
Augmented World Models Facilitate Zero-Shot Dynamics Generalization From a Single Offline Environment

Philip J. Ball^{*1} Cong Lu^{*1} Jack Parker-Holder¹ Stephen Roberts¹

Abstract

Reinforcement learning from large-scale offline datasets provides us with the ability to learn policies without potentially unsafe or impractical exploration. Significant progress has been made in the past few years in dealing with the challenge of correcting for differing behavior between the data collection and learned policies. However, little attention has been paid to potentially changing dynamics when transferring a policy to the online setting, where performance can be up to 90% reduced for existing methods. In this paper we address this problem with *Augmented World Models* (AugWM). We augment a learned dynamics model with simple transformations that seek to capture potential changes in physical properties of the robot, leading to more robust policies. We not only train our policy in this new setting, but also provide it with the sampled augmentation as a context, allowing it to adapt to changes in the environment. At test time we *learn the context* in a self-supervised fashion by approximating the augmentation which corresponds to the new environment. We rigorously evaluate our approach on over 100 different changed dynamics settings, and show that this simple approach can significantly improve the zero-shot generalization of a recent state-of-the-art baseline, often achieving successful policies where the baseline fails.

1. Introduction

Offline reinforcement learning (RL) describes the problem setting where RL agents learn policies solely from previously collected experience without further interaction with the environment (Fujimoto et al., 2019; Levine et al., 2020). This could have tremendous implications for real world

problems (Dulac-Arnold et al., 2019), with the potential to leverage rich datasets of past experience where exploration is either not feasible (e.g. a Mars Rover) or unsafe (e.g. in medical settings). As such, interest in offline RL has surged in recent times.

This work focuses on model-based offline RL, which has achieved state-of-the-art performance through the use of uncertainty penalized updates (Yu et al., 2020; Kidambi et al., 2020). However, existing work only addresses the issue of transferring from different behavior policies in the *same environment*, ignoring any possibility of distribution shift. Consider the case where it is expensive to collect data, and we have access to a single dataset from a robot. Using existing methods we would be unable to make any changes that impact the dynamics, such as using a newer model of the robot or deploying it in a different room.

A related setting is the Sim2Real problem which considers transferring an agent from a simulated environment to the real world. A popular recent approach is domain randomization (Tobin et al., 2017; James et al., 2017), the process of randomizing non-essential regions of the observation space to make agents robust to ‘observational overfitting’ (Song et al., 2020). Indeed, methods seeking to generalize to novel dynamics have also shown promise (Peng et al., 2018), by randomizing physical properties such as the mass of the agent. A significant limitation of these approaches is the requirement for a simulator, which may not be available.

In this work we take inspiration from Sim2Real to generalize solely from an offline dataset, in a learned simulator or *World Model* (WM). We therefore describe our problem setting as follows: an agent must learn to generalize to unseen test-time dynamics whilst having access to offline data from *only a single environment*; we call this “dynamics generalization from a single offline environment”.

In this paper we concentrate on the zero-shot performance of our agents to unseen dynamics, as it may not be practical (nor safe) to perform multiple rollouts at test time. To tackle this problem, we propose a novel form of data augmentation: rather than augment observations, we focus on *augmenting the dynamics*. We first learn a world model of the environment, and then augment the transition function at

^{*}Equal contribution ¹University of Oxford, United Kingdom. Correspondence to: Philip Ball <ball@robots.ox.ac.uk>, Cong Lu <cong.lu@stats.ox.ac.uk>.




Figure 1. An illustration of our approach, each figure shows a transition $P(s, a) \rightarrow s'$. In a) we show the World Model dynamics \hat{P} , trained from \mathcal{D}_{env} , a single offline dataset. In b) we show the Augmented World Model, blue represents \hat{P} , while each green agent illustrates an instantiation of augmented dynamics, which is sampled at each timestep: $\hat{P}_z, z \sim \mathcal{Z}$. The goal is to approximate c) where we show an unseen test environment, with transition dynamics P^* .

policy training time, making the agent train under different imagined dynamics. In addition, our agent is given access to the augmentation itself as part of the observations, allowing it to consider the context of modified dynamics.

At test time we propose a simple, yet surprisingly effective, self-supervised approach to learning an agent’s augmentation context. We learn a linear dynamics model which is then used to approximate the dynamics augmentation induced by the modified environment. This context is then given to the agent, allowing it to adapt on the fly to the new dynamics within a single episode (i.e. zero-shot). We show that our approach is capable of training agents that can vastly outperform existing Offline RL methods on the “dynamics generalization from a single offline environment” problem. We also note that this approach does *not* require access to environment rewards at test time. This facilitates application to Sim2Real problems whereby test time rewards may not be available.

Our contributions are twofold: 1) As far as we are aware, we are the first to propose **dynamics augmentation for model based RL**, allowing us to generalize to changing dynamics despite only training on a single setting. We do this without access to any environment parameters or prior knowledge. 2) We propose a simple **self-supervised context adaptation** reward-free algorithm, which allows our policy to use information from interactions in the environment to vary its behavior in a single episode, increasing zero-shot performance. We believe both of these approaches are not only novel, but offer significant improvement v.s. state-of-the-art methods, improving generalization and providing a promising approach for using offline RL in the real world.

2. Related Work

In this work we focus on Model Based RL (MBRL). A key challenge in MBRL is that an inaccurate model can be exploited by the policy, leading to behaviors that fail to transfer to the real environment. As such, a swathe of recent works have made use of model ensembles to improve robustness (Kurutach et al., 2018; Chua et al., 2018; Clavera et al., 2018; Ball et al., 2020; Janner et al., 2019). With increased accuracy, MBRL has recently been shown to be

competitive with model free methods in continuous control (Ha and Schmidhuber, 2018; Chua et al., 2018; Janner et al., 2019) and games (Schrittwieser et al., 2019; Kaiser et al., 2020). We make use of the ensemble of probabilistic dynamics models, first introduced in (Chua et al., 2018), and subsequently used in (Janner et al., 2019).

