They found that planning primarily boosts the policy network’s learning by generating more informative data and constructing better training targets. Surprisingly, their findings also reveal that, in most domains, planning at evaluation time does not significantly improve performance compared to using the policy network alone, even with large search budgets. One explanation is that the policy network has converged to the optimal policy, rendering planning less useful. Furthermore, Danihelka et al. [5] show that improving MuZero’s planning enables strong performance even with extremely small search budgets ( $n = 2$ , or  $n = 4$ ) in Go. This raises the question, what is the contribution of MuZero’s learned model.

In this study, we aim to bridge the gap in our understanding of MuZero’s learned model by exploring two fundamental questions:

1. To what extent does MuZero learn a value-equivalent model?
2. To what extent does MuZero learn a model that supports effective policy improvement (through planning)?

Learning a truly value-equivalent model is essential for model-based credit assignment, which directly influences the potential for improving existing policies through model-based planning. The more effectively we can improve existing policies through planning with the learned model, the greater the sample efficiency achieved by an MBRL method. As such, addressing these questions can help us better understand MuZero’s empirical success and inform the design of future algorithms or extensions.

Through our empirical analysis, we find that MuZero’s learned model is generally not accurate enough for policy evaluation, and the accuracy of the model decreases as the policy to evaluate deviates further from MuZero’s data collection policy. Consequently, this limits the extent to which we can find a good policy via planning. However, we find that MuZero’s incorporation of the policy prior in MCTS alleviates this problem, which biases the search towards actions where the model is more accurate. Based on these findings, we speculate that the role of the model in MuZero may be similar to that in model-free planning, providing a more powerful representation of value and policy functions, as the extent to which it can support policy improvement is rather limited. Moreover, using the policy prior indirectly takes model uncertainty into account during planning, which results in a form of regularized policy optimization with the learned model.

In Section 2, we introduce the essential background. Section 3 studies trained MuZero models in the policy *evaluation* setting (the objective for which the models were trained). Section 4 extends the analysis to the policy *improvement* setting (planning). Finally, Section 5 discusses the limitations and outlook of this work.

## 2 Background

### 2.1 Markov Decision Processes

A Markov decision process (MDP) is a model that describes the interaction between a decision-making agent and an environment. Formally, an infinite-horizon discounted MDP is a 6-tuple  $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \mu, \gamma)$  where  $\mathcal{S}$  represents the state space of the environment and  $\mathcal{A}$  denotes the set of actions that the agent can take. The initial state of the environment,  $s_0$ , follows a distribution  $\mu \in \Delta(\mathcal{S})$ . At each time step  $t$ , the agent observes the environment state  $s_t$  and selects an action  $a_t$ , causing the environment to transition to a new state  $s_{t+1} \sim \mathcal{T}(\cdot|s_t, a_t)$  and return a numerical reward  $r_t = \mathcal{R}(s_t, a_t)$ . Given a stationary policy  $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$ , the value function  $V$  measures the expected discounted sum of future rewards from a state following  $\pi$  afterward,  $V^\pi(s) = E[\sum_{t=0}^{\infty} \gamma^t r_t | \pi, s_0 = s]$ , where  $\gamma \in [0, 1)$  is the discount factor. In this work, we specifically focus on environments with deterministic dynamics, where the next state  $s_{t+1}$  is solely determined by the current state  $s_t$  and action  $a_t$  with probability 1:  $s_{t+1} = \mathcal{T}(s_t, a_t)$ .

### 2.2 MuZero

MuZero [41] is a recent MBRL method that has achieved state-of-the-art performance in Atari games and matched the superhuman performance of AlphaZero [47] in Go, chess, and shogi.

**Components** Similar to other MBRL methods, MuZero learns a deterministic world model that consists of a representation function  $h_\theta$  and a dynamics function  $g_\theta$ . The representation function encodes an environment state  $s_t$  into a latent state  $z_t^0 = h_\theta(s_t)$ . Here, the subscript denotes the real time step at which the encoding occurs, and the superscript denotes the number of time steps that have been spent in the learned model since then. Given a latent state  $z_t^k$  and an action  $a_{t+k}$ , the dynamics function predicts the next latent state  $z_t^{k+1}$  and the reward  $u_t^k$ ,  $(z_t^{k+1}, u_t^k) = g_\theta(z_t^k, a_{t+k})$ . Apart from the representation and dynamics functions, MuZero uses a prediction function to predict the value and policy at a latent state,  $\pi_t^k, v_t^k = f_\theta(z_t^k)$ . For ease of notation, we split the prediction function into a policy function  $\pi_\theta(z_t^k)$  and a value function  $v_\theta(z_t^k)$ . To distinguish between the policy function  $\pi_\theta$  and the MuZero policy  $\pi^{\text{MuZero}}$ , which

runs MCTS, we will refer to the former as the policy prior and the latter as MuZero’s behavior policy.

