

---

# Goal-Aware Prediction: Learning to Model What Matters

---

Suraj Nair<sup>1</sup> Silvio Savarese<sup>1</sup> Chelsea Finn<sup>1</sup>

## Abstract

Learned dynamics models combined with both planning and policy learning algorithms have shown promise in enabling artificial agents to learn to perform many diverse tasks with limited supervision. However, one of the fundamental challenges in using a learned forward dynamics model is the mismatch between the objective of the learned model (future state reconstruction), and that of the downstream planner or policy (completing a specified task). This issue is exacerbated by vision-based control tasks in diverse real-world environments, where the complexity of the real world dwarfs model capacity. In this paper, we propose to direct prediction towards task relevant information, enabling the model to be aware of the current task and encouraging it to only model relevant quantities of the state space, resulting in a learning objective that more closely matches the downstream task. Further, we do so in an entirely self-supervised manner, without the need for a reward function or image labels. We find that our method more effectively models the relevant parts of the scene conditioned on the goal, and as a result outperforms standard task-agnostic dynamics models and model-free reinforcement learning.

## 1 Introduction

Enabling artificial agents to learn from their prior experience and generalize their knowledge to new tasks and environments remains an open and challenging problem. Unlike humans, who have the remarkable ability to quickly generalize skills to new objects and task variations, current methods in multi-task reinforcement learning require heavy supervision across many tasks before they can even begin to generalize well. One way to reduce the dependence on

heavy supervision is to leverage data that the agent can collect autonomously without rewards or labels, termed *self-supervision*. One of the more promising directions in learning transferable knowledge from this unlabeled data lies in learning the dynamics of the environment, as the physics underlying the world are often consistent across scenes and tasks. However learned dynamics models do not always translate to good downstream task performance, an issue which we study and attempt to mitigate in this work.

While learning dynamics in low-dimensional state spaces has shown promising results (McAllister and Rasmussen, 2016; Deisenroth and Rasmussen, 2011; Chua et al., 2018; Amos et al., 2018a), scaling to high dimensional states, such as image observations, poses numerous challenges. One key challenge is that, in high dimensional spaces, learning a perfect model is often impossible due to limited model capacity, and as a result downstream task-specific planners/policies struggle from inaccurate model predictions. Specifically, a learned planner/policy will often exploit errors in the model that make it drastically overestimate its performance. Furthermore, depending on the nature of the downstream task, prediction accuracy on certain states may be more important, which is not captured by the next state reconstruction objective used to train forward dynamics models.

Our primary hypothesis is that this “objective mismatch” between the training objective of the learned model (future state reconstruction), and the downstream planner or policy (completing a specified task) is one of the primary limitations in learning models of high dimensional states. In other words, the learned model is encouraged to predict large portions of the state which may be irrelevant to the task at hand. Consider for example the task of picking up a pen from a cluttered desk. The standard training objective of the learned model would encourage it to equally weigh modeling the pen and all objects on the table, when given the downstream task, modeling the pen and adjacent objects precisely is clearly the most critical.

To that end, we propose goal-aware prediction (GAP), a framework for learning forward dynamics models that direct their capacity differently conditioned on the task, resulting in a model that is more accurate on trajectories most relevant

---

<sup>1</sup>Stanford University. Correspondence to: Suraj Nair <surajn@stanford.edu>.

to the downstream task. Specifically, we propose to learn a latent representation of not just the state, but both the state and goal, and to learn dynamics in this latent space. Furthermore, we can learn this latent space in a way that focuses primarily on parts of the state relative to achieving the goal, namely by reconstructing the *goal-state residual* instead of the full state. We find that this modification combined with training via goal-relabeling (Andrychowicz et al., 2017) allows us to learn expressive, task-conditioned dynamics models in an *entirely self-supervised* manner. We observe that GAP learns dynamics that achieve significantly lower error on task relevant states, and as a result outperforms standard latent dynamics model learning and self-supervised model-free reinforcement learning (Nair et al., 2018) across a range of vision based control tasks.

## 2 Related Work

Recent years have seen impressive results from reinforcement learning (Sutton and Barto, 2018) applied to challenging problems such as video games (Mnih et al., 2015; OpenAI, 2018), Go (Silver et al., 2016), and robotics (Levine et al., 2016; OpenAI et al., 2018; Kalashnikov et al., 2018). However, the dependence on large quantities of labeled data can limit the applicability of these methods in the real world. One approach is to leverage self-supervision, where an agent only uses data that it can collect autonomously.

**Self-Supervised Reinforcement Learning:** Self-supervised reinforcement learning explores how RL can leverage data which the agent can collect autonomously to learn meaningful behaviors, without dependence on task specific reward labels, with promising results on tasks such as robotic grasping and object re-positioning (Pinto and Gupta, 2016; Ebert et al., 2018; Zeng et al., 2018). One approach to self-supervised RL has been combining goal-conditioned policy learning (Kaelbling, 1993; Schaul et al., 2015; Codevilla et al., 2017) with goal re-labeling (Andrychowicz et al., 2017) or sampling goals (Nair et al., 2018; 2019). While there are numerous ways to leverage self-supervised data, ranging from learning distance metrics (Yu et al., 2019b; Hartikainen et al., 2020), generative models over the state space (Kurutach et al., 2018; Fang et al., 2019; Eysenbach et al., 2019; Liu et al., 2020; Nair and Finn, 2020), and representations (Veerapaneni et al., 2019), one of the most heavily utilized techniques is learning the dynamics of the environment (Watter et al., 2015; Finn and Levine, 2017; Agrawal et al., 2016).

**Model-Based Reinforcement Learning:** Learning a model of the dynamics of the environment and using it to complete tasks has been a well studied approach to solving reinforcement learning problems, either through planning with the model (Deisenroth and Rasmussen, 2011; Watter et al., 2015; McAllister and Rasmussen, 2016; Banijamali et al., 2017; Chua et al., 2018; Amos et al., 2018a; Hafner

et al., 2019b; Nagabandi et al., 2019) or optimizing a policy in the model (Racanière et al., 2017; Ha and Schmidhuber, 2018; Łukasz Kaiser et al., 2020; Lee et al., 2019; Janner et al., 2019; Wang and Ba, 2019; Hafner et al., 2019a; Gregor et al., 2019; Byravan et al., 2019). Numerous works have explored how these methods might leverage deep neural networks to extend to high dimensional problem settings, such as images. One technique has been to learn large video prediction models (Finn and Levine, 2017; Babaeizadeh et al., 2018; Ebert et al., 2017; 2018; Paxton et al., 2018; Lee et al., 2018; Villegas et al., 2019; Xie et al., 2019), however model under-fitting remains an issue for these approaches (Dasari et al., 2019). Similarly, many works have explored learning low dimensional latent representations of high dimensional states (Watter et al., 2015; Dosovitskiy and Koltun, 2016; Zhang et al., 2018; Hafner et al., 2019b; Kurutach et al., 2018; Ichter and Pavone, 2019; Wang et al., 2019; Lee et al., 2019; Gelada et al., 2019) and learning the dynamics in the latent space. Unlike these works, we aim to make the problem easier by encouraging the network to predict only task-relevant quantities, while also changing the objective, and hence the distribution of prediction errors, in a task-driven way. This allows the prediction problem to be more directly connected to the downstream use-case of task-driven planning.