In this paper we focus on Model-Based *offline* RL, where MOPO (Yu et al., 2020) and MOREL (Kidambi et al., 2020) have recently demonstrated the effectiveness of learned dynamics models, using model uncertainty to constrain policy optimization. We build upon this approach for zero-shot dynamics generalization from offline data. There have also been successes in off policy methods for offline RL (Wu et al., 2019; Fujimoto et al., 2019; Kumar et al., 2020; Rudner et al., 2021) and context based approaches (Ajay et al., 2021), although these works only consider tasks within the support of the offline dataset. Finally, MBOP (Argenson and Dulac-Arnold, 2021) addresses the problem of goal-conditioned zero-shot transfer from offline datasets. However, their goal-conditioning relies on unchanged dynamics in the test environment.

In online RL, recent work has achieved strong dynamics generalization with a learned model (Seo et al., 2020). However, this required training under varied dynamics, assigning different experiences to models. In addition, this work used MPC whereas we train a policy inside the model, which is significantly faster at deployment time. Also related is (Clavera et al., 2019; Nagabandi et al., 2019), where the model is trained to quickly to adapt to new dynamics $P(s'|s, a)$, however both these works place more emphasis on model-adaptation rather than zero-shot policy performance. Furthermore, access to an underlying task distribution is required, something we do not have in our offline setting. Also similar to our work is the recently proposed Policy Adaptation during Deployment (PAD, (Hansen et al., 2021)) approach. Our approach differs in that we learn a context, whereas PAD uses an auxiliary objective to adapt its features. In addition, PAD considers the *online model free* setting, while our method is *offline and model based*.

Sim2Real is the setting where an agent trained in a simulator must transfer to the real world. A common approach to

solve this problem is through domain randomization (Tobin et al., 2017; James et al., 2017), whereby parameters in the simulator are varied during training. This has shown to be effective for dynamics generalization (Andrychowicz et al., 2020; Antonova et al., 2017; Peng et al., 2018; Yu et al., 2017; Zhou et al., 2019; OpenAI et al., 2019), but requires access to a simulator which we do not have. Another form of domain randomization, data augmentation, has proved to be effective for training RL policies (Laskin et al., 2020a;b; Kostrikov et al., 2021; Raileanu et al., 2020), resulting in improved efficiency and generalization. So far, these works have focused on online model free methods, and used data augmentation on the *state* space, reducing *observational* overfitting (Song et al., 2020). In contrast, we focus on offline MBRL and instead augment the *dynamics*.

We also note clear links to contextual MDPs (Hallak et al., 2015; Modi et al., 2018) and hidden parameter MDPs (HiP-MDP) (Doshi-Velez and Konidaris, 2016; Killian et al., 2017; Zhang et al., 2021) settings, whereby our self-supervised dynamics embedding can be considered as a context/hidden parameter. However we note that in these settings the embedding is chosen at the beginning of each episode, and is fixed throughout, whereas our embedding varies per timestep.

We are not the first to propose data augmentation in the MBRL setting, (Pitis et al., 2020) proposed Counterfactual data augmentation for improving performance in the context of locally factored tasks. Approaches to ensuring adversarial robustness can include data augmentations that assist with out-of-domain generalization, as opposed to observational overfitting. In (Volpi et al., 2018) this is done without a simulator and from a single source of data, however they only work on supervised learning problems and require an adversary to be learned, adding computational complexity.

3. Preliminaries

We consider a Markov Decision Process (MDP), defined as a tuple $M = (\mathcal{S}, \mathcal{A}, P, R, \rho_0, \gamma)$, where \mathcal{S} and \mathcal{A} denote the state space and action space respectively, $P(s'|s, a)$ the transition dynamics, $R(s, a)$ the reward function, ρ_0 the initial state distribution, and $\gamma \in (0, 1)$ the discount factor. The goal in RL is to optimize a policy $\pi(a|s)$ that maximizes the expected discounted return $\mathbb{E}_{\pi, P, \rho_0} [\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t)]$. The value function $V^\pi(s) := \mathbb{E}_{\pi, P} [\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t) | s_0 = s]$ gives the expected discounted return under π when starting from state s . In *offline RL*, the policy is not deployed in the environment until test time. Instead, the algorithm only has access to a static dataset $\mathcal{D}_{env} = \{(s, a, r, s')\}$, collected by one or more behavioral policies π_b . We borrow notation from (Yu et al., 2020) and refer to the distribution from which \mathcal{D}_{env} was sampled as the *behavioral distribution*.

When training a model, we follow MBPO (Janner et al., 2019) and MOPO (Yu et al., 2020) and train an ensemble of N probabilistic dynamics models (Nix and Weigend, 1994). Each of the models learns to predict both the next state s' and reward r from a state-action pair, using \mathcal{D}_{env} in a supervised fashion. Concretely, each of the N models output a Gaussian $\hat{P}_i(s_{t+1}, r_t | s_t, a_t) = \mathcal{N}(\mu(s_t, a_t), \Sigma(s_t, a_t))$. The resulting dynamics model \hat{P} defines a *model MDP* $\hat{M} = (\mathcal{S}, \mathcal{A}, \hat{P}, \hat{R}, \rho_0, \gamma)$, where \hat{R} refers to the learned reward model.