**Acting** MuZero makes decisions by planning with the learned model. At each time step  $t$ , MuZero encodes the environment state  $s_t$  into a latent state  $z_t^0$  and uses it as the root node to perform MCTS. As the result of the search, MuZero selects the action  $a_t$  by sampling from a distribution that is constructed using the visit counts at the root node and a temperature parameter  $T \in (0, \infty)$ :

$$\pi^{\text{MuZero}}(a|s_t) = \frac{N(z_t^0, a)^{1/T}}{\sum_b N(z_t^0, b)^{1/T}} \quad (1)$$

MuZero’s planning differs from traditional MCTS methods like UCT [27] in two key ways. First, rather than employing random rollouts for leaf node value estimation, MuZero uses its learned value function  $v_\theta$ . Second, MuZero integrates the policy prior into its action selection, guiding the simulation of actions at tree nodes:

$$\arg \max_a \left[ Q(z, a) + c \cdot \pi_\theta(a|z) \cdot \frac{\sqrt{\sum_b N(z, b)}}{1 + N(z, a)} \right] \quad (2)$$

where  $c = c_1 + \log\left(\frac{\sum_b N(z, b) + c_2 + 1}{c_2}\right)$  with  $c_1 = 1.25$  and  $c_2 = 19652$ .  $Q$  and  $N$  are the estimated values and visit counts of actions. MuZero inherits much of this search mechanism from AlphaZero, including the use of policy prior and value functions. A recent study by Grill et al. [11] shows that the use of policy prior in MCTS effectively makes the action visit distribution at the root node track the solution of a local regularized policy optimization problem.

**Training** The key difference between MuZero and many prior works in MBRL lies in their approach to learning the model. In Figure 1, we illustrate the loss function of MuZero. Essentially, given a segment of a real episode that starts from state  $s_t$ , MuZero unrolls its model for  $K$  simulated time steps (below the dotted line) and compares these to real experiences (above the dotted line). This comparison results in a loss consisting of three terms for each of the  $K$  steps. The first term is the per-step reward prediction loss, with the target being the real reward received. The second term is the policy prediction loss, with the target being the action visit distribution of the MCTS at the root node. The third term is the value prediction loss, with the target being the discounted sum of  $n$ -step real rewards plus a value estimation bootstrapped from MCTS  $n$  steps into the future  $v_t^{\text{target}} = \sum_{k=0}^{n-1} \gamma^k r_{t+k} + \gamma^n v_{t+n}^{\text{MCTS}}$ . For simplicity, we omit the links for value targets in Figure 1. All components of MuZero are trained jointly end-to-end by minimizing the aggregated loss. In practice, MuZero incorporates various additional techniques from the literature, such as prioritized experience replay, to improve training. It also introduces a novel algorithm called ‘reanalyse’ to generate fresh targets from old trajectories by re-running MCTS on them using the latest network. In our experiments, we followed the original paper and used all the aforementioned techniques to train our MuZero agents.

### 2.3 The value equivalence principle

The value equivalence principle, introduced by Grimm et al. [12, 13, 14], is motivated by the consideration that the construction of a model should take into account the final use of the model. It defines the order- $k$  value equivalence class  $\mathcal{M}^k(\Pi, \mathcal{V})$  as the subset of all models that can predict the correct  $k$ -step Bellman update for any  $\pi \in \Pi$  and  $v \in \mathcal{V}$  in a set of policies  $\Pi$  and functions  $\mathcal{V}$ . For  $k \rightarrow \infty$ , an additional proper value equivalence class  $\mathcal{M}^\infty(\Pi)$  is defined, which excludes the set of functions from the specification. In essence, a model in the proper value equivalence class must

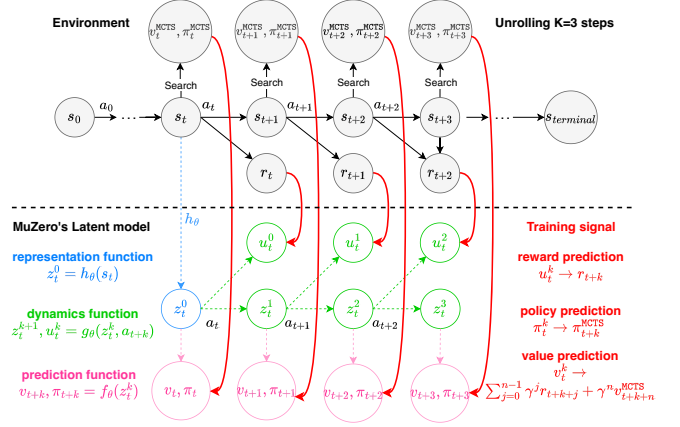


Figure 1: An illustration of MuZero’s loss function.

have the true value function as the fixed point of the Bellman operators for all policies in this set. As such, these models can be seen as the  $Q^\Pi$ -irrelevance abstractions for the MDP [30]. Notably, the largest proper value equivalence class that guarantees optimal planning is  $\mathcal{M}^\infty(\Pi^{\text{DET}})$ , where  $\Pi^{\text{DET}}$  is the set of all deterministic policies. Moreover, Grimm et al. [13] show that a simplified version of MuZero’s loss upper bounds the proper value equivalence loss, making an explicit connection between the theory of value equivalence principle and the strong empirical performance of MuZero.

## 3 Policy Evaluation Experiments


We are interested in the extent to which MuZero’s learned model supports additional policy improvement, which is crucial for MuZero’s sample efficiency as an MBRL method. Our hypothesis is that, since MuZero’s model is trained on data collected by previous policies, it is not generally value equivalent for all policies, especially those that have not been executed. As accurate policy evaluation is the basis for effective policy improvement, this will limit the extent to which MuZero can additionally improve its policy through planning. In this section, we validate this hypothesis.

### 3.1 Training MuZero agents

For our empirical analysis, we used three fully observable deterministic environments, as MuZero was designed for this setting: Cart Pole, a deterministic version of Lunar Lander, and Atari Breakout [2]. We trained 30 MuZero agents with different random seeds for Cart Pole and Lunar Lander, and 20 for Atari Breakout. For each agent, we saved the model weights at different training steps. In the figures below, we aggregate results from different seeds/agents and report their means as well as the corresponding standard errors.