**Addressing Model Errors:** Other works have also studied the problem of model error and exploitation. Approaches such as ensembles (Chua et al., 2018; Thananjeyan et al., 2019) have been leveraged to measure uncertainty in model predictions. Similarly, Janner et al. (2019) explore only leveraging the learned model over finite horizons where it has accurate predictions and Levine et al. (2016) use local models. Exploration techniques can also be used to collect more data where the model is uncertain (Pathak et al., 2017).

Most similar to our proposed approach are techniques which explicitly change the models objective to optimize for performance on downstream tasks. (Schrittwieser et al., 2019; Havens et al., 2020) explore only predicting future reward to learn a latent space in which they learn dynamics, Freeman et al. (2019) learn a model with the objective of having a policy achieve high reward from training in it, and Amos et al. (2018b); Srinivas et al. (2018) embed a model/planner inside a neural network. Similarly, Farahmand et al. (2017); D’Oro et al. (2020); Lambert et al. (2020) explore how model training can be re-weighted using value functions, policy gradients, or expert trajectories to emphasize task performance. Unlike these works, which depend heavily on task-specific supervision, our approach can be learned on purely self-supervised data, and generalize to unseen tasks.

## 3 Goal-Aware Prediction

We consider a goal-conditioned RL problem setting (described next), for which we utilize a model-based reinforce-

ment learning approach. The key insight of this work stems from the idea that the distribution of model errors greatly affects task performance and that, when faced with limited model capacity, we can control the distribution of errors to achieve better task performance. We theoretically and empirically investigate this effect in Sections 3.2 and 3.3 before describing our approach for skewing the distribution of model errors in Section 3.4.

### 3.1 Preliminaries

We formalize our problem setting as a goal-conditioned Markov decision process (MDP) defined by the tuple  $(\mathcal{S}, \mathcal{A}, p, \mathcal{G}, \lambda)$  where  $s \in \mathcal{S}$  is the state space,  $a \in \mathcal{A}$  is the action space,  $p(s_{t+1}|s_t, a_t)$  governs the environment dynamics,  $p(s_0)$  corresponds to the initial state distribution,  $\mathcal{G} \subset \mathcal{S}$  represents the *unknown* set of goal states which is a subset of possible states, and  $\lambda$  is the discount factor. Note that this is simply a special case of a Markov decision process, where we do not have access to extrinsic reward (i.e. it is self-supervised), and where we separate the state and goal for notational clarity.

We will assume that the agent has collected an unlabeled dataset  $\mathcal{D}$  that consists of  $N$  trajectories  $[\tau_1, \dots, \tau_N]$ , and each trajectory  $\tau$  consists of a sequence of state action pairs  $[(s_0, a_0), (s_1, a_1), \dots, (s_T)]$ . We will denote the estimated distance between two states as  $C(s_t, s_g) = \|s_t - s_g\|_2^2$ , which may not accurately reflect the true distance, e.g. when states correspond to images. At test time, the agent is initialized at a start state  $s_0 \sim p(s_0)$  with a goal state  $s_g$  sampled at random from  $\mathcal{G}$ , and must minimize cost  $\mathcal{C}(s_t, s_g)$ . We assume that for any states  $s_t, s_g$  we can measure  $\mathcal{C}(s_t, s_g)$  as the distance between the states, for example in image space  $\mathcal{C}$  would be pixel distance. Success is measured as reaching within some true distance of  $s_g$ .

In the model-based RL setting we consider here, the agent aims to solve the RL problem by learning a model of the dynamics  $p_\theta(s_{t+1}|s_t, a_t)$  from experience, and using that model to plan a sequence of actions or optimize a policy.

### 3.2 Understanding the Effect of Model Error on Task Performance

A key challenge in model-based RL is that dynamics prediction error does not directly correspond to task performance. Specifically, for good task performance, certain model errors may be more costly than others, and if errors are simply distributed uniformly over dynamics predictions, errors in these critical areas may be exploited when selecting actions downstream. Intuitively, when optimizing actions for a given task, we would like our model to to give accurate predictions for actions that are important for completing the task, while the model likely does not need to be as accurate on trajectories that are completely unrelated to the task. In this section, we formalize this intuition.

Suppose the model is used by a policy to select from  $N$  action sequences  $a_{1:T}^i$ , each with expected final cost  $c_i^* = E_{p(s_{t+1}|s_t, a_t), a_{1:T}^i}[\mathcal{C}(s_T, s_g)]$ . Without loss of generality, let  $c_1^* \leq c_2^* \dots \leq c_N^*$ , i.e. the order of action sequences is sorted by their cost under the true model, which is unknown to the agent. Denote  $\hat{c}_i$  as the predicted final cost of action sequence  $a_{1:T}^i$  under the learned model, i.e.  $\hat{c}_i = E_{p_\theta(s_{t+1}|s_t, a_t), a_{1:T}^i}[\mathcal{C}(\hat{s}_T, s_g)]$ . Moreover, we consider a policy that simply selects the action sequence with lowest cost under the model:  $\hat{a} = \arg \min_{a_{1:T}^i} \hat{c}_i$ . Let the policies behavior be  $\epsilon$ -optimal if the cost of the selected action sequence  $a_{1:T}^{\hat{a}}$  has cost  $c_{\hat{a}}^* \leq c_1^* + \epsilon$ . Under this set-up, we now analyze how model error affects policy performance.

