
Bridging Worlds in Reinforcement Learning with Model-Advantage

Nirbhay Modhe^{*1} Harish Kamath^{*1} Dhruv Batra^{1,2} Ashwin Kalyan¹

Abstract

Despite the breakthroughs achieved by Reinforcement Learning (RL) in recent years, RL agents often fail to perform well in unseen environments. This inability to *generalize* to new environments prevents their deployment in the real world. To help measure this gap in performance, we introduce *model-advantage* - a quantity similar to the well-known (policy) advantage function. First, we show relationships between the proposed model-advantage and generalization in RL — using which we provide guarantees on the gap in performance of an agent in new environments. Further, we conduct toy experiments to show that even a sub-optimal policy (learnt with minimal interactions with the target environment) can help *predict* if a training environment (say, a simulator) helps learn policies that generalize. We then show connections with Model Based RL.

1. Introduction

Reinforcement Learning (RL) has emerged as a promising learning paradigm owing to its successes in applications like strategy games (Mnih et al., 2015; Silver et al., 2016; 2017) and robotics (Levine et al., 2016; Gu et al., 2016). Due to the high cost of interacting with the real world – e.g. accidents in the case of physical robots or loss of revenue in recommendation systems – it is common to train RL agents in a more accessible environment. For example, if we want to train a self-driving agent to drive in Maine, a good starting point is to train using a simulator or in a different environment which is more accessible such as Nevada. However, an agent with good performance in this training environment may or may not achieve similar performance in the test environment (the *real world*). This *lack of generalization* of RL agents prevents its reliable deployment in the real world.

^{*}Equal contribution ¹Georgia Tech ²Facebook AI Research. Correspondence to: Nirbhay Modhe <nirbhaym@gatech.edu>, Harish Kamath <hkamath@gatech.edu>.

Intuitively, better performance can be expected when the training environment informs the agent of the trajectories it is likely to take in the target environment. Ideally, we would like that in both the training and target MDP, taking actions at states lead to the same transitions as well as rewards. To formalize this notion, we introduce *model-advantage* – a quantity similar to the well-known advantage function in RL. The standard advantage function – which we refer to as policy-advantage – evaluates the *advantage* of playing a particular action as opposed to the action of a reference policy. Similarly, we define model-advantage as the advantage of transitioning to a state as opposed to transitioning according to an MDP M , while acting according to some policy. Specifically, we are interested in the expected advantage of transitioning according to one MDP with respect to another one as reference, which allows us to evaluate the effectiveness of using one model in lieu of the other – much like how policy-advantage helps compare two different policies.

2. Model-Advantage

In this section, we formally introduce *model-advantage* and definitions associated with it. Recall that $\mathbb{A}^\pi(s, a)$ i.e. policy-advantage defined over states $s \in \mathcal{S}$ and actions $a \in \mathcal{A}$ measures the *utility* of taking action a as opposed to acting according to policy π . However, unlike policy-advantage that measures the difference in utility of taking an action, we are interested in knowing the utility difference of transitioning to a particular state, while following a single policy. Specifically, *model-advantage* denoted by $\mathbf{A}_M^\pi(s, s')$ compares the utility of moving to state s' and following the trajectory governed by model M as opposed to doing it from state s ; both under policy π .

We define the model-dependent value function as follows, where we make explicit the MDP M (and hence, the transition function \mathcal{P}_M) used to generate trajectories.

$$V_M^\pi(s) = \mathbb{E}_{\rho_M^\pi} \left[\sum_{t=0}^{\infty} \gamma^t \mathcal{R}_M(s_t, a_t) \mid \pi, M, s_0 = s \right]$$

Here, ρ_π is the distribution of trajectories (s_0, a_0, s_1, \dots) , $s_0 \sim \mathcal{P}_0$, when acting according to policy π . Now, the intuition of model-advantage ¹ is given by:

¹A $Q(s, s')$ function can also be defined; however, it requires

$$\mathbf{A}_M^\pi(s, s') = \gamma [V_M^\pi(s') - \mathbb{E}_{s'' \sim P_M(s, \pi)} V_M^\pi(s'')] \quad (1)$$

Analogous to policy-advantage that compares different policies in the same environment, model-advantage helps compare the same policy acting in two different environments. For such a comparison, we need to look at the quantity $\mathbb{E}_{s' \sim M'} [\mathbf{A}_M^\pi(s, s')] -$ the expected model-advantage (evaluated at π) when the next state s' is obtained from the MDP M' . We formalize this by the following (*model performance difference*) lemma. The proof resembles the proof by (Kakade & Langford, 2002) and is provided in Appendix A.