In Cart Pole and Lunar Lander, we extensively trained the MuZero agents, for 100K and 1M steps. This way, we can conduct our analysis on the trained agents throughout their lifecycle of learning. For Atari Breakout, we adopted the same setup as EfficientZero [52] but extended the training from 100K steps to 500K steps.

Figure 2 shows the online performance of MuZero agents at various training steps, which make decisions by performing MCTS planning (Equation 1). For reference, we also plot the performance of the policy prior of the same agent, which samples actions directly from the policy network,  $a_t \sim \pi_\theta(\cdot|h_\theta(s_t))$ . Throughout training, planning enables MuZero to achieve a substantially better performance than directly using the policy prior. It is important to note that the improved performance may not only come from the learned model



**Figure 2:** Online performance of MuZero agents in Cart Pole (Left), Lunar Lander (Middle) and Atari Breakout (Right).

but may also come from the value network, which is used to estimate the value of leaf nodes in MuZero’s MCTS. In the following, we will take a deeper look at how much the learned model contributes to this.

### 3.2 Evaluating the learned model

To assess MuZero’s learned model in the context of policy evaluation, we will compare it to the ground-truth model. Through experiments, we aim to answer questions in the format of *how much value prediction error should we expect when using MuZero’s learned model to evaluate a policy  $\pi$* ?

Considering that MuZero does not train or employ the model for infinite-horizon rollouts, we impose a limit on the evaluation horizon when assessing the value prediction error. For each state  $s_t$ , we define the discounted sum of future rewards for taking an action sequence  $(a_t, \dots, a_{t+h-1})$  as:

$$v^{a_t:t+h-1}(s_t) = \sum_{k=0}^{h-1} \gamma^k r_{t+k} \quad (3)$$

where  $r_{t+k} = \mathcal{R}(s_{t+k}, a_{t+k})$  and  $s_{t+k+1} = \mathcal{T}(s_{t+k}, a_{t+k})$ , assuming the environment is deterministic. As mentioned in Section 2.2, MuZero predicts the value of this action sequence by first encoding the state into a latent state  $z_t^0 = h_\theta(s_t)$  and then rolling out the model from the latent state:

$$\hat{v}^{a_t:t+h-1}(s_t) = \sum_{k=0}^{h-1} \gamma^k u_t^k \quad (4)$$

where  $z_t^{k+1}, u_t^k = g_\theta(z_t^k, a_{t+k})$ . Then, we can define the value prediction error of the learned model for the action sequence  $a_{t:t+h-1}$ .

**Definition 1.** *The value prediction error of using the learned model to evaluate the action sequence  $a_{t:t+h-1}$  at state  $s_t$  is:*

$$|v^{a_t:t+h-1}(s_t) - \hat{v}^{a_t:t+h-1}(s_t)| \quad (5)$$

As a stationary policy  $\pi : \mathcal{S} \rightarrow \Delta(\mathcal{A})$  that operates in the original environment defines a distribution over such action sequences:

$$\Pr(a_{t:t+h-1}|s_t, \pi) = \prod_{k=0}^{h-1} \pi(a_{t+k}|s_{t+k} = \mathcal{T}(s_{t+k-1}, a_{t+k-1})) \quad (6)$$


We can define the value prediction error for evaluating  $\pi$ .

**Definition 2.** *The value prediction error of using the learned model to evaluate a stationary policy  $\pi$ , which operates in the original environment, for horizon  $h$  is:*

$$|v_h^\pi(s_t) - \hat{v}_h^\pi(s_t)| = \left| \mathbb{E}_{\substack{a_{t:t+k-1} \\ \sim \Pr(\cdot|s_t, \pi)}} [v^{a_t:t+h-1}(s_t) - \hat{v}^{a_t:t+h-1}(s_t)] \right| \quad (7)$$

### 3.3 How accurately can MuZero’s learned model predict the value of its own behavior policy?

As the model is trained on data collected by MuZero’s behavior policy  $\pi^{\text{MuZero}}$ , we expect it to be at least accurate on this data collection policy. Therefore, we begin our investigation by examining the model’s prediction performance on this policy. Due to the continuous state spaces of our environments, it is not feasible to enumerate all states and compute the error for each one of them. Instead, we sample states from MuZero’s on-policy state distribution  $d_{\pi^{\text{MuZero}}}$ . For each sampled state, we conduct the evaluation and aggregate the errors. To facilitate tractable evaluation, at every state  $s$ , we use Monte Carlo sampling to estimate both the true value  $v_h^{\pi^{\text{MuZero}}}(s)$  and the value predicted by MuZero’s model  $\hat{v}_h^{\pi^{\text{MuZero}}}(s)$ .