**Theorem 3.1.** *The policy will remain  $\epsilon$ -optimal, that is,*

$$c_{i'}^* \leq c_1^* + \epsilon \quad i' = \arg \min_i \hat{c}_i \quad (1)$$

*if the following two conditions are met: first, that the model prediction error on the **best** action sequence  $a_{1:T}^1$  is bounded such that*

$$|c_1^* - \hat{c}_1| < \epsilon \quad (2)$$

*and second, that the errors of sub-optimal actions sequences  $a_{1:T}^i$  are bounded by*

$$|c_i^* - \hat{c}_i| < (c_i^* - c_1^*) - \epsilon \quad \forall i \mid c_i^* > c_1^* + \epsilon \quad (3)$$

*Proof.* For the specified policy, violating  $\epsilon$ -optimality will only occur if the cost of the best action sequence  $a_{1:T}^1$  is overestimated or if the cost of a sub-optimal action sequence ( $i \mid c_i^* > c_1^* + \epsilon$ ) is underestimated. Thus, let us define the "worst case" cost predictions as the ones for which  $c_1^*$  is most overestimated and  $c_i^* \quad \forall i \mid c_i^* > c_1^* + \epsilon$  are most underestimated (while still satisfying Equations 2 and 3). Concretely we write the worst case cost estimates as

$$\tilde{c}_i := \min \hat{c}_i \quad \forall i \mid c_i^* > c_1^* + \epsilon$$

$$\tilde{c}_1 := \max \hat{c}_1$$

s.t. Eq. 2 and 3 hold. We will now show that  $\tilde{c}_1 < \tilde{c}_i \quad \forall i \mid c_i^* > c_1^* + \epsilon$ . First, since  $\tilde{c}_i$  satisfies Eq. 3, we have that

$$\tilde{c}_i > c_i^* - (c_i^* - c_1^*) + \epsilon$$

Similarly, since  $\tilde{c}_1$  satisfies Eq. 2, we have that

$$\tilde{c}_1 < c_1^* + \epsilon$$

Substituting, we see that

$$\tilde{c}_i > c_i^* - (c_i^* - c_1^*) + \epsilon = c_1^* + \epsilon > \tilde{c}_1 \quad \forall i \mid c_i^* > c_1^* + \epsilon \quad (4)$$

Hence even in the worst case, Equations 2 and 3 ensure that  $\hat{c}_i > \hat{c}_1 \quad \forall i \mid c_i^* > c_1^* + \epsilon$ , and thus no action sequence  $i$  for which  $c_i^* > c_1^* + \epsilon$  will be selected and the policy will remain  $\epsilon$ -optimal. Note that action sequences besides  $i = 1$  for which  $c_i^* \leq c_1^* + \epsilon$  costs are unbounded, as it is ok for them to be significantly underestimated since selecting them still allows the policy to be  $\epsilon$ -optimal.  $\square$

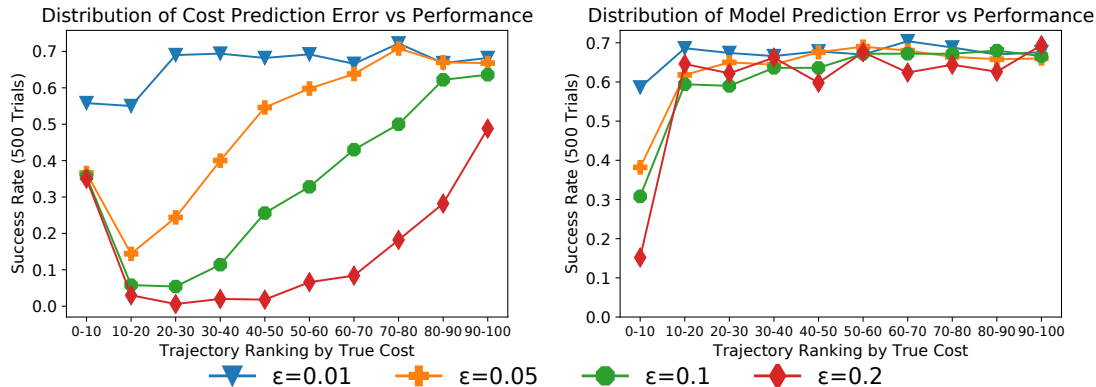


Figure 1: **Distribution of model errors vs. performance:** We validate how the distribution of model errors affects performance on a simple 2D navigation domain, by adding noise to cost predictions (**left**) or model predictions (**right**). We add varying amounts of noise with magnitude up to  $\epsilon$  to the predictions of the 10 lowest true cost trajectories (**0-10**) to the 10 highest true cost trajectories (**90-100**). We observe that adding noise to low true cost trajectories dramatically reduces performance, while adding noise to the high true cost trajectories has no nearly no impact on performance.

Theorem 3.1 suggests that, for good task performance, model error must be low for good trajectories, and we can afford higher model error for trajectories with higher cost. That is, **the greater the trajectory cost, the more model error we can afford**. Specifically, we see that the allowable error bound on cost of an action sequence from a learned model scales linearly with how far from optimal that action sequence is, in order to maintain the optimal policy for the downstream task. Note, that while Theorem 3.1 relates cost prediction error (not explicitly dynamics prediction error) to planning performance, we can expect dynamics prediction error to relate to the resulting cost prediction error. We also verify this empirically in the next section.

### 3.3 Verifying Theorem 3.1 Experimentally

We now verify the above analysis through a controlled study of how prediction error affects task performance. To do so, we will use the true model of an environment and true cost of an action sequence for planning, but will artificially add noise to the cost/model predictions to generate model error.

Consider a 2 dimensional navigation task, where the agent is initialized at  $s_0 = [0.5, 0.5]$  and is randomly assigned a goal  $s_g \in [0, 1]^2$ . Assume we have access to the underlying model of the environment, and cost defined as  $\mathcal{C}(s_t, s_g) = \|s_t - s_g\|_2$ . We can run the policy described in Section 3.2, specifically sampling  $N = 100$  action sequences, and selecting the one with lowest predicted cost, where we consider 2 cases: (1) predicted cost is using the true model, but with noise  $\alpha$  added to the true cost  $\hat{c}_i = c_i^* + \alpha$  of some subset of action sequences, and (2) predicted cost is true cost, but with noise  $\alpha$  added to the model predictions  $s_{t+1} = \bar{s}_{t+1} + \alpha$  where  $\bar{s}_{t+1} \sim p(s_{t+1}|s_t, a_t)$  of some subset of action sequences. The first case relates directly to Theorem 3.1, while the second case relates to what we can control when training a self-supervised dynamics model. When selectively adding noise, we will use uniform

noise  $\alpha \sim \mathcal{U}(-\epsilon, \epsilon)$ . We specifically study the difference in task performance when adding noise  $\alpha$  to model predictions for the first 10% of trajectories with lowest true cost, the second 10% lowest true cost trajectories, etc., up to the 10% of trajectories with highest true cost. Here “true cost” refers to the cost of the action sequence under the true model and cost function without noise. For each noise augmented model we measure the task performance, specifically the success rate (reaching within 0.1 of the goal), over 500 random trials.