To train the policy, we use k step rollouts inside \hat{M} , adding experience to a replay buffer \mathcal{D}_{env} to learn an action-value function and a policy, using Soft Actor Critic (SAC (Haarnoja et al., 2018)). SAC alternates between a soft policy evaluation step, which estimates $Q^\pi(s, a) = \mathbb{E}_{\pi} [\sum_{t=0}^{\infty} \gamma^t (R(s_t, a_t) + \alpha \mathcal{H}(\pi(\cdot | s_k)))] | s_0 = s, a_0 = a]$ using Bellman backups (where α is a temperature parameter for policy entropy \mathcal{H}), and a policy improvement which learns a policy by minimizing the expected KL divergence $J_\pi(\phi, \mathcal{D}) = \mathbb{E}_{s_t \sim \mathcal{D}_{env}} [D_{KL}(\pi \| \exp\{\frac{1}{\alpha} \{Q^\pi - V^\pi\}\})]$. Note that the SAC algorithm is unchanged from the model-free setting, aside from the environment being a learned model, and the rollout horizon k being truncated. Perhaps surprisingly, this approach alone produces strong results in the offline setting. However MOPO (Yu et al., 2020) and MOREl (Kidambi et al., 2020) show improvement in performance by penalizing rewards in regions of the state space where the ensemble of probabilistic models is less certain. This implicitly reduces reliance on samples which deviate beyond the support of \mathcal{D}_{env} .

While MOPO and MOREl have addressed the issue of training a policy in \mathcal{D}_{env} , and transferring to the true environment M , they only consider where the data in \mathcal{D}_{env} is actually drawn from P . However, sometimes this may not be sufficient for deployment. For example, a robot could fail to walk when learning from data that was collected by a different version of the robot (with different mass), or if the same robot collected data but in a different room to deployment (with varied friction). It is this setting, where dynamics may vary at test time, that is the focus of our work. To learn successfully we propose a novel approach to training robust context-dependent policies.

4. Augmented World Models with Self Supervised Policy Adaptation

In this section we introduce our algorithm: Augmented World Models (AugWM). We first discuss our training procedure (Fig. 2) before moving onto our self-supervised approach to online context learning (Fig. 3).




Figure 2. Training policies in Augmented World Models. For each state-action pair sampled from the buffer, a new augmentation $z_t \sim \mathcal{Z}$ is sampled to produce an augmentation operator \mathcal{T}_z , which is applied to the transition. The policy is then trained with the new tuple of data, with the context concatenated to the state.




Figure 3. Self-supervised policy adaptation via learned context. At each timestep, the state s_t is fed into a linear model (trained online at each timestep) to predict the change in state δ_t , and also passed to the fixed World Model (trained on the offline data) to predict the change in state under \hat{P} , $\hat{\delta}_t$. The approximate context is then $z_t = \delta_t / \hat{\delta}_t$, which is concatenated with s_t and passed to the agent to produce an action.

4.1. Augmented World Models

Rather than seeking to transfer our policy from \hat{M} to M , we instead wish to transfer from \hat{M} to M^* , where $M^* = (\mathcal{S}, \mathcal{A}, P^*, R^*, \rho_0, \gamma)$ is an unseen environment with *different dynamics*. Thus, our emphasis shifts to producing experience inside the world model such that our agent is able to generalize to unseen, out of distribution dynamics. We approach this problem by training our policy in an *Augmented World Model*. Formally, we denote an augmentation as $z_t \sim \mathcal{Z}, z \in \mathbb{R}^{|\mathcal{S}|}$, which is sampled at each timestep to produce an augmentation operator \mathcal{T}_z . \mathcal{T}_z is applied to (s, a, r, s') tuples from the dataset \mathcal{D} , and when used with tuples sampled from \mathcal{D}_{env} indirectly induces an augmented distribution \hat{P}_z . In principle, we wish to produce augmentations such that the true modified environment dynamics P^* lies in the support of the distribution of augmented dynamics, i.e.

$$\inf_z D(\hat{P}_z(s, a) \| P^*(s, a)) \leq \epsilon \quad (1)$$

for all s, a , some small $\epsilon > 0$, and suitable distance/divergence metric D . We consider several augmentations, beginning with existing works before moving to new approaches which specifically target the problem of dynamics generalization. We begin with **Random Amplitude Scaling** as in (Laskin et al., 2020a), which we refer to as RAD. RAD scales both s_t and s_{t+1} as follows:

$$\mathcal{T}_z : (s_t, a_t, r_t, s_{t+1}) \mapsto (z \odot s_t, a_t, r_t, z \odot s_{t+1}) \quad (2)$$

for $z \sim \text{Unif}([a, b]^{|\mathcal{S}|})$. Given that our focus is on changing dynamics (vs. observational overfitting), we also propose to scale only the next state, i.e., **Random Amplitude**

Nextstate Scaling (RANS):

$$\mathcal{T}_z : (s_t, a_t, r_t, s_{t+1}) \mapsto (s_t, a_t, r_t, z \odot s_{t+1}) \quad (3)$$

for $z \sim \text{Unif}([a, b]^{|\mathcal{S}|})$. Note that while RANS is more focused on augmenting dynamics than RAS, it still suffers from a dependence on the magnitude of s_{t+1} . As such, we further propose a more targeted augmentation, which we call **Dynamics Amplitude Sampling (DAS)**. Rather than directly scale the state, DAS scales the *change in the state* which we denote with $\delta_t = s_{t+1} - s_t$, as follows:

$$\mathcal{T}_z : (s_t, a_t, r_t, s_{t+1}) \mapsto (s_t, a_t, r_t, s_t + z \odot \delta_t) \quad (4)$$

for $z \sim \text{Unif}([a, b]^{|\mathcal{S}|})$. The full training procedure is shown in Algorithm 1.

Algorithm 1 Augmented World Models: Training

Input: Offline data \mathcal{D}_{env} , Penalty λ , horizon h , batchsize B , augmentation \mathcal{Z} .

Initialize: Ensemble of N dynamics models \hat{P} , policy π .