Lemma 2.1. (*Model Performance Difference Lemma*) Let M and M' be two different MDPs. Further, define $\mathcal{R}^\pi(s) = \mathbb{E}_{a \sim \pi(\cdot|s)} [\mathcal{R}(s, a)]$, $\mathcal{R}_\epsilon^\pi(s) = \mathcal{R}_M^\pi(s) - \mathcal{R}_{M'}^\pi(s)$ and $J(\pi) = \mathbb{E}_{s \sim \mathcal{P}_0} [V^\pi(s)]$. For any policy $\pi \in \Pi$ we have:

$$J_M(\pi) - J_{M'}(\pi) = \mathbb{E}_{s \sim d_{M, \pi}} [\mathcal{R}_\epsilon(s)] + \frac{1}{1-\gamma} \mathbb{E}_{s \sim d_{M, \pi}} \mathbb{E}_{s' \sim \mathcal{P}_M(s'|s, \pi)} [\mathbf{A}_{M'}^\pi(s, s')] \quad (2)$$

Here, we use a model-dependent stationary state distribution $d_{M, \pi}(s)$ where the dynamics \mathcal{P}_M are used, assuming a start state distribution \mathcal{P}_0 . Compared to its policy counterpart the (model) performance difference lemma involves an additional reward error term $\mathbb{E}_{d_{M, \pi}} \mathcal{R}_\epsilon$ that vanishes when the two MDPs differ only in the transition probabilities.

We can also use the Bellman evaluation operator for policy π to obtain an equivalent formalization (see Appendix A):

$$J_M(\pi) = J_{M'}(\pi) + \frac{1}{1-\gamma} \mathbb{E}_{d_{M, \pi}} \underbrace{[\mathcal{T}_M^\pi V_M^\pi - \mathcal{T}_{M'}^\pi V_M^\pi]}_{\text{deviation error}} \quad (3)$$

The term $\mathcal{T}_M^\pi V_M^\pi - \mathcal{T}_{M'}^\pi V_M^\pi$, which we denote by $\delta_{M, M'}^\pi$ ² represents the *deviation* in the value function when acted upon by Bellman operators corresponding to two different MDPs. This term is exactly equal to model-advantage when the reward functions of the two MDPs are the same. We can form an upper bound on the extrinsic error as:

$$J_M(\pi) - J_{M'}(\pi) \leq \frac{1}{1-\gamma} \|\mathcal{T}_M^\pi V^\pi - \mathcal{T}_{M'}^\pi V^\pi\|_\infty \leq \frac{1}{1-\gamma} [\epsilon_R + \gamma \epsilon_P \|V^\pi\|_\infty] \quad (4)$$

where the rewards and the dynamics themselves are individually bounded *i.e.* $\max_s \max_a |R_M(s, a) - R_{M'}(s, a)| \leq \epsilon_R$

additional formalism not necessary for our exposition. See concurrent work (Edwards et al., 2020) for a detailed discussion.

²Optionally, when using the optimality operators corresponding to MDPs M and M' , we drop the π and denote the deviation error as $\delta_{M, M'}(V(s))$.

and $\max_s \max_a \|P_M(s, a) - P_{M'}(s, a)\|_1 \leq \epsilon_P$. Of course, note that $\|V^\pi\|_\infty$ is trivially bounded by $\frac{1}{1-\gamma} R_{\max}$, assuming rewards are bounded by R_{\max} .

3. Generalization in RL

An RL problem is characterized by an MDP M and is considered solved when a policy $\pi \in \Pi$ that maximizes the expected (discounted) return is found. However, in practice, the RL agent may then be deployed in a slightly different environment characterized by an MDP M' . Therefore, we are interested in how well an RL agent performs in an unseen environment – in other words, its ability to *generalize*.