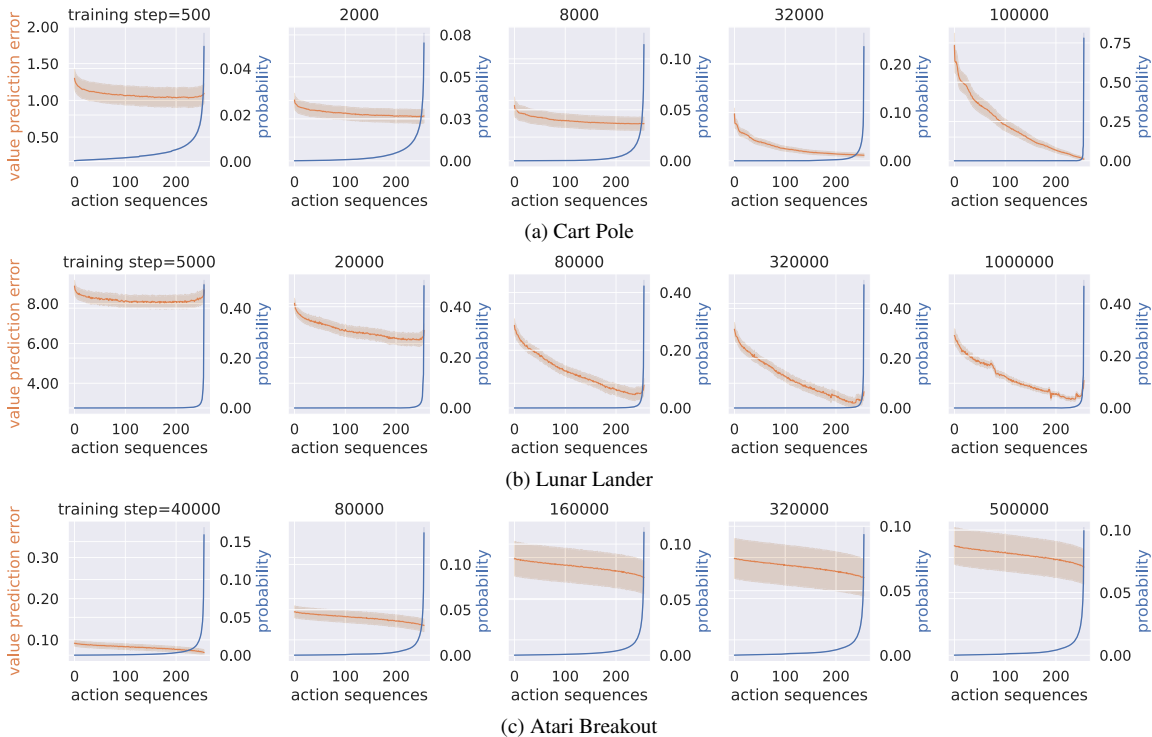
**Figure 3:** Value prediction error of using MuZero’s learned model to evaluate its own behavior policy.

In Figure 3, we report the value prediction error (Y-axis) under various evaluation horizons (lines) across different training steps (X-axis). The maximum evaluation horizons are set to the number of unrolling steps during training, which are  $\{10, 10, 5\}$  for Cart Pole, Lunar Lander, and Atari Breakout, respectively. In all environments, value predictions for short horizons are highly accurate, with errors approaching zero. However, as the evaluation horizon increases, the errors consistently grow larger. On the one hand, this is not surprising because learned models are known prone to accumulate errors during long rollouts [29]. On the other hand, this shows that models learned by MuZero cannot be fully value equivalent as they are not even accurate enough to predict values for the policy that collects the training data. Errors at different training steps are generally not comparable due to the different state distributions, but the decreasing errors observed at the end of training suggest the convergence of the policy as a possible explanation.

### 3.4 How accurately can MuZero’s learned model evaluate policies that are different from the behavior policy?

To assess whether MuZero’s learned model effectively supports planning, we will investigate its capacity to generalize and accurately predict values beyond its own data collection policy. Our hypothesis is that the model will exhibit increasing inaccuracies when evaluating policies that differ significantly from the behavior policy, which is responsible for collecting the training data.

To test this hypothesis, we conduct an experiment focusing on the relationship between the value prediction error for an action sequence  $|v^{a_t:t+h-1}(s_t) - \hat{v}^{a_t:t+h-1}(s_t)|$  and the probability of the behavior policy selecting this action sequence  $\Pr(a_{t:t+h-1}|s_t, \pi^{\text{MuZero}})$ . We again sample states from MuZero’s on-policy state distribution  $d_{\pi^{\text{MuZero}}}$  and compute the probabilities and



**Figure 4:** X-axis: action sequences sorted by their probabilities of being taken by the behavior policy (from unlikely to likely reading from left to right). Y-axis: the probabilities (blue, from small to large) and the corresponding value prediction errors (yellow, from small to large). Action sequences with higher probabilities to be taken by MuZero’s behavior policy correlate with lower value prediction errors by MuZero’s learned model. This implies that the learned model is less accurate for policies that are different from the current data collection policy.

value prediction errors for *all* action sequences of length  $\{8, 4, 4\}$  for Cart Pole, Lunar Lander and Breakout (limited by computation budget). Considering that both probabilities and errors are real-valued and non-uniformly distributed, we aggregate results as follows: first, for each state, we rank action sequences by their probabilities of being chosen by MuZero’s behavior policy. Then, we compile statistics for action sequences with the same ranks. Finally, we combine results from different agents and report their means and standard errors.


In Figure 4, we present the results. Here, the X-axis represents the action sequences that are sorted by their likelihood of being taken by the current behavior policy, ranging from unlikely to likely. On the Y-axis, we plot both the probabilities (blue) and the corresponding value prediction errors (yellow). The results clearly show that, as the likelihood of the behavior policy selecting the action sequence decreases, the value prediction error for that action sequence increases. This trend is consistently observed across different environments and training steps. Moreover, it seems to become more evident with more training, possibly because the behavior policy becomes more deterministic. This finding supports our hypothesis that the model is more reliable in predicting values for the behavior policy, which collects the training data. Consequently, evaluating a policy using MuZero’s learned model may yield increasingly inaccurate results as the policy to evaluate deviates further from the behavior policy.