Replay buffer \mathcal{D}_{env}

1. Train \hat{P} in a supervised fashion using \mathcal{D}_{env} .

for epoch = 1, 2, ... **do**

 (i) Sample initial states: $\{s_1^1, \dots, s_1^B\} \sim \mathcal{D}_{\text{env}}$

 (ii) Rollout policy (in parallel), using λ -penalized reward, storing all data in \mathcal{D}_{env}

 (iii) Train policy using $\mathcal{D}_{\text{env}} \cup \mathcal{D}_{\text{env}}$. For each (s, a, r, s') , sample $z \sim \mathcal{Z}$ and apply \mathcal{T}_z .

end

Return: Policy π

One crucial addition to our method is the use of context. Concretely, when we are optimizing the policy using a batch of data, we concatenate the next state with the augmentation

vector z . This allows our policy to be informed of the specific augmentation that was applied to the environment and thus behave accordingly. However, at test time we do not know z , so what can we use? Next we propose a solution to this problem, learning the context on the fly.

4.2. Self-Supervised Policy Selection

In the meta-learning literature there have been many recent successes making use of a *learned context* to adapt a policy at test time to a new environment (Rakelly et al., 2019; Zintgraf et al., 2020; 2021), typically using a blackbox model with a latent state. Crucially, these approaches require several episodes to adapt at test time, making them unfeasible in our zero-shot setting. What makes our setting unique is we explicitly know what the context represents: a linear transformation of s_t , or δ_t . Using this insight, we are able to learn an effective context on the fly at test time. Concretely, we observe that from a state s_t drawn from M^* , we can sample an action $a_t \sim \pi$ and then compute an approximate \hat{s}_{t+1} using our model \hat{P} . With \hat{s}_{t+1} , we have a sample estimate of the *state change under \hat{P}* , i.e. $\hat{\delta}_t = \hat{s}_{t+1} - s_t$. We can make this approximation of the next state without interacting with the environment, but once we do take the action a_t in the environment, we then receive the *true next state* s_{t+1} and can store the true difference $\delta_t = s_{t+1} - s_t$. Using the DAS augmentation, we can approximate z as δ/δ .

Algorithm 2 Augmented World Models: Testing

Input: Initial state s_1 , horizon H , policy π , world model \hat{P} , initial context $\hat{z} = \mathbf{1}^{|S|}$

Initialize: Linear model f_ψ , dataset $\mathcal{D} = \emptyset$, return $R_0 = 0$.

for $step = 1, 2, \dots, H$ **do**

Select action: $a_t \sim \pi(s_t, \hat{z}_t)$

Take action: $s_{t+1} \sim P^*(s_t, a_t)$

Update return: $R_{t+1} = R_t + R^*(s_t, a_t, s_{t+1})$

Update Dataset: $\mathcal{D} \cup (s_t, \delta_t)$.

Update Linear model by minimizing $\mathcal{L}_{\text{MSE}}(\psi, \mathcal{D})$.

Predict new context using $\hat{z}_t = \delta/\delta$.

end

Return: R_T

This however is retrospective; we can only approximate z having already seen the next state, by which time our agent has already acted. Furthermore, we believe under changed dynamics the true z likely depends on s , thus we cannot use a previous z for future timesteps. Therefore, we learn a forward dynamics model using the data collected *during* the test rollout. After h timesteps in the environment, we have the following dataset: $\mathcal{D} = \{(s_1, s_2), \dots, (s_{h-1}, s_h)\}$. This allows us to learn a simple dynamics model $f_\psi : (s_t) \mapsto \delta_t = s_{t+1} - s_t$, by minimizing the mean squared error $\mathcal{L}_{\text{MSE}}(\psi, \mathcal{D})$. Notably, since we never actually plan with this model, it does not need to be as accurate as a typi-

cal dynamics model in MBRL. Instead, it is crucial that the model learns quickly enough such that we can use it in a zero-shot evaluation. Thus, we choose to use a simple linear model for f . To show the effectiveness of this, in Fig. 4 we show the mean R-squared of linear models learned on the fly during evaluation rollouts.




Figure 4. Mean R^2 of the linear model across 20 rollouts.

We observe that in less than 100 timesteps the linear model achieves high accuracy on the test data. Subsequently, equipped with f_ψ , we can approximate δ_t , and predict the augmentation as $\hat{z}_t = \delta/\delta$. We then provide the agent with \hat{z}_t to compute action a_t . The full procedure is shown in Algorithm 2.

5. Experiments

In our experiments we aim to investigate the effectiveness of our approach for zero-shot dynamics generalization from a single offline dataset. To assess this, we will answer a series of questions, beginning with a question on the necessity of our method:

Do we really need to develop methods specifically for dynamics generalization?

To answer this, we train MOPO using offline data from a single environment, and test it under changed dynamics. We consider the HalfCheetah environment from the OpenAI Gym (Brockman et al., 2016), using offline data from D4RL (Fu et al., 2021). We train a MOPO agent using the mixed dataset, using our own implementation of the algorithm (but using the same hyperparameters as the original authors). To test the trained policy, we vary both the mass of the agent and damping coefficient by a multiplicative factor.¹ In this work we consider a grid of the following values for HalfCheetah: $\{0.25, 0.5, 0.75, 1.0, 1.25, 1.5, 1.75\}$ and $\{0.5, 0.75, 1.0, 1.25, 1.5\}$ for Walker2d, representing a significant out-of-distribution shift.

The results (Fig. 5), show that MOPO performance is clearly impacted by changing dynamics. We see in the central cell, that performance for our version of MOPO matches the author results (Yu et al., 2020), and in some cases we even see small gains (e.g. mass = 0.75, damping = 1.0). However,

¹The standard environment (both in Gym and D4RL) corresponds to both these values being set to 1.0.