We see in Figure 1 that for multiple values of noise  $\epsilon$ , when adding noise to the better (lower true cost) trajectories we see a significant drop in task performance, while when adding noise to the worse (higher true cost) trajectories task performance remains relatively unchanged (except for the case with very large  $\epsilon$ ). In particular, we notice that when adding noise to cost predictions, performance scales almost linearly as we add noise to worse trajectories. Note there is one exception to this trend: if we add noise only to the top 10% of trajectories, performance is not optimal, but reasonable because the best few trajectories will occasionally be assigned a lower cost under the noise model.

In the case of model error, we see a much steeper increase in performance, where adding model error to the best 10 trajectories significantly hurts performance, while adding to the others does not. This is because, in this environment, noise added to model predictions generally makes the cost of those predictions worse; so if no noise is added to the best trajectories, the best action sequence is still likely to be selected. The exact relationship between model prediction error and cost prediction error depends on the domain and task. But, we can see that in both cases in Figure 1, the conclusion from Theorem 3.1 holds true: accuracy on good action sequences matters much more than accuracy on bad action sequences.



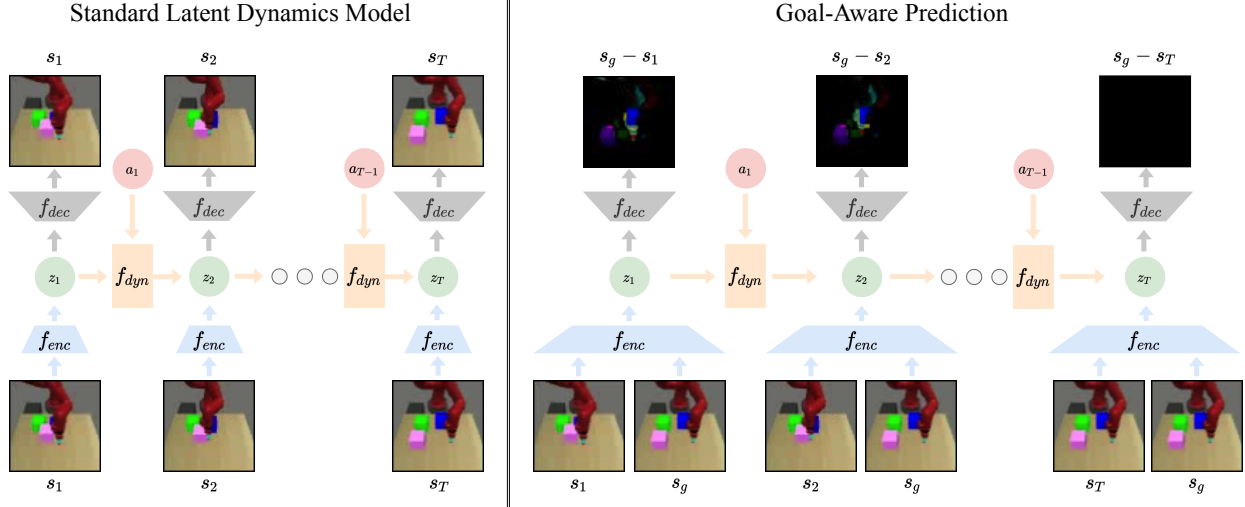


Figure 2: **Goal-Aware Prediction:** Compared to a standard latent dynamics model (**left**), our proposed method, goal-aware prediction (GAP), (**right**) encodes both the current state  $s_t$  and goal  $s_g$  into a single latent space  $z_t$ . Samples from the distribution of  $z_t$  are then used to reconstruct the residual between the current state and goal  $s_g - s_t$ . Simultaneously, we learn the forward dynamics in the latent space  $z$ , specifically, learning to predict  $z_{t+1}$  from  $z_t$  and  $a_t$ . Using this approach, we obtain 2 favorable properties: (1) the latent space only needs to capture components of the scene relevant to the goal, and (2) the prediction task becomes easier (the residual approaches 0) for states closer to the goal.

### 3.4 Redistributing Model Errors with Goal Aware Prediction

Our analysis above suggests that distributing errors uniformly across action sequences will not lead to good task performance. Yet, standard model learning objectives will encourage just that. In this section, we aim to change our model learning approach in an aim to redistribute errors more appropriately.

Ideally, we would like to encourage the model to have more accurate predictions on the trajectories which are relevant to the task. However, actually identifying how relevant a trajectory is to a specific goal  $s_g$  can be challenging.

One potential approach to doing this would be to re-weight the training loss of the model on transitions  $p_\theta(s_{t+1}|s_t, a_t)$  *inversely* by the cost  $\mathcal{C}(s_t, s_g)$ , such that low cost trajectories are weighted more heavily. While this simple approach may work when the cost function  $\mathcal{C}(s_t, s_g)$  is accurate, the distance metric  $\mathcal{C}(s_t, s_g)$  for high-dimensional states is often sparse and not particularly meaningful; when states are images,  $\mathcal{C}$  amounts to  $\ell_2$  distance in pixel space.

An alternative way of approaching this problem is, rather than focusing on how to re-weight the model’s predictions, instead ask, “what exactly should the model be predicting?” If the downstream task involves using the model to plan to reach a goal state, then intuitively the model should only need to focus on predicting goal relevant parts of the scene. Moreover, if the model is trained to focus on parts of the scene relevant to the goal, it will naturally be biased towards higher accuracy in states relevant to the task, re-distributing model error favourably for downstream performance.

To that end, we propose goal-aware prediction (GAP) as a technique to re-distribute model error by learning a model that, in addition to the current state and action,  $s_t$  and  $a_t$ , is conditioned on the goal state  $s_g$ , and instead of reconstructing the future state  $s_{t+1}$ , reconstructs the difference between the future state and the goal state, that is:  $p_\theta((s_g - s_{t+1})|s_t, s_g, a_t)$ . Critically, to train GAP effectively, we need action sequences that are relevant to the corresponding goal. To accomplish this, we can choose to set the goal state for a given action sequence as the final state of that trajectory, i.e. using hindsight relabeling (Andrychowicz et al., 2017). Specifically, given a trajectory  $[(s_1, a_1), (s_2, a_2), \dots, (s_T)]$ , the goal is assigned to be the last state in the trajectory  $s_g = s_T$ , and for all states  $\{s_t | 1 \leq t \leq T - 1\}$ ,  $p_\theta(s_t, s_g, a_t)$  is trained to reconstruct the delta to the goal  $s_g - s_{t+1}$ .

Our proposed GAP method has two clear benefits over standard dynamics models. First, assuming that the agent is not in a highly dynamic scene with significant background distractor motion, by modeling the delta between  $s_g$  and  $s_t$ ,  $p_\theta$  only needs to model components of the state which are relevant to the current goal. This is particularly important in high dimensional settings where there may be large components of the state which are irrelevant to the task, and need not be modeled. Second, states  $s_t$  that are temporally close to the goal state  $s_g$  will have a smaller delta  $s_g - s_t$ , approaching zero along the trajectory until  $s_t = s_g$ . As a result, states closer to the goal will be easier to predict, biasing the model towards low error near states relevant to the goal. In light of our analysis of model error in the previous sections, we hypothesize that this model will lead to

better downstream task performance compared to a standard model that distributes errors uniformly across trajectories.