Related Work. Studying and benchmarking generalization properties of RL agents via large-scale experiments has been the focus of many works in the recent years (Zhang et al., 2018a;b; Cobbe et al., 2018). Additionally, (Slaoui et al., 2019) derives a lemma comparable to lemma 2.1 for policies that are Lipschitz continuous over a set of state-representations. Importantly, (Wang et al., 2019a) formally define generalization gap which we adopt in this work (see eq. (5)) They give formal bounds on this gap in the setting of *reparametrizable RL* while making additional assumptions like Lipschitz continuity on value functions. While we do not require any such assumptions, it is important to note that it is not possible to guarantee tighter bounds without them.

3.1. Generalization Gap.

Let $J(\pi) := \mathbb{E}[\sum_t \mathcal{R}(s_t, a_t)]$ denote the cost function, where the stochasticity is due to the policy and transition dynamics; let $\hat{J}_n(\pi)$ denote its empirical estimate with n samples. Given an RL agent trained in MDP M with finite data, we are interested in its performance in a different MDP M' . We can formally write this *generalization gap* as:

$$\Phi = \underbrace{|\hat{J}_n(\pi) - J(\pi)|}_{\text{intrinsic error}} + \underbrace{|J(\pi) - J'(\pi)|}_{\text{extrinsic error}} \quad (5)$$

The generalization gap can be bounded with two different sources of error as indicated in eq. (5). Following (Wang et al., 2019a), we call them *intrinsic* error and *extrinsic* error to denote the error due to learning from finite samples and the error due to mismatch in training and deployment environments. The intrinsic error decreases, typically as $\mathcal{O}(1/\sqrt{n})$, with more samples; this is well-studied in RL literature (Kakade et al., 2003; Azar et al., 2012; Jin et al., 2018). The extrinsic error on the other hand is an artifact of training and deploying the RL agent in different environments and therefore, cannot be avoided.

When does π generalize? Observe that the extrinsic error is nothing but the difference in performance due to

model mismatch. From lemma 2.1, we know that this is equal to the model-advantage, allowing us to both estimate and bound this error term. In other words, if the model-advantage is bounded by ϵ (see eq. (4)) *i.e.*

$$|J_M(\pi_M) - J_{M'}(\pi_M)| \leq \epsilon$$

we can say that π_M , the policy learnt with experiences from MDP M achieves similar performance in the target MDP M' . As model of the “test” environment is not known, a reasonable estimate of the model-advantage can be obtained with enough samples – allowing one to *predict* the extent to which the policy performs in the novel environment.

How good is π really? However, note that the above generalization gap still does not provide the complete picture. Ideally, we would like the policy π_M to have performance comparable to $\pi_{M'}^*$, the optimal policy in the target MDP M' *i.e.*

$$\begin{aligned} & |J_{M'}(\pi_M) - J_{M'}(\pi_{M'}^*)| \\ & \leq \underbrace{|J_{M'}(\pi_M) - J_M(\pi_M)|}_{\text{term-I}} + \underbrace{|J_M(\pi_M) - J_{M'}(\pi_{M'}^*)|}_{\text{term-II}} \end{aligned} \quad (6)$$

It is easy to see that term-I is nothing but the extrinsic error in eq. (5) and is related to the model-advantage (evaluated under policy π_M) through lemma 2.1. Intuitively, this term corresponds to the cost of transferring π_M learnt in the seen MDP M to the novel MDP M' .

In the rest of this section, we will bound term-II for specific instantiations of obtaining policy π_M – specifically, Value Iteration (VI) and Fitted Q-Iteration (FQI), with the former being a model-based and the latter, a model-free approach to solve MDPs.

3.2. Generalization with Value Iteration and Fitted Q iteration

Value Iteration. When the dynamics and the reward functions are known, Value Iteration (VI) and its variants are often employed to arrive at the optimal policy. VI is an iterative algorithm that applies the Bellman optimality operator \mathcal{T}_M ³ at each step *i.e.* $V^{(n)} = \mathcal{T}_M V^{(n-1)}$. The obtained iterates converge to V_M^* , the value function of π_M^* , asymptotically as \mathcal{T}_M is a contraction in the infinity-norm. We can bound the difference in value from training on another MDP with the following theorem:

Theorem 3.1. *Let M, M' be two MDPs s.t. $\max_s \max_a |\mathcal{R}_M(s, a) - \mathcal{R}_{M'}(s, a)| \leq \epsilon_R$ and $\max_s \max_a \|p_M(s, a) - p_{M'}(s, a)\|_1 \leq \epsilon_P$. Let π_{n+1} be the policy obtained after n VI iterations on MDP M and let*

³ Assume optimality operator by default if policy is not explicitly defined

$\|V_M^{(n+1)} - V_M^{(n)}\|_\infty \leq \epsilon^{(n)}$ Then we have,

$$\|V_{M'}^{\pi_{n+1}} - V_{M'}^*\|_\infty \leq \frac{1}{1-\gamma} \left[\gamma \epsilon^{(n)} + 2\epsilon_R + \frac{2\epsilon_P R_{max}}{1-\gamma} \right]$$

We provide a similar bound when following a Fitted Q iteration (FQI) method, which is more effective when dealing with a large (or infinite) state space or unknown dynamics/reward functions. The statement and proof of the bound can be found in Appendix B.2.

3.3. Sim2Real: Is My Simulator Good?

The fundamental bottleneck preventing the usage of RL to train agents in the real-world is *exploration*. As the model of the environment is not available, finding the optimal policy not only requires exploring a large search space but is also *costly*. A common strategy to avoid this issue is to learn a *coarse* policy using a simulator and then *fine-tune* it upon deployment. But how does one build the simulator in the first place? It either requires considerable domain expertise or a large number of samples from the real-world, and we must know a priori that the simulator can *express* all variations feasible in the real world. We are left with the question: *Given a set of simulators, which one is likely to “generalize” best to the real-world?*

Predicting Generalization with Model-Advantage. Recall that (model) performance difference lemma 2.1 allows us to compare two models given a policy. Given M , the simulator MDP and M' , the real-world MDP and π_{exp} , an expert policy for M' , we can write:

$$|J_M(\pi_{exp}) - J_{M'}(\pi_{exp})| = \left| \mathbb{E}_{(s,s') \sim M} [A_{M'}^{\pi_{exp}}(s, s')] \right|$$

To compute the advantage function $A_{M'}^{\pi_{exp}}$, the expert policy has to be executed in the real-world. Alternately, such an “expert” policy and its corresponding value function in MDP M' can be learnt by collecting a finite set of data from the real-world – for instance, by running FQI on the collected dataset⁴. After paying this one-time cost of interacting with the real-world, the model-advantage can be estimated in an inexpensive manner for every simulator with finite samples. In our experiments we will show that using an approximately optimal policy is sufficient for comparing model advantage across simulators, alleviating the need for an expert policy.

Grid World Experiments. We consider the toy environment of FrozenLake available as part of OpenAI Gym⁵ to illustrate the effectiveness of the proposed model-advantage

⁴Note that convergence to optimal value-function is guaranteed only if the distribution used to sample states is exploratory. While the “tax” of exploration cannot be waived, the hope here is that a small amount of data is sufficient to learn a sub-optimal expert.

⁵<https://gym.openai.com/>

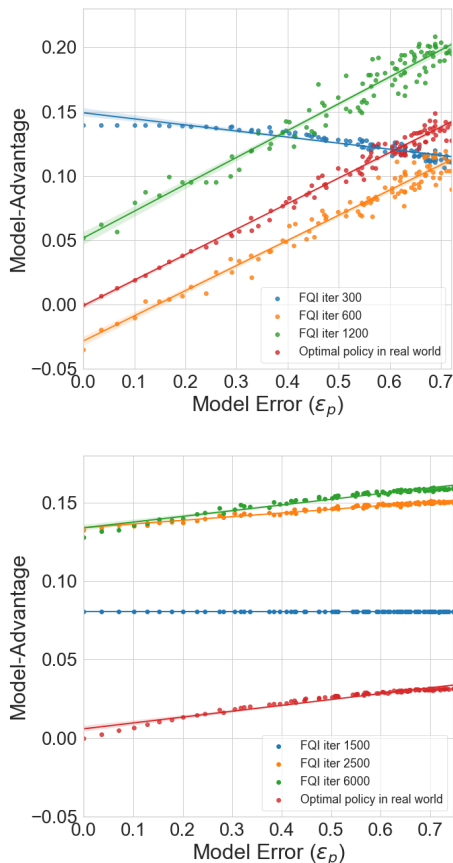


Figure 1. Generalization gap on sub-optimal real world policies on FrozenLake environments (FrozenLake 4x4 on the left and 8x8 on the right). Even a sub-optimal policy obtained with minimal interactions with the real-world is sufficient to use model-advantage and compare different *training* environments or simulators.