Mass Scale	0.25	0.5	0.75	1	1.25	1.5	1.75
0.25	561	545	905	1715	2311	2790	3217
0.5	785	1723	2421	2786	5579	5146	4517
0.75	1448	2442	4358	6482	6196	5554	4894
1	2886	4582	5518	6336	5955	5299	4336
1.25	3247	5356	5957	5732	5104	4194	3520
1.5	3720	5206	5370	4599	3810	3259	2867
1.75	2968	4323	3843	3496	3124	2682	2305

Figure 5. Mean performance for 5 seeds for MOPO on HalfCheetah environments with varying dynamics. Note the central cell (1,1) corresponds to the in-sample data.

on the top left we see dramatically weaker performance, often below 1k, indicating the robot is failing to achieve locomotion. Before evaluating AugWM, we first test whether training with the ‘correct’ augmentation improves generalization performance. In short, we ask:

Is augmenting dynamics even worthwhile?

To answer this, we train SAC for 1×10^5 steps and save the states visited in the ‘true’ environment. We then use these starting states to train a policy using an offline MBPO² approach with AugWM. However, instead of sampling $z_t \sim \mathcal{Z}$, we provide the actual $z = \delta/\delta$ as we have access to the ‘true’ and ‘modified’ environments; we refer to this as an *oracle* version of our method, and is designed to assess the viability of our approach. Note that we *do not* augment the ‘true’ environment rewards. We consider two baselines: a) offline MBPO in the ‘true’ environment; b) online SAC in the ‘modified’ environment. We train MBPO until convergence, and SAC for 1×10^5 steps.




Figure 6. Mean performance for 3 seeds for MBPO with a dynamics oracle on HalfCheetah with $1.5 \times$ mass and damping.

As shown Fig. 6, when provided with the true z , AugWM outperforms both baselines. The SAC result is surprising: the baseline agent was trained directly on the ‘modified’ environment for the same number of steps as the policy that generated our oracle starting states. One explanation

²Since we have access to the true environment, there is no need for the MOPO penalty.

is the greater exploration induced by the ‘easier’ dynamics of the ‘true’ environment. This validates our approach; if we augment the dynamics \hat{P} from a model correctly, we can generalize to unseen dynamics. In other words, neither the starting states nor rewards need to be from the ‘modified’ environment. With this in mind, our next question is a simple one:

Which augmentation strategy is most effective?

To test this, we train as in Algorithm 1, *without* context, to isolate the effectiveness of the training process. We use the HalfCheetah Mixed dataset and train a MOPO agent, augmenting either both s and s' (RAD), just s' (RANS) or just δ (DAS).

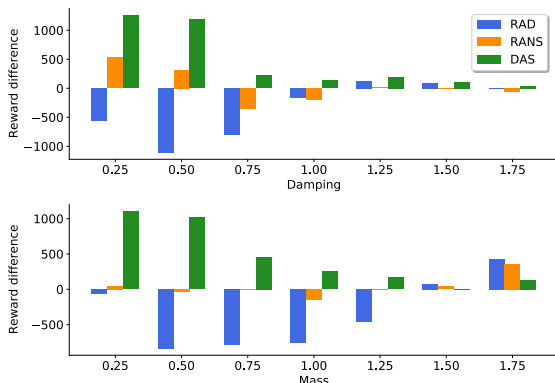


Figure 7. Average performance gains of different World Model augmentations over base MOPO for different levels of damping and mass in the HalfCheetah test environment (5 seeds).

The results are shown in Fig. 7. As we see, the RAD augmentation fails to improve dynamics generalization, actually leading to worse performance overall. RANS does improve performance on unseen dynamics, as we are influencing the *dynamics*, not just the observation. However, DAS clearly provides the strongest performance. As a result, we use DAS for AugWM. Our final algorithm design question is as follows:

Does training with context improve performance?

To answer this question, we return to the HalfCheetah Mixed setting from Fig. 7, taking the policy trained with DAS. We now train two additional agents: 1) Default Context: at *train* time the agent is provided with the DAS augmentation z as context, at *test* time it is provided with a vector of ones, $\mathbf{1}^{|\mathcal{S}|}$; 2) Learned Context: trained as in 1), but context is learned online using Algorithm 2.

The results are shown in in Fig. 8. We observe that training with context (orange) improves performance on average, while adapting the context on the fly (green) leads to further gains (+80 on average). These methods combine to produce our AugWM algorithm. We are now ready for the final question:




Figure 8. Average performance gains of adding contexts over base MOPO for different levels of damping and mass in the HalfCheetah test environment (5 seeds).

Can Augmented World Models improve zero-shot generalization?

To answer this question we perform a rigorous analysis, using multiple benchmarks from the D4RL dataset (Fu et al., 2021). Namely, we consider the Random, Medium, Mixed and Medium-Expert datasets for both Walker2d and HalfCheetah. In each setting, we compare AugWM against base MOPO on zero-shot performance, training entirely on the data provided, but not seeing the test environment until evaluation. The results are shown as a change v.s. MOPO, averaged over one dimension in Fig. 11, and as a total return number averaged over both dimensions in Table 1. For additional implementation details (e.g., hyperparameters) see Appendix B. AugWM provides *statistically significant* improvements in zero-shot performance v.s. MOPO in many cases, achieving successful policies where MOPO fails.

Table 1. Each entry for HalfCheetah is the mean of 49 different dynamics, while for Walker2d it is over 25 dynamics. Results are mean \pm 1std. \star indicates $p < 0.05$ for Welch’s t-test for gain over MOPO (5 seeds).