*When do we expect GAP to improve performance on downstream tasks?* We expect GAP to be most effective when the goals involve changing a subset  $d < D$  state dimensions from the initial states  $s_t \in \mathbb{R}_D$ . Under these conditions, GAP only needs to predict the dynamics of the  $d$  dimensions, while standard latent dynamics need to predicts all  $D$  making GAP an easier problem.

## 4 Implementing Goal-Aware Prediction

We implement GAP with a latent dynamics model, as shown in Figure 2. Given a dataset of trajectories  $[\tau_1, \dots, \tau_N]$ , we sample sequences of states  $[(s_1, a_1), \dots, (s_T)]$  where we re-label goal for the trajectory as  $s_g = s_T$ .

The GAP model consists of three components, (1) an encoder  $f_{enc}(z_t|s_t, s_g; \theta_{enc})$  that encodes the state  $s_t$  and goal  $s_g$  into a latent space  $z_t$ , (2) a decoder  $f_{dec}(s_g - s_t|z_t; \theta_{dec})$  that decodes samples from the latent distribution into  $s_g - s_t$ , and (3) a forward dynamics model in the latent space  $f_{dyn}(z_{t+1}|z_t, a_t; \theta_{dyn})$  which learns to predict the future latent distribution over  $z_{t+1}$  from  $z_t$  and action  $a_t$ . In our experiments we work in the setting where states are images, so  $f_{enc}(z_t|s_t, s_g)$  and  $f_{dec}(s_g - s_t|z_t)$  are convolutional neural networks, and  $f_{dyn}(z_{t+1}|z_t, a_t)$  is a fully-connected network. The full set of parameters  $\theta = \{\theta_{enc}, \theta_{dec}, \theta_{dyn}\}$  are jointly optimized. Exact architecture and training details for all modules can be found in the supplement. Following prior works (Finn et al., 2016; Finn and Levine, 2017; Amos et al., 2018a), we train for multi-step prediction. More specifically, given  $s_t, a_{t:t+H}, s_g$ , the model is trained to reconstruct  $(s_g - s_t), \dots, (s_g - s_{t+H})$ , shown in Figure 2.

**Data Collection and Model Training:** In our *self-supervised* setting, data collection simply corresponds to rolling out a random exploration policy in the environment. Specifically, we sample uniformly from the agent’s action space, and collect 2000 episodes, each of length 50, for a total of 100,000 frames of data.

During training, sub-trajectories of length 30 time steps are sampled from the data set, with the last timestep labeled as the goal  $s_g = s_{30}$ . Depending on the current value of  $H$ , loss is computed over  $H$  step predictions starting from states  $s_{t:(t+H)}$ . We use a curriculum when training all models, where  $H$  starts at 0, and is incremented by 1 every 50,000 training iterations. All models are trained to convergence, for about 300,000 iterations on the same dataset.

**Planning with GAP:** For all trained models, when given a new goal at test time  $s_g$ , we plan using model predictive control (MPC) in the latent space of the model. Specifically, both the current state  $s_t$  and  $s_g$  are encoded into their respective latent spaces  $z_t$  and  $z_g$  (Algorithm 1, Line 3).

---

### Algorithm 1 Latent MPC( $f_{enc}, f_{dyn}, s_t, s_g$ )

---

```

1: Let  $D = 1000, D^* = 10, H = 15$ 
2: Receive current state  $s_t$  and goal state  $s_g$ 
3: Encode  $z_t \sim f_{enc}(s_t, s_g), z_g \sim f_{enc}(s_g, s_g)$ 
4: Initialize  $N(\mu, \sigma^2) = \mathcal{N}(0, 1)$ 
5: Let cost function  $\mathcal{C}(z_i, z_j) = \|z_i - z_j\|_2^2$ 
6: while iterations  $\leq 3$  do
7:    $a_{t:H}^1, \dots, a_{t:H}^D \sim N(\mu, \sigma^2)$ 
8:    $z_{t+1:t+H}^1, \dots, z_{t+1:t+H}^D \sim f_{dyn}(z_t, a_{t:H}^1, \dots, a_{t:H}^D)$ 
9:    $\hat{c}^1, \dots, \hat{c}^D = [\sum_{h=1}^H \mathcal{C}(z_{t+h}^1, z_g), \dots, \sum_{h=1}^H \mathcal{C}(z_{t+h}^D, z_g)]$ 
10:   $a_{sorted} = \text{Sort}([a_{t:H}^1, \dots, a_{t:H}^D])$  by  $\hat{c}$ 
11:  Refit  $\mu, \sigma^2$  to  $a_{sorted}[1 : D^*]$ 
12: end while
13: Return  $\hat{c}_{sorted}[1], a_{sorted}[1]$ 

```

---

Then using the model  $f_{dyn}(z_{t+1}|z_t, a_t)$ , the agent plans a sequence of  $H$  actions to minimize cost  $\sum_{h=0}^H \|z_g - z_{t+h}\|_2^2$  (Algorithm 1, Lines 4-11). Following prior works (Finn and Levine, 2017; Hafner et al., 2019b), we use the cross-entropy method (Rubinstein and Kroese, 2004) as the planning optimizer. Finally, the best sequence of actions is returned and executed in the environment (Algorithm 1, Line 13).

While executing the plan, our model re-plans every  $H$  timesteps. That is, it starts at state  $s_t$ , uses Latent MPC (Algorithm 1) to first plan a sequence of  $H$  actions, executes them in the environment resulting in a state  $s_{t+H}$ , then re-plans an additional  $H$  actions, and executes them resulting in a final state  $s_T$ . Success is computed based the difference between  $s_T$  and  $s_g$ .

## 5 Experiments

In our experiments, we investigate three primary questions

(1) Does using our proposed technique for goal-aware prediction (GAP) re-distribute model error such that predictions are more accurate on good trajectories?


(2) Does re-distributing model errors using GAP result in better performance in downstream tasks?

(3) Can GAP be combined with large video prediction models to scale to the complexity of real world images?

We design our experimental set-up with these questions in mind in Section 5.1, then examine each of the questions in Sections 5.2, 5.3, and 5.4 respectively.

### 5.1 Experimental Domains and Comparisons


**Experimental Domains:** Our primary experimental domain is a simulated tabletop manipulation task built off of the Meta-World suite of environments (Yu et al., 2019a). Specifically, it consists of a simulated Sawyer robot, and 3 blocks on a tabletop. In the *self-supervised* data collection phase, the agent executes a random policy for 2,000 episodes, collecting 100,000 frames worth of data. Then, after learning a model, the agent is tested on 4 previously unseen tasks, where the task is specified by a goal image.