### 3.5 How accurately can MuZero’s learned model from one training step evaluate the behavior policy of the same agent from another training step?

In our previous experiment, we investigated the generalization capability of the learned model in predicting values for policies that differ

from the behavior policy. We accomplished this by considering all action sequences of length  $h$  at each state, possibly including actions that are unlikely to be taken by any sensible policy. In this experiment, we conduct a similar analysis but with a focus on more interesting policies. Specifically, we assess the accuracy with which the model at training step  $X$  (row) can evaluate the policy (of the same agent) at training step  $Y$  (column). The idea is that if a model cannot accurately evaluate high-performing future policies, planning with it would not be effective in finding a good policy. For this experiment, we set the evaluation horizon for each environment to the number of unrolling steps during training. Results are aggregated over states sampled from MuZero’s on-policy state distribution at training step  $X$  (same per row as the model). Errors at different model steps (rows) are not directly comparable due to the different state distributions.

It is clear from Figure 5 that models at all training steps are most accurate when evaluating policies at the same training steps, as the error per row is smallest at the diagonal. This aligns with our earlier finding: learned models are more accurate when assessing action sequences with a higher selection likelihood by the current behavior policy. Notably, we can see that using early models to evaluate future policies results in large errors in Cart Pole and Lunar Lander. This may have important implications for policy improvement: *if the model at the current training step cannot accurately evaluate a future policy that performs better, then the extent to which we can improve our current policy will be limited.* Intuitively, if the model does not know a high-performing policy is good, then a deeper search in the model would not help us find the policy. This phenomenon is less evident in Breakout, possibly because the evaluation horizon is too short in this environment to allow for an observable difference in rewards across the policies (5 vs 10 in the other environments).



**Figure 5:** Cross model policy evaluation. We evaluate MuZero’s behavior policy at training step  $Y$  (column) with the learned model at training step  $X$  (row) and measure the value prediction error. Results are aggregated over states sampled from MuZero’s on-policy state distribution at training step  $X$  (same as the model).

## 4 Policy Improvement Experiments

Our policy evaluation experiments indicate that MuZero’s learned model can become increasingly inaccurate when evaluating policies differing from the data collection policy, particularly those unseen during training. In this section, we investigate the natural follow-up question, the question that we are most interested in: what is the effect of this on policy improvement through planning?

Intuitively, if the agent is given the ground-truth model and has an infinite budget for planning, it can act optimally by exhaustively searching with the model. In this case, sample efficiency is maximized because the agent does not need any sample from the environment to learn. However, in real-world scenarios, the planning budget is always constrained, both during training and deployment. To improve planning with these constraints, AlphaZero and MuZero employ a learned policy prior to guide the action simulation in MCTS. If the policy prior is well-trained, it can greatly enhance planning, but if not, it might diminish efficiency.

We evaluate planning using MuZero’s learned model both with and without the guidance of the policy prior. In the latter case, we replace the policy prior with a uniform prior, allowing for a form of “free search”. To focus on evaluating the contribution of the learned model, we modify MuZero’s MCTS by replacing the value network’s predictions at leaf nodes with random rollouts in the model, a standard approach to estimating the value of a leaf node in MCTS. To address the computational costs of long rollouts in neural models, we restrict the planning horizon to  $\{16, 128, 32\}$  in Cart Pole, Lunar Lander, and Atari Breakout. For comparison, we include the baselines of the policy prior and planning with the ground-truth model using the same MCTS. The goal of this experiment is to answer two questions: (a) How effectively does MuZero’s learned model support free search? (b) To what extent can MuZero improve its policy by planning with the learned model?


We present our results in Figure 7. In Cart Pole, MuZero’s learned model exhibits some degree of support for free search (orange dashed) after some training. However, it significantly underperforms compared to free search with the ground-truth model (green dashed) in terms of both planning efficiency and asymptotic performance. In Lunar Lander and Atari Breakout, free search with the learned model (orange dashed) fails completely. Consequently, the potential for finding a good policy through planning with the learned model will be effectively limited.

In Cart Pole and Lunar Lander, with enough MCTS simulations, planning with the learned model improves performance over using