Dataset Type	Environment	MOPO	AugWM (Ours)
Random	HalfCheetah	2303 \pm 112	2818 \pm 197 \star
Random	Walker2d	569 \pm 103	706 \pm 139
Mixed	HalfCheetah	3447 \pm 218	3948 \pm 122 \star
Mixed	Walker2d	946 \pm 95	1317 \pm 206 \star
Medium	HalfCheetah	2954 \pm 89	2967 \pm 106
Medium	Walker2d	1477 \pm 337	1614 \pm 440
Med-Expert	HalfCheetah	1590 \pm 766	2885 \pm 432 \star
Med-Expert	Walker2d	1062 \pm 334	2521 \pm 316 \star

By now we have provided significant evidence that AugWM can significantly improve performance for HalfCheetah and Walker2d with varied mass and damping. However, this is only a small subset of possible dynamics changes. We next consider several significantly harder settings. We test increased dimensionality, using the Ant environment from MOPO (Yu et al., 2020), and consider additional types of dynamics changes (e.g., Ant with crippled legs, HalfCheetah

with changed limb sizes from (Henderson et al., 2017)). We illustrate the impact of the crippled leg Ant environment on baseline agent performance, and the improvement provided by AugWM, in Fig. 12. We show the mean results over each of these factors of variation in Table 2, where again AugWM provides a non-trivial improvement over a strong baseline. For more details see Appendix B.

Table 2. Mean performance for MOPO, AugWM with the default context, and AugWM with learned context (LM). Entries are mean zero-shot reward for all dynamics. Bold = highest (5 seeds).

Setting	MOPO	AugWM (Default)	AugWM (LM)
Ant: Mass/Damp	1634	1715	1804
Ant: One Crippled Leg	1370	1572	1680
Ant: Two Crippled Legs	700	697	795
HalfCheetah: Big	4891	5194	4968
HalfCheetah: Small	5151	5488	5263

Finally, we note that dynamics may change *during* an episode; consider a robot that suffers a motor fault, reducing the power delivered to its joints. Evidently the underlying dynamics have been altered, and being robust to such changes when only training from a single dataset of offline experience is challenging. To illustrate this, we perform a 1500 step rollout in the HalfCheetah environment, starting with offline dynamics (mass/damping = 1), before changing to mass = 0.75, damping = 0.5 after 500 steps; performance is shown in Fig. 9. Observe that after 500 steps, MOPO performance is dramatically reduced. This is because the agent continues to apply the same force and thus falls forward with lighter mass. For our AugWM agent, performance initially drops, then when the new context is learned we achieve *higher* performance than before, making use of the lighter torso.




Figure 9. Performance under changing dynamics. Top = cumulative returns, Bottom = rolling average single step reward. Both averaged over twenty seeds, shaded area shows \pm 1std.




Figure 10. Mean improvement for Augmented World Models over MOPO for the HalfCheetah environment, averaged over five seeds. Top row (blue) = damping scale, bottom row (green) = mass scale. Dotted line is the mean, the same value for both damping and mass.




Figure 11. Mean improvement for Augmented World Models over MOPO for the Walker2d environment, averaged over five seeds. Top row (blue) = damping scale, bottom row (green) = mass scale. Dotted line is the mean, the same value for both damping and mass.




Figure 12. Comparison between MOPO and AugWM on the crippled leg Ant environment. The crippled leg is highlighted in red. The MOPO agent flips over while the AugWM agent is able to adapt to changed dynamics and successfully run. MOPO on the regular Ant environment is provided as a reference.

Discussion We believe that our experiments provide significant support to the claim that training with AugWM improves zero-shot dynamics generalization. In a broad set of commonly used datasets, and with a wide range of out-of-distribution dynamics, our algorithm learns good policies where a state-of-the-art baseline fails.³ This is due to a number of novel contributions: 1) using *dynamics augmentation* rather than *observation augmentation*; 2) training and testing with a context-based policy. Regarding limitations, we note that training inside the WM with context generally takes longer to converge (Appendix B). Furthermore,

³For videos see: <https://sites.google.com/view/augmentedworldmodels/>

in more nonlinear settings such as the HalfCheetah modified body part setting we saw a reduced performance for the learned context. This could be because the dynamics changes are out of the distribution of DAS augmentations (violating Eqn. 1), or due to the difficulty of modeling the task with a linear model. We note that linear models have achieved success in planning (Gu et al., 2016) and meta learning (Peng et al., 2021), and are effective in our case due to their data efficiency, but can be replaced by more flexible models to deal with different augmentations. Indeed, given our work is the first of its kind, we believe significant improvements are possible, such as using more complex and problem-specific augmentations.

6. Conclusion and Future Work

In this paper we propose Augmented World Models (AugWM), which we show sufficiently simulates changes in dynamics such that agents can generalize in a zero-shot manner. We believe that we are the first to propose this problem setting, and our results show a significant improvement over existing state-of-the-art methods which ignore this problem.

A promising line of future work would be to meta-train a policy over AugWM such that it can quickly adapt to new dynamics in the few-shot setting. There is evidence that data augmentation can improve robustness in meta-learning (Rajendran et al., 2020), and could extend to strong performance in out-of-distribution tasks. We also wish to consider varying goals at test time, and other potential sources of non-stationarity which may impact policy performance (Igl et al., 2021). It may also be possible to extend AugWM to pixel-based tasks, which have received a great deal of recent attention (Hafner et al., 2019; 2020). We believe that our transition based augmentations will be applicable to a latent representation, as commonly used in state-of-the-art vision MBRL approaches. Thus we think that extending our work to this setting, while a considerable feat of engineering, should not require significant methodological changes.

Acknowledgements

Philip Ball is funded through the Willowgrove Studentship. Cong Lu is funded by the Engineering and Physical Sciences Research Council (EPSRC). We are grateful to Taylor Killian for useful discussions on contextual/HIP-MDPs. The authors would also like to thank the anonymous ICML reviewers and the area chair for constructive feedback which helped to improve the paper.