**Figure 3: Distribution of Model Errors:** We examine the distribution of model prediction errors of GAP compared to prior methods over 1000 random action sequences, evaluated on the “Task 1” domain. The y-axis corresponds to model mean-squared error (with standard error bars), and the x-axis corresponds to number of time steps predicted forward. Naturally, we observe that model error increases as the prediction horizon increases, for all approaches. However, although all approaches have a similar error over all 1000 action sequences (left), GAP achieves significantly lower error on the *best* 10 trajectories (right). This suggests that changing the model objective through predicting the goal-state residual leads to more accurate predictions in areas that matter in downstream tasks.

**Task 1** consists of pushing the green, pink, or blue block to a goal position, while the more challenging **Task 2** requires the robot to push 2 blocks to each of their respective goal positions (see Figure 4). Task success is defined as being within 0.1 of the goal positions. **Task 3 and 4** involve closing and opening the door respectively with distractor objects on the table, where success is defined as being within  $\pi/6$  radians of the goal position. The agent receives  $64 \times 64$  RGB camera observations of the tabletop. We also study model error on real robot data from the BAIR Robot Dataset (Ebert et al., 2017) and RoboNet dataset (Dasari et al., 2019) in Section 5.4.

**Comparisons:** We compare to several model variants in our experiments. **GAP** is our approach of learning dynamics in a latent space conditioned on the current state and goal, and reconstructing the residual between the current state and goal state, as described in Section 3.4. **GAP (-Goal Cond)** is an ablation of GAP that does not use goal conditioning. Instead of conditioning on the goal and predicting the residual to the goal, it is conditioned on the initial state, and predicts the residual to the initial state. This is representative of prior works (e.g. Nagabandi et al. (2019)) that predict residuals for model-based RL. **GAP (-Residual)** is



**Figure 4: Evaluation Tasks:** Sample initial & goal states for each of the simulated manipulation tasks. Tasks involve manipulating blocks or a door, with the task specified by a goal image.

another ablation of GAP that is conditioned on the goal but maintains the standard reconstruction objective instead of the residual. This is similar to prior work on goal conditioned video prediction (Rybkin et al., 2020). **Standard** refers to a standard latent dynamics model, representative of approaches such as PlaNet (Hafner et al., 2019b), but without reward prediction since we are in the self-supervised setting.

When studying task performance, we also compare to two alternative self-supervised reinforcement learning approaches. First, we compare to an **Inverse Model**, which is a latent dynamics model where the latent space is learned via an action prediction loss (instead of image reconstruction), as done in Pathak et al. (2017). Second, we compare to a model-free approach: reinforcement learning with imagined goals (**RIG**) (Nair et al., 2018), where we train a VAE on the same pre-collected dataset as the other models, then train a policy in the latent space of the VAE to reach goals sampled from the VAE. Further implementation details can be found in the supplement.

## 5.2 Experiment 1: Does GAP Favorably Redistribute Model Error?

In our first set of experiments, we study how GAP affects the distribution of model errors, and if it leads to lower model error on task relevant trajectories. We sample 1000 random action sequences of length 15 in the Task 1 domain. We compute the true next states  $s_{1:H}^1, \dots, s_{1:H}^{1000}$  and costs  $c^1, \dots, c^{1000}$  for each action sequence by feeding it through the true simulation environment. We then get the predicted next states from our learned models, including GAP as well as the comparisons outlined above. We then examine the model error of each approach, and how it changes when looking at all trajectories, versus the lowest cost trajectories.

We present our analysis in Figure 3. We specifically look at the model error on all 1000 action sequences, the top 100

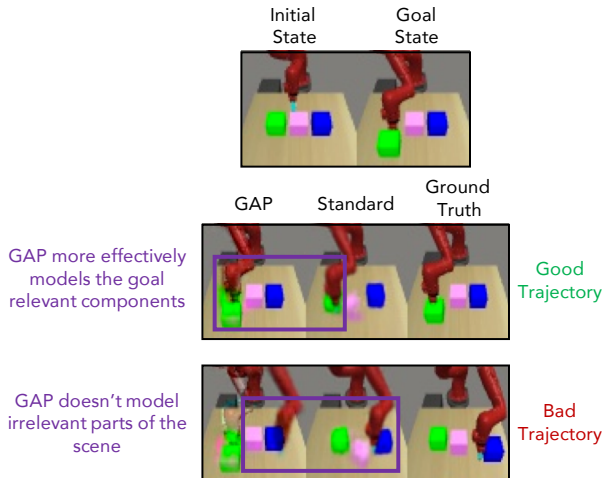


Figure 5: **GAP Predictions on Good/Bad Trajectories.** Here we show qualitatively how GAP focuses on the task relevant parts of the scene. Note, for GAP predictions we add back the goal image to the predicted goal-state residual. Given the task specified by pushing the green block (**top**), consider a good action sequence (**middle**) and bad action sequence (**bottom**). On the good action sequence GAP more effectively models the goal relevant parts of the scene (the green block) than the standard model. Additionally, on the bad trajectory, GAP ignores the irrelevant objects and does not model their dynamics at all, while the standard model does.

action sequences, and the top 10 action sequences. First, we observe that model error increases with the prediction horizon, which is expected due to compounding model error. More interestingly, however, we observe that while our proposed GAP approach has the highest error averaged across all 1000 action sequences, it has by far the lowest error on the top 10. This suggests that the goal conditioned prediction of the goal-state residual indeed encourages low model error in the relevant parts of the state space. Furthermore, we see that the conditioning on and reconstructing the difference to the *actual goal* is in fact critical, as the ablation GAP (-Goal Cond) which instead is conditioned on and predicts the residual to the first frame actually gets worse error on the lowest cost trajectories.

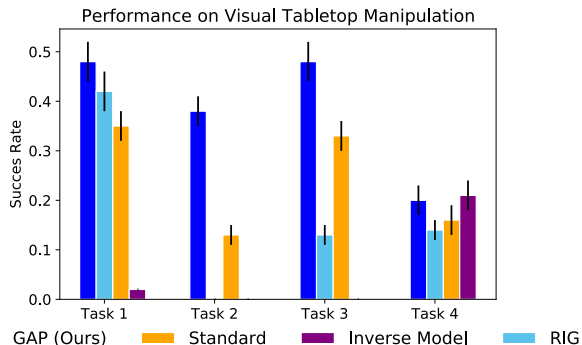


Figure 6: **Success rate on tabletop manipulation.** On the tasks proposed in Section 5.1, we find that GAP outperforms the comparisons. Specifically on the harder 2 block manipulation task, GAP has a significantly higher success rate.