term in evaluating simulators. We treat the original setting as M , the real-world MDP and then, corrupt the transition dynamics with various levels of random noise to obtain a set of “simulators” $\{M'_i\}_{i=1}^K$. We then run DQN (Mnih et al., 2013) that uses a single hidden-layer MLP to learn Q-values in M and obtain a sub-optimal “expert” Q-function by not running the training to completion. As can be seen from Figure 1, we see that even for Q-values far from optimal, the model-advantage increases with increasing modeling error ϵ_P – the same trend exhibited by the optimal Q-function in the real-world.

3.4. Connections to Model-Based RL

Model Based Reinforcement Learning (MBRL) is a dominant RL framework that first learns an *approximate* model of the real-world, and then runs standard planning algorithms like VI to find a suitable policy. Variants of the model performance difference lemma (see lemma 2.1) involving the approximated model have been derived in this context – for

e.g. see (Kearns & Singh, 2002; Kakade et al., 2003) and concurrent works, (Rajeswaran et al., 2020; Kidambi et al., 2020; Xu et al., 2018). In this section, we introduce a novel perspective of MBRL and MBRL techniques.

The objective of MBRL is to construct a model that mimics the dynamics and reward functions of the real-world as accurately as possible. As seen from eq. (4), the model-advantage is upper bounded as a function of ϵ_R and ϵ_P , the error in reward and transition dynamics functions. Therefore, many previous works minimize this upper bound (Azar et al., 2012; 2013; 2017). The drawback is that when the upper bound is fairly loose, it leads to a “dynamics bottleneck” (Wang et al., 2019b; Lambert et al., 2020).

Directly Minimizing Model-Advantage. From lemma 2.1 and eq. (4), a tighter bound on the model-advantage can be minimized (Wang et al., 2019b; Lambert et al., 2020) –

$$\begin{aligned} |J_M(\pi) - J_{M'}(\pi)| &= \left| \mathbb{E}_{(s,s') \sim M} [A_{M'}^\pi(s, s')] \right| \\ &\leq \frac{\epsilon_R}{1 - \gamma} + \frac{\gamma \epsilon_P R_{\max}}{1 - \gamma} \end{aligned}$$

This is explored by the Value-Aware Model Learning framework (Farahmand, 2018a;b); specifically, they minimize the $\mathbb{E} \left| \delta_{(M, M')}^\pi \right|^2$ for a worst-case policy that seeks to maximize the deviation. The crux of MBRL algorithms is to assume access to the real-world – either through trajectories collected by running a policy or with a generative model of the MDP. These samples correspond to *expert* dynamics (or rewards) and can be learnt by using the Imitation Learning (IL) framework. This naturally leads to an algorithm that alternately updates a model to minimize the dynamics and reward error with imitation learning, and updates a policy to maximize reward and further collect samples from the true MDP. We defer the reader to Appendix C for guarantees of such an algorithm with MBRL and IL.

4. Conclusion

In this work, we proposed *model-advantage* that helps compare two models, similar to policy advantage that can be used to compare two policies. This term provides a theoretical framework for understanding generalization in RL – specifically, we provide formal guarantees on the generalization ability of policies learnt via Value Iteration (VI) and Fitted Q-Iteration (FQI). Further, we conduct toy experiments to show that even a sub-optimal policy, learnt from minimal interactions with the target environment, can help identify the training environment that facilitates maximum generalization. Finally, we discuss connections between model-advantage and Model Based Reinforcement Learning (MBRL), and formally establish connections between MBRL and Imitation Learning using the proposed model-advantage terms.