References

- Ajay, A., Kumar, A., Agrawal, P., Levine, S., and Nachum, O. (2021). OPAL: Offline primitive discovery for accelerating offline reinforcement learning. In *International Conference on Learning Representations*.
- Andrychowicz, O. M., Baker, B., Chociej, M., Józefowicz, R., McGrew, B., Pachocki, J., Petron, A., Plappert, M., Powell, G., Ray, A., Schneider, J., Sidor, S., Tobin, J., Welinder, P., Weng, L., and Zaremba, W. (2020). Learning dexterous in-hand manipulation. *The International Journal of Robotics Research*, 39(1):3–20.
- Antonova, R., Cruciani, S., Smith, C., and Kragic, D. (2017). Reinforcement learning for pivoting task.
- Argenson, A. and Dulac-Arnold, G. (2021). Model-based offline planning. In *International Conference on Learning Representations*.
- Ball, P., Parker-Holder, J., Pacchiano, A., Choromanski, K., and Roberts, S. (2020). Ready policy one: World building through active learning. In *Proceedings of the 37th International Conference on Machine Learning, ICML*.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. (2016). Openai gym. *arXiv preprint arXiv:1606.01540*.
- Chua, K., Calandra, R., McAllister, R., and Levine, S. (2018). Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems 31*, pages 4754–4765.
- Clavera, I., Nagabandi, A., Liu, S., Fearing, R. S., Abbeel, P., Levine, S., and Finn, C. (2019). Learning to adapt in dynamic, real-world environments through meta-reinforcement learning. In *International Conference on Learning Representations*.
- Clavera, I., Rothfuss, J., Schulman, J., Fujita, Y., Asfour, T., and Abbeel, P. (2018). Model-based reinforcement learning via meta-policy optimization. In *2nd Annual Conference on Robot Learning, CoRL 2018, Zürich, Switzerland, 29-31 October 2018, Proceedings*, volume 87 of *Proceedings of Machine Learning Research*, pages 617–629. PMLR.
- Doshi-Velez, F. and Konidaris, G. (2016). Hidden parameter markov decision processes: A semiparametric regression approach for discovering latent task parametrizations. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI’16*, page 1432–1440. AAAI Press.
- Dulac-Arnold, G., Mankowitz, D. J., and Hester, T. (2019). Challenges of real-world reinforcement learning. *CoRR*, abs/1904.12901.
- Fu, J., Kumar, A., Nachum, O., Tucker, G., and Levine, S. (2021). D4{rl}: Datasets for deep data-driven reinforcement learning.
- Fujimoto, S., Meger, D., and Precup, D. (2019). Off-policy deep reinforcement learning without exploration. In *International Conference on Machine Learning*, pages 2052–2062.
- Gu, S., Lillicrap, T., Sutskever, I., and Levine, S. (2016). Continuous deep q-learning with model-based acceleration. In Balcan, M. F. and Weinberger, K. Q., editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2829–2838, New York, New York, USA. PMLR.

- Ha, D. and Schmidhuber, J. (2018). Recurrent world models facilitate policy evolution. In *Proceedings of the 32Nd International Conference on Neural Information Processing Systems*, NeurIPS’18, pages 2455–2467.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., and Levine, S. (2018). Soft actor-critic algorithms and applications. *CoRR*, abs/1812.05905.
- Hafner, D., Lillicrap, T., Ba, J., and Norouzi, M. (2020). Dream to control: Learning behaviors by latent imagination. In *International Conference on Learning Representations*.
- Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., and Davidson, J. (2019). Learning latent dynamics for planning from pixels. In Chaudhuri, K. and Salakhutdinov, R., editors, *Proceedings of the 36th International Conference on Machine Learning*, pages 2555–2565.
- Hallak, A., Castro, D. D., and Mannor, S. (2015). Contextual markov decision processes.
- Hansen, N., Jangir, R., Sun, Y., Alenyà, G., Abbeel, P., Efros, A. A., Pinto, L., and Wang, X. (2021). Self-supervised policy adaptation during deployment. In *International Conference on Learning Representations*.
- Henderson, P., Chang, W.-D., Shkurti, F., Hansen, J., Meger, D., and Dudek, G. (2017). Benchmark environments for multitask learning in continuous domains. *ICML Lifelong Learning: A Reinforcement Learning Approach Workshop*.
- Igl, M., Farquhar, G., Luketina, J., Boehmer, W., and White-son, S. (2021). Transient non-stationarity and generalisation in deep reinforcement learning. In *International Conference on Learning Representations*.
- James, S., Davison, A. J., and Johns, E. (2017). Transferring end-to-end visuomotor control from simulation to real world for a multi-stage task. In *1st Conference on Robot Learning*.
- Janner, M., Fu, J., Zhang, M., and Levine, S. (2019). When to trust your model: Model-based policy optimization. In *Advances in Neural Information Processing Systems*.
- Kaiser, L., Babaeizadeh, M., Milos, P., Osinski, B., Campbell, R. H., Czechowski, K., Erhan, D., Finn, C., Koza-kowski, P., Levine, S., Mohiuddin, A., Sepassi, R., Tucker, G., and Michalewski, H. (2020). Model based reinforcement learning for Atari. In *International Conference on Learning Representations*.
- Kidambi, R., Rajeswaran, A., Netrapalli, P., and Joachims, T. (2020). Morel : Model-based offline reinforcement learning. In *Advances in Neural Information Processing Systems*.
- Killian, T., Konidaris, G., and Doshi-Velez, F. (2017). Robust and efficient transfer learning with hidden parameter markov decision processes. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1).
- Kostrikov, I., Yarats, D., and Fergus, R. (2021). Image augmentation is all you need: Regularizing deep reinforcement learning from pixels. In *International Conference on Learning Representations*.
- Kumar, A., Zhou, A., Tucker, G., and Levine, S. (2020). Conservative q-learning for offline reinforcement learning. In *Advances in Neural Information Processing Systems*.
- Kurutach, T., Clavera, I., Duan, Y., Tamar, A., and Abbeel, P. (2018). Model-ensemble trust-region policy optimization. In *International Conference on Learning Representations*.
- Laskin, M., Lee, K., Stooke, A., Pinto, L., Abbeel, P., and Srinivas, A. (2020a). Reinforcement learning with augmented data. In *Advances in Neural Information Processing Systems 33*.
- Laskin, M., Srinivas, A., and Abbeel, P. (2020b). CURL: Contrastive unsupervised representations for reinforcement learning. In *Proceedings of the 37th International Conference on Machine Learning*.
- Levine, S., Kumar, A., Tucker, G., and Fu, J. (2020). Offline reinforcement learning: Tutorial, review, and perspectives on open problems.
- Modi, A., Jiang, N., Singh, S., and Tewari, A. (2018). Markov decision processes with continuous side information. In Janoos, F., Mohri, M., and Sridharan, K., editors, *Proceedings of Algorithmic Learning Theory*, volume 83 of *Proceedings of Machine Learning Research*, pages 597–618. PMLR.
- Nagabandi, A., Finn, C., and Levine, S. (2019). Deep online learning via meta-learning: Continual adaptation for model-based RL. In *International Conference on Learning Representations*.
- Nix, D. A. and Weigend, A. S. (1994). Estimating the mean and variance of the target probability distribution. In *Proceedings of 1994 IEEE International Conference on Neural Networks (ICNN’94)*, volume 1, pages 55–60 vol.1.
- OpenAI, Akkaya, I., Andrychowicz, M., Chociej, M., Litwin, M., McGrew, B., Petron, A., Paino, A., Plappert, M., Powell, G., Ribas, R., Schneider, J., Tezak, N.,