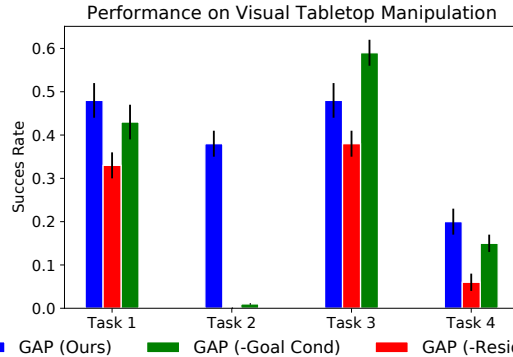


Figure 7: **Success rate on tabletop manipulation (ablation).** We compare the success rate of GAP to ablations on the task described in Section 5.1. We find that in all tasks except Task 3 both goal conditioning and residual are important for good performance.

This indicates that GAP successfully re-distributes error such that it has the most accurate predictions on low-cost trajectories. We also observe this qualitatively in Figure 5. For a given initial state and goal state from Task 1, GAP effectively models the target object (the green block) on a good action sequence that reaches the goal, while the standard model struggles. On a poor action sequence that hits the non-target blocks, the Standard approach models them, while **GAP does not model interaction with these blocks at all**, suggesting that GAP does not model irrelevant parts of the scene. In the next section, we examine if this error redistribution translates to better task performance.

### 5.3 Experiment 2: Does GAP Lead to Better Downstream Task Performance?

To study downstream task performance, we test on the tabletop manipulation tasks described in Section 5.1. We perform planning over 30 timesteps with the learned models as described in Section 3, and report the final success rate of each task over 200 trials in Figure 6. We see that in all tasks GAP outperforms the comparisons, especially in the most challenging 2 block manipulation task in Task 2 (where precise modeling of the relevant objects is especially important).

We make a similar comparison, but to the ablations of GAP, in Figure 7. Again we see that GAP is performant, and in all tasks except Task 3 both goal conditioning and residual are important for good performance. Interestingly, we observe that GAP (-Goal Cond) is competitive on the door manipulation tasks.

Hence, we can conclude that GAP not only enables lower model error in task relevant states, but by doing so, also achieves a 10-20% absolute performance improvement over baselines on 3 out of 4 tasks.

### 5.4 Experiment 3: Does GAP scale to real, cluttered visual scenes?

Lastly, we study whether our proposed GAP method extends to real, cluttered visual scenes. To do so we combine it with






Figure 8: **GAP+SVG Video Prediction (BAIR Robot Dataset):** Here we present qualitative examples of action-conditioned SVG with and without GAP on the BAIR robot dataset, predicting on goal-reaching trajectories. Note, in the GAP predictions the goal is added back to the predicted goal-state residual. In this case the goal is the rightmost frame. We see that GAP is able to more accurately predict the objects relevant to the goal, for example the small spoon highlighted in the red box.

an action-conditioned version of the video generation model, SVG (Denton and Fergus, 2018). Specifically, we condition the SVG encoder on the goal, and the current goal-state residual, and predict the next goal-state residual.

We compare the prediction error of SVG to SVG+GAP on goal reaching trajectories (Figure 9) from real robot datasets, namely the BAIR Robot Dataset (Ebert et al., 2017) and the RoboNet Dataset (Dasari et al., 2019). We see that action-conditioned SVG combined with GAP as well as the ablation without residual prediction both have lower prediction error on goal reaching trajectories than standard action-conditioned SVG.

Qualitatively, we also observe that SVG+GAP is able to more effectively capture goal relevant components, as shown in Figure 8. We see that GAP is able to capture the motion of the small spoon, while correctly modeling the dynamics of the arm, while SVG ignores the spoon.

As a result, we conclude that GAP can effectively be combined with large video prediction models, and scaled to challenging real visual scenes.

## 6 Discussion and Limitations

In this paper, we studied the role of model error in task performance. Motivated by our analysis, we proposed goal-aware prediction, a self-supervised framework for learning dynamics that are conditioned on a goal and structured in a way that favorably re-distributes model error to be low in goal-relevant states. In visual control domains, we verified that GAP (1) enables lower model error on task relevant states, (2) improves downstream task performance, and (3) scales to real, cluttered visual scenes.

While GAP demonstrated significant gains, multiple limitations and open questions remain. Our theoretical analysis suggests that we should re-distribute model errors according to their planning cost. While GAP provides one way to do that in a *self-supervised* manner, there are likely many other approaches that can be informed by our analysis, including




Figure 9: **Model Errors (Real Robot Data):** We examine the model error of SVG combined with GAP on **unseen, goal-reaching trajectories** from two real robot datasets (the BAIR Dataset (Ebert et al., 2017) and the RoboNet Dataset (Dasari et al., 2019)). We see that action-conditioned SVG combined with GAP has lower prediction error on the goal reaching trajectories than standard action-conditioned SVG. We observe that the GAP ablation which also conditions on the goals, but predicts residuals is equally effective in this setting.

approaches that leverage human supervision. For example, we anticipate that GAP-based models may be less suitable for environments with dynamic distractors such as changing lighting conditions and moving distractor objects, since GAP would be still encouraged to model these events. To effectively solve this case, an agent would likely require human supervision to indicate the axes of variation that are relevant to the goal. Incorporating such supervision is outside the scope of this work, but an exciting avenue for future investigation.

Additionally, while in this work we found GAP to work well with goals selected at the end of sampled trajectories, there may be more effective ways to sample goals. Studying the relationship between how exactly goals are sampled and learning performance, as well as how best to sample and re-label goals is an exciting direction for future work.

## Acknowledgments

We would like to thank Ashwin Balakrishna, Oleh Rybkin, and members of the IRIS lab for many valuable discussions. This work was supported in part by Schmidt Futures and an NSF graduate fellowship. Chelsea Finn is a CIFAR Fellow in the Learning in Machines & Brains program.