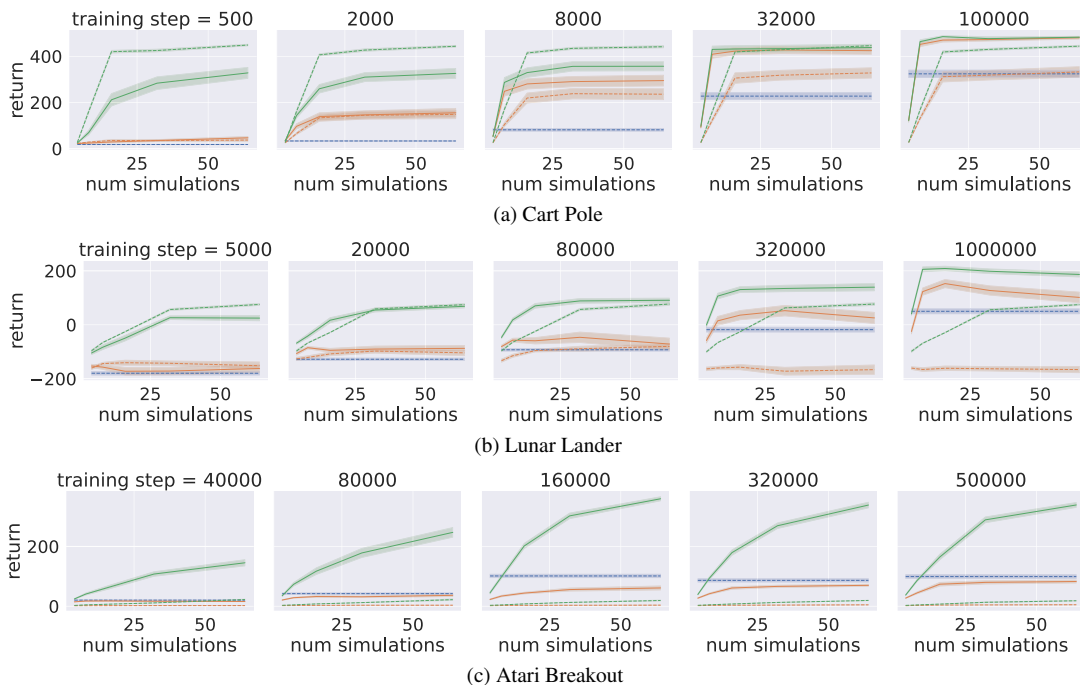
the policy prior alone. However, compared to the ground-truth model, it is evident that the extent to which we can improve upon the policy prior via planning is still very limited. This is a clear sign that the model error is restricting the extent to which we can further improve the policy by planning. In Atari Breakout, while planning with more simulations using the learned model improves the agent’s performance, it cannot outperform the policy prior.

Interestingly, when comparing planning with the policy prior to planning with the uniform prior, the learned model consistently shows a larger gap compared to the ground-truth model except in Atari Breakout. Furthermore, unlike the ground-truth model, the gap for the learned model does not appear to diminish quickly with more simulations, suggesting an additional role of the policy prior to accelerating MCTS. As shown in Figure 4, action sequences with higher probabilities of being selected by the behavior policy tend to be more accurately evaluated by the learned model. As the policy prior is directly learned to match the behavior policy, it is reasonable to assume that the values of action sequences favored by the policy prior can also be more accurately predicted by the learned model.

To verify this hypothesis, in Figure 6, we plot the learned model’s value prediction error for MCTS’s simulated trajectories under the guidance of the policy prior (blue) and uniform prior (orange) in CartPole after 2000 training steps. Meanwhile, we also plot the total variation  $\text{TV}[\pi_\theta, \hat{\pi}]$  and KL divergence  $\text{KL}[\pi_\theta, \hat{\pi}]$  between the policy prior  $\pi_\theta$  and MCTS’s empirical visit distribution  $\hat{\pi} = \frac{1+N(a)}{\sum_b |A|+N(b)}$  at the root node. Clearly, the policy prior regularizes MCTS to visit actions that are more favored by it, as suggested by the lower total variation and KL divergence (blue, middle, and right), which results in a smaller value prediction error (blue, left), when compared to the uniform prior, which explores more out-of-the-box (orange, middle, and right) and incurs more value prediction errors (orange, left). This suggests that *apart from biasing the search, the policy prior may also serve to prevent the search from exploring directions where the learned model is less accurate.*



**Figure 6:** Value prediction error of MCTS’s simulated trajectories using the learned model, total variation, and KL divergence (see text).



**Figure 7:** MCTS planning with (i) MuZero’s learned model (orange) and (ii) the ground-truth model (green), under the guidance of (a) MuZero’s policy prior (solid line) and (b) the uniform prior (dotted line)). The X and Y axes are the number of simulations per step and the return.

## 5 Discussion

**Limitations** In this study, we analyzed MuZero, one of the most successful DMBRL methods that is grounded in the value equivalence principle. The conclusions we draw can, therefore, not be directly generalized to other DMBRL methods, certainly not those that employ other auxiliary loss functions on top, e.g., [10, 50, 52]. However, our analysis clearly shows that, despite its simplicity, using the value equivalence principle does not imply that we learn models that are actually value equivalent. Even when restricting the predictions to its own behavior policy, MuZero’s learned model leads to prediction errors that quickly grow with the horizon of prediction. As such, our study does serve as a warning for other DMBRL methods that aim to use value equivalence as their guiding principle.

**Outlook** We have demonstrated that MuZero’s model, trained with a value-equivalence-based loss, struggles to generalize and predict values accurately for unseen or unfamiliar policies. This raises the question of how different losses for model learning would behave and compare in our analysis, such as the reconstruction-based loss [26, 18, 19, 20] and the temporal-consistency [50, 6, 52]. Specifically, we hypothesize that, while the value equivalence loss is the most flexible loss as it does not directly impose requirements on state representations, other losses may have an advantage for generalization in low-data regimes due to a richer supervision signal. We consider this potential trade-off between representation complexity and generalization an exciting direction for future research.

## 6 Conclusion

In this work, we empirically studied the models learned by MuZero, which are trained based on the value equivalence principle. Our analysis focused on two fundamental questions: (1) To what extent does MuZero learn a value-equivalent model? and (2) To what extent does the learned model support effective policy improvement by planning? We find that MuZero’s learned model cannot generally evaluate policies accurately, especially those that further deviate from the data collection policy. Consequently, the failure of the model to predict values for policies that are out of the training distribution prevents effective planning from scratch, which limits the extent to which MuZero can additionally improve its policy via planning. Moreover, we uncover that apart from accelerating the search as in AlphaZero, the policy prior in MuZero serves another crucial function: it regularizes the search towards areas where the learned model is more accurate, which effectively reduces the model error that is accumulated into planning. From these findings, we speculate that because MuZero’s model itself is limited in its ability for policy improvement, the empirical success of MuZero may be a result of the model providing the algorithm with a more powerful representation of the values and policies compared to single-lookahead methods like deep Q-learning [36]. This implies a role for the model that is similar to that in model-free planning but extends it with the policy prior to make planning conservative [28].