- Tworek, J., Welinder, P., Weng, L., Yuan, Q., Zaremba, W., and Zhang, L. (2019). Solving rubik’s cube with a robot hand. *CoRR*, abs/1910.07113.
- Peng, M., Zhu, B., and Jiao, J. (2021). Linear representation meta-reinforcement learning for instant adaptation. *CoRR*.
- Peng, X. B., Andrychowicz, M., Zaremba, W., and Abbeel, P. (2018). Sim-to-real transfer of robotic control with dynamics randomization. In *IEEE International Conference on Robotics and Automation, ICRA*.
- Pitis, S., Creager, E., and Garg, A. (2020). Counterfactual data augmentation using locally factored dynamics. In *Advances in Neural Information Processing Systems*.
- Raileanu, R., Goldstein, M., Yarats, D., Kostrikov, I., and Fergus, R. (2020). Automatic data augmentation for generalization in deep reinforcement learning. *CoRR*, abs/2006.12862.
- Rajendran, J., Irpan, A., and Jang, E. (2020). Meta-learning requires meta-augmentation. In *Advances in Neural Information Processing Systems*.
- Rakelly, K., Zhou, A., Finn, C., Levine, S., and Quillen, D. (2019). Efficient off-policy meta-reinforcement learning via probabilistic context variables. In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97, pages 5331–5340. PMLR.
- Rudner, T. G. J., Lu, C., Osborne, M. A., Gal, Y., and Teh, Y. W. (2021). On pathologies in kl-regularized reinforcement learning from expert demonstrations. *ICLR RobustML Workshop*.
- Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., Guez, A., Lockhart, E., Hassabis, D., Graepel, T., Lillicrap, T. P., and Silver, D. (2019). Mastering Atari, Go, Chess and Shogi by Planning with a Learned Model. *CoRR*, abs/1911.08265.
- Seo, Y., Lee, K., Clavera, I., Kurutach, T., Shin, J., and Abbeel, P. (2020). Trajectory-wise multiple choice learning for dynamics generalization in reinforcement learning. In *Advances in Neural Information Processing Systems*.
- Song, X., Jiang, Y., Tu, S., Du, Y., and Neyshabur, B. (2020). Observational overfitting in reinforcement learning. In *International Conference on Learning Representations*.
- Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., and Abbeel, P. (2017). Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 23–30.
- Volpi, R., Namkoong, H., Sener, O., Duchi, J. C., Murino, V., and Savarese, S. (2018). Generalizing to unseen domains via adversarial data augmentation. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 31, pages 5334–5344. Curran Associates, Inc.
- Wu, Y., Tucker, G., and Nachum, O. (2019). Behavior regularized offline reinforcement learning. *arXiv preprint arXiv:1911.11361*.
- Yu, T., Thomas, G., Yu, L., Ermon, S., Zou, J., Levine, S., Finn, C., and Ma, T. (2020). Mopo: Model-based offline policy optimization. In *Advances in Neural Information Processing Systems*.
- Yu, W., Tan, J., Liu, C. K., and Turk, G. (2017). Preparing for the unknown: Learning a universal policy with online system identification. In Amato, N. M., Srinivasa, S. S., Ayanian, N., and Kuindersma, S., editors, *Robotics: Science and Systems XIII*.
- Zhang, A., Sodhani, S., Khetarpal, K., and Pineau, J. (2021). Learning robust state abstractions for hidden-parameter block {mdp}s. In *International Conference on Learning Representations*.
- Zhou, W., Pinto, L., and Gupta, A. (2019). Environment probing interaction policies. In *International Conference on Learning Representations*.
- Zintgraf, L., Feng, L., Lu, C., Igl, M., Hartikainen, K., Hofmann, K., and Whiteson, S. (2021). Exploration in approximate hyper-state space for meta reinforcement learning. In *International Conference on Machine Learning*.
- Zintgraf, L., Shiarlis, K., Igl, M., Schulze, S., Gal, Y., Hofmann, K., and Whiteson, S. (2020). Varibad: A very good method for bayes-adaptive deep rl via meta-learning. In *International Conference on Learning Representations*.