References

- Azar, M. G., Munos, R., and Kappen, B. On the sample complexity of reinforcement learning with a generative model. *arXiv preprint arXiv:1206.6461*, 2012.
- Azar, M. G., Munos, R., and Kappen, H. J. Minimax pac bounds on the sample complexity of reinforcement learning with a generative model. *Machine learning*, 91(3):325–349, 2013.
- Azar, M. G., Osband, I., and Munos, R. Minimax regret bounds for reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 263–272. JMLR. org, 2017.
- Cobbe, K., Klimov, O., Hesse, C., Kim, T., and Schulman, J. Quantifying generalization in reinforcement learning. *arXiv preprint arXiv:1812.02341*, 2018.
- Edwards, A. D., Sahni, H., Liu, R., Hung, J., Jain, A., Wang, R., Ecoffet, A., Miconi, T., Isbell, C., and Yosinski, J. Estimating $q(s, s')$ with deep deterministic dynamics gradients. *arXiv preprint arXiv:2002.09505*, 2020.
- Farahmand, A.-m. Iterative value-aware model learning. In *Advances in Neural Information Processing Systems*, pp. 9072–9083, 2018a.
- Farahmand, A.-m. Iterative value-aware model learning. In *Advances in Neural Information Processing Systems*, pp. 9072–9083, 2018b.
- Gu, S., Lillicrap, T., Sutskever, I., and Levine, S. Continuous deep q-learning with model-based acceleration. In *International Conference on Machine Learning*, pp. 2829–2838, 2016.
- Jin, C., Allen-Zhu, Z., Bubeck, S., and Jordan, M. I. Is q-learning provably efficient? In *Advances in Neural Information Processing Systems*, pp. 4863–4873, 2018.
- Kakade, S. and Langford, J. Approximately optimal approximate reinforcement learning. In *ICML*, volume 2, pp. 267–274, 2002.
- Kakade, S., Kearns, M. J., and Langford, J. Exploration in metric state spaces. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pp. 306–312, 2003.
- Kearns, M. and Singh, S. Near-optimal reinforcement learning in polynomial time. *Machine learning*, 49(2-3):209–232, 2002.
- Kidambi, R., Rajeswaran, A., Netrapalli, P., and Joachims, T. Morel: Model-based offline reinforcement learning. *arXiv preprint arXiv:2005.05951*, 2020.
- Lambert, N., Amos, B., Yadan, O., and Calandra, R. Objective mismatch in model-based reinforcement learning. *arXiv preprint arXiv:2002.04523*, 2020.
- Levine, S., Finn, C., Darrell, T., and Abbeel, P. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., and Riedmiller, M. Playing atari with deep reinforcement learning. December 2013. URL <http://arxiv.org/abs/1312.5602>.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.
- Munos, R. and Szepesvári, C. Finite-time bounds for fitted value iteration. *Journal of Machine Learning Research*, 9(May):815–857, 2008.
- Rajeswaran, A., Mordatch, I., and Kumar, V. A game theoretic framework for model based reinforcement learning. *arXiv preprint arXiv:2004.07804*, 2020.
- Ross, S. and Bagnell, J. A. Agnostic system identification for model-based reinforcement learning. *arXiv preprint arXiv:1203.1007*, 2012.
- Ross, S., Gordon, G., and Bagnell, D. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 627–635, 2011.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneer-shelvam, V., Lanctot, M., et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484–489, 2016.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017.
- Slaoui, R. B., Clements, W. R., Foerster, J. N., and Toth, S. Robust domain randomization for reinforcement learning. *arXiv preprint arXiv:1910.10537*, 2019.
- Wang, H., Zheng, S., Xiong, C., and Socher, R. On the generalization gap in reparameterizable reinforcement learning. In *International Conference on Machine Learning*, pp. 6648–6658, 2019a.
- Wang, T., Bao, X., Clavera, I., Hoang, J., Wen, Y., Langlois, E., Zhang, S., Zhang, G., Abbeel, P., and Ba, J. Benchmarking model-based reinforcement learning. *arXiv preprint arXiv:1907.02057*, 2019b.
- Xu, H., Li, Y., Tian, Y., Darrell, T., and Ma, T. Algorithmic framework for model-based reinforcement learning with theoretical guarantees. *arXiv preprint arXiv:1807.03858*, 2018.
- Zhang, A., Ballas, N., and Pineau, J. A dissection of overfitting and generalization in continuous reinforcement learning. *arXiv preprint arXiv:1806.07937*, 2018a.
- Zhang, A., Wu, Y., and Pineau, J. Natural environment benchmarks for reinforcement learning. *arXiv preprint arXiv:1811.06032*, 2018b.