## Acknowledgements

The authors acknowledge the use of computational resources provided by the Delft AI Cluster (<https://daic.tudelft.nl>) and the Delft High Performance Computing Centre (<https://www.tudelft.nl/dhpc>).

## References

- [1] I. Antonoglou, J. Schrittwieser, S. Ozair, T. K. Hubert, and D. Silver. Planning in Stochastic Environments with a Learned Model. In *ICLR*, 2021.
- [2] M. G. Bellemare, J. Veness, and M. Bowling. The Arcade Learning Environment: An Evaluation Platform for General Agents. *JAIR*, 2013.
- [3] R. I. Brafman and M. Tennenholtz. R-MAX-A General Polynomial Time Algorithm for Near-Optimal Reinforcement Learning. *JMLR*, 2002.
- [4] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton. A survey of Monte Carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games*, 2012.
- [5] I. Danihelka, A. Guez, J. Schrittwieser, and D. Silver. Policy improvement by planning with Gumbel. In *ICLR*, 2021.
- [6] J. de Vries, K. Voskuil, T. M. Moerland, and A. Plaat. Visualizing MuZero Models. In *ICML 2021 Workshop on Unsupervised Reinforcement Learning*, 2021.
- [7] J. Degraeve, F. Felici, J. Buchli, M. Neunert, B. Tracey, F. Carpanese, T. Ewalds, R. Hafner, A. Abdolmaleki, D. de las Casas, C. Donner, L. Fritz, C. Galperti, A. Huber, J. Keeling, M. Tsimpanoukelli, J. Kay, A. Merle, and J.-M. Moret, et al. Magnetic control of tokamak plasmas through deep reinforcement learning. *Nature*, 2022.
- [8] G. Farquhar, T. Rocktäschel, M. Igl, and S. Whiteson. TreeQN and ATreeC: Differentiable Tree-Structured Models for Deep Reinforcement Learning. In *ICLR*, 2018.
- [9] V. François-Lavet, P. Henderson, R. Islam, M. G. Bellemare, and J. Pineau. An Introduction to Deep Reinforcement Learning. *Foundations and Trends® in Machine Learning*, 2018.
- [10] C. Gelada, S. Kumar, J. Buckman, O. Nachum, and M. G. Bellemare. DeepMDP: Learning Continuous Latent Space Models for Representation Learning. In *ICML*, 2019.
- [11] J.-B. Grill, F. Althé, Y. Tang, T. Hubert, M. Valko, I. Antonoglou, and R. Munos. Monte-Carlo Tree Search as Regularized Policy Optimization. In *ICML*, 2020.
- [12] C. Grimm, A. Barreto, S. Singh, and D. Silver. The Value Equivalence Principle for Model-Based Reinforcement Learning. In *NeurIPS*, 2020.
- [13] C. Grimm, A. Barreto, G. Farquhar, D. Silver, and S. Singh. Proper value equivalence. In *NeurIPS*, 2021.
- [14] C. Grimm, A. Barreto, and S. Singh. Approximate Value Equivalence. In *NeurIPS*, volume 35, 2022.
- [15] A. Guez, T. Weber, I. Antonoglou, K. Simonyan, O. Vinyals, D. Wierstra, R. Munos, and D. Silver. Learning to search with MCTSnets. In *ICML*, 2018.
- [16] A. Guez, M. Mirza, K. Gregor, R. Kabra, S. Racaniere, T. Weber, D. Rapposo, A. Santoro, L. Orseau, T. Eccles, G. Wayne, D. Silver, and T. Lillicrap. An Investigation of Model-Free Planning. In *ICML*, 2019.
- [17] D. Ha and J. Schmidhuber. Recurrent World Models Facilitate Policy Evolution. In *NeurIPS*, 2018.
- [18] D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi. Dream to Control: Learning Behaviors by Latent Imagination. In *ICLR*, 2019.
- [19] D. Hafner, T. P. Lillicrap, M. Norouzi, and J. Ba. Mastering Atari with Discrete World Models. In *ICLR*, 2020.
- [20] D. Hafner, J. Pasukonis, J. Ba, and T. Lillicrap. Mastering Diverse Domains through World Models. *arXiv preprint arXiv:2301.04104*, 2023.
- [21] J. B. Hamrick, A. L. Friesen, F. Behbahani, A. Guez, F. Viola, S. Witherspoon, T. Anthony, L. H. Buesing, P. Veličković, and T. Weber. On the role of planning in model-based deep reinforcement learning. In *ICLR*, 2022.
- [22] M. Henaff. Explicit Explore-Exploit Algorithms in Continuous State Spaces. In *NeurIPS*, 2019.
- [23] M. Hessel, I. Danihelka, F. Viola, A. Guez, S. Schmitt, L. Sifre, T. Weber, D. Silver, and H. van Hasselt. Muesli: Combining Improvements in Policy Optimization. *arXiv preprint arXiv:2104.06159*, 2022.
- [24] T. Hubert, J. Schrittwieser, I. Antonoglou, M. Barekatin, S. Schmitt, and D. Silver. Learning and Planning in Complex Action Spaces. In *ICML*, 2021.
- [25] M. Jaderberg, V. Mnih, W. M. Czarnecki, T. Schaul, J. Z. Leibo, D. Silver, and K. Kavukcuoglu. Reinforcement Learning with Unsupervised Auxiliary Tasks. In *ICLR*, 2022.
- [26] Ł. Kaiser, M. Babaeizadeh, P. Miłoś, B. Osiniński, R. H. Campbell, K. Czechowski, D. Erhan, C. Finn, P. Kozakowski, S. Levine, A. Mohiuddin, R. Sepassi, G. Tucker, and H. Michalewski. Model Based Reinforcement Learning for Atari. In *ICLR*, 2020.
- [27] L. Kocsis and C. Szepesvári. Bandit based Monte-Carlo planning. In *ECML*, 2006.
- [28] A. Kumar, A. Zhou, G. Tucker, and S. Levine. Conservative Q-Learning for Offline Reinforcement Learning. In *NeurIPS*, 2020.
- [29] N. Lambert, K. Pister, and R. Calandra. Investigating Compounding Prediction Errors in Learned Dynamics Models. *arXiv preprint arXiv:2203.09637*, 2022.
- [30] L. Li, T. J. Walsh, and M. L. Littman. Towards a Unified Theory of State Abstraction for MDPs. In *ISAAC*, 2006.
- [31] K. Lowrey, A. Rajeswaran, S. Kakade, E. Todorov, and I. Mordatch. Plan Online, Learn Offline: Efficient Learning and Exploration via Model-Based Control. In *ICLR*, 2018.
- [32] D. Madeka, K. Torkkola, C. Eisenach, A. Luo, D. P. Foster, and S. M. Kakade. Deep Inventory Management. *arXiv preprint arXiv:2210.03137*, 2022.
- [33] A. Mandhane, A. Zhernov, M. Rauh, C. Gu, M. Wang, F. Xue, W. Shang, D. Pang, R. Claus, C.-H. Chiang, C. Chen, J. Han, A. Chen, D. J. Mankowitz, J. Broshear, J. Schrittwieser, T. Hubert, O. Vinyals, and T. Mann. MuZero with Self-competition for Rate Control in VP9 Video Compression. *arXiv preprint arXiv:2202.06626*, 2022.
- [34] A. K. McCallum. *Reinforcement Learning with Selective Perception and Hidden State*. PhD thesis, University of Rochester, 1996.
- [35] A. Mirhoseini, A. Goldie, M. Yazgan, J. W. Jiang, E. Songhori, S. Wang, Y.-J. Lee, E. Johnson, O. Pathak, A. Nazi, J. Pak, A. Tong, K. Srinivasa, W. Hang, E. Tuncer, Q. V. Le, J. Laudon, R. Ho, and R. Carpenter, et al. A graph placement methodology for fast chip design. *Nature*, 2021.
- [36] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis. Human-level control through deep reinforcement learning. *Nature*, 2015.
- [37] T. M. Moerland, J. Broekens, A. Plaat, and C. M. Jonker. Model-based Reinforcement Learning: A Survey. *Foundations and Trends® in Machine Learning*, 2023.
- [38] J. Oh, S. Singh, and H. Lee. Value Prediction Network. In *NeurIPS*, 2017.
- [39] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell. Curiosity-driven Exploration by Self-supervised Prediction. *arXiv preprint arXiv:1705.05363*, 2017.
- [40] D. Pathak, D. Gandhi, and A. Gupta. Self-Supervised Exploration via Disagreement. In *ICML*, 2019.
- [41] J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, and T. Graepel. Mastering atari, go, chess and shogi by planning with a learned model. *Nature*, 2020.
- [42] R. Sekar, O. Rybkin, K. Daniilidis, P. Abbeel, D. Hafner, and D. Pathak. Planning to Explore via Self-Supervised World Models. In *ICML*, 2020.
- [43] P. Shyam, W. Jaśkowski, and F. Gomez. Model-Based Active Exploration. In *ICML*, 2019.
- [44] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, and T. Graepel, et al. Mastering the game of Go with deep neural networks and tree search. *Nature*, 2016.
- [45] D. Silver, H. Hasselt, M. Hessel, T. Schaul, A. Guez, T. Harley, G. Dulac-Arnold, D. Reichert, N. Rabinowitz, A. Barreto, and T. Degris. The Predictron: End-To-End Learning and Planning. In *ICML*, 2017.
- [46] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, L. Sifre, G. Van Den Driessche, T. Graepel, and D. Hassabis. Mastering the game of Go without human knowledge. *Nature*, 2017.
- [47] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, T. Lillicrap, K. Simonyan, and D. Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 2018.
- [48] R. S. Sutton. Dyna, an integrated architecture for learning, planning, and reacting. *ACM SIGART Bulletin*, 1991.
- [49] A. Tamar, YI. WU, G. Thomas, S. Levine, and P. Abbeel. Value Iteration Networks. In *NeurIPS*, 2016.
- [50] E. van der Pol, T. Kipf, F. A. Oliehoek, and M. Welling. Plannable Approximations to MDP Homomorphisms: Equivariance under Actions. In *AAMAS*, 2020.
- [51] H. P. van Hasselt, M. Hessel, and J. Aslanides. When to use parametric models in reinforcement learning? In *NeurIPS*, 2019.
- [52] W. Ye, S. Liu, T. Kurutach, P. Abbeel, and Y. Gao. Mastering Atari Games with Limited Data. In *NeurIPS*, 2021.