## References

- P. Agrawal, A. V. Nair, P. Abbeel, J. Malik, and S. Levine. Learning to poke by poking: Experiential learning of intuitive physics. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 5074–5082. Curran Associates, Inc., 2016.
- B. Amos, L. Dinh, S. Cabi, T. Rothörl, A. Muldal, T. Erez, Y. Tassa, N. de Freitas, and M. Denil. Learning awareness models. In *International Conference on Learning Representations*, 2018a. URL <https://openreview.net/forum?id=r1HhRfWRZ>.
- B. Amos, I. Jimenez, J. Sacks, B. Boots, and J. Z. Kolter. Differentiable mpc for end-to-end planning and control. In *Advances in Neural Information Processing Systems*, pages 8289–8300, 2018b.
- M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel, and W. Zaremba. Hindsight experience replay. *CoRR*, abs/1707.01495, 2017.
- M. Babaeizadeh, C. Finn, D. Erhan, R. H. Campbell, and S. Levine. Stochastic variational video prediction. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=rk49Mg-CW>.
- E. Banijamali, R. Shu, M. Ghavamzadeh, H. H. Bui, and A. Ghodsi. Robust locally-linear controllable embedding. *ArXiv*, abs/1710.05373, 2017.
- A. Byravan, J. T. Springenberg, A. Abdolmaleki, R. Hafner, M. Neunert, T. Lampe, N. Siegel, N. M. O. Heess, and M. A. Riedmiller. Imagined value gradients: Model-based policy optimization with transferable latent dynamics models. *ArXiv*, abs/1910.04142, 2019.
- K. Chua, R. Calandra, R. McAllister, and S. Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, pages 4754–4765, 2018.
- F. Codevilla, M. Müller, A. Dosovitskiy, A. López, and V. Koltun. End-to-end driving via conditional imitation learning. *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–9, 2017.
- S. Dasari, F. Ebert, S. Tian, S. Nair, B. Bucher, K. Schmeckpeper, S. Singh, S. Levine, and C. Finn. Robonet: Large-scale multi-robot learning. *ArXiv*, abs/1910.11215, 2019.
- M. Deisenroth and C. E. Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pages 465–472, 2011.
- E. L. Denton and R. Fergus. Stochastic video generation with a learned prior. In *International Conference on Machine Learning*, 2018.
- P. D’Oro, A. M. Metelli, A. Tirinzoni, M. Papini, and M. Restelli. Gradient-aware model-based policy search. In *Thirty-Fourth AAAI Conference on Artificial Intelligence*, 2020.
- A. Dosovitskiy and V. Koltun. Learning to act by predicting the future. *ArXiv*, abs/1611.01779, 2016.
- F. Ebert, C. Finn, A. X. Lee, and S. Levine. Self-supervised visual planning with temporal skip connections. *CoRR*, abs/1710.05268, 2017.
- F. Ebert, C. Finn, S. Dasari, A. Xie, A. X. Lee, and S. Levine. Visual foresight: Model-based deep reinforcement learning for vision-based robotic control. *CoRR*, abs/1812.00568, 2018.
- B. Eysenbach, R. R. Salakhutdinov, and S. Levine. Search on the replay buffer: Bridging planning and reinforcement learning. In *Advances in Neural Information Processing Systems*, pages 15246–15257, 2019.
- K. Fang, Y. Zhu, A. Garg, S. Savarese, and L. Fei-Fei. Dynamics learning with cascaded variational inference for multi-step manipulation. *ArXiv*, abs/1910.13395, 2019.
- A.-M. Farahmand, A. Barreto, and D. Nikovski. Value-Aware Loss Function for Model-based Reinforcement Learning. In A. Singh and J. Zhu, editors, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 1486–1494, Fort Lauderdale, FL, USA, 20–22 Apr 2017. PMLR. URL <http://proceedings.mlr.press/v54/farahmand17a.html>.
- C. Finn and S. Levine. Deep visual foresight for planning robot motion. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2786–2793. IEEE, 2017.
- C. Finn, I. J. Goodfellow, and S. Levine. Unsupervised learning for physical interaction through video prediction. In *NIPS*, 2016.

- D. Freeman, D. Ha, and L. Metz. Learning to predict without looking ahead: World models without forward prediction. In *Advances in Neural Information Processing Systems*, pages 5379–5390, 2019.
- C. Gelada, S. Kumar, J. Buckman, O. Nachum, and M. G. Bellemare. Deepmdp: Learning continuous latent space models for representation learning. *CoRR*, abs/1906.02736, 2019. URL <http://arxiv.org/abs/1906.02736>.
- K. Gregor, D. J. Rezende, F. Besse, Y. Wu, H. Merzic, and A. van den Oord. Shaping belief states with generative environment models for rl. In *NeurIPS*, 2019.
- D. Ha and J. Schmidhuber. World models. *ArXiv*, abs/1803.10122, 2018.
- D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019a.
- D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson. Learning latent dynamics for planning from pixels. In *International Conference on Machine Learning*, pages 2555–2565, 2019b.
- K. Hartikainen, X. Geng, T. Haarnoja, and S. Levine. Dynamical distance learning for semi-supervised and unsupervised skill discovery. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=H1lmhaVtvr>.
- A. Havens, Y. Ouyang, P. Nagarajan, and Y. Fujita. Learning latent state spaces for planning through reward prediction, 2020. URL <https://openreview.net/forum?id=ByxJj1HKwr>.
- B. Ichter and M. Pavone. Robot motion planning in learned latent spaces. *IEEE Robotics and Automation Letters*, 4(3):2407–2414, 2019.
- M. Janner, J. Fu, M. Zhang, and S. Levine. When to trust your model: Model-based policy optimization. In *NeurIPS*, 2019.
- L. P. Kaelbling. Learning to achieve goals. In *IJCAI*, pages 1094–1098, 1993.
- D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, and S. Levine. Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation. *arxiv:Preprint*, 2018.
- T. Kurutach, A. Tamar, G. Yang, S. J. Russell, and P. Abbeel. Learning plannable representations with causal infogan. In *Advances in Neural Information Processing Systems*, pages 8733–8744, 2018.
- N. G. Lambert, B. Amos, O. Yadan, and R. Calandra. Objective mismatch in model-based reinforcement learning. *ArXiv*, abs/2002.04523, 2020.
- A. X. Lee, R. Zhang, F. Ebert, P. Abbeel, C. Finn, and S. Levine. Stochastic adversarial video prediction. *CoRR*, abs/1804.01523, 2018.
- A. X. Lee, A. Nagabandi, P. Abbeel, and S. Levine. Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model. *ArXiv*, abs/1907.00953, 2019.
- S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.
- K. Liu, T. Kurutach, P. Abbeel, and A. Tamar. Hallucinative topological memory for zero-shot visual planning, 2020. URL <https://openreview.net/forum?id=BkgF4kSFPB>.
- R. McAllister and C. E. Rasmussen. Improving pilco with bayesian neural network dynamics models. 2016.
- V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. A. 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*, 518:529–533, 2015.
- A. Nagabandi, K. Konoglie, S. Levine, and V. Kumar. Deep dynamics models for learning dexterous manipulation. *ArXiv*, abs/1909.11652, 2019.
- A. Nair, S. Bahl, A. Khazatsky, V. Pong, G. Berseth, and S. Levine. Contextual imagined goals for self-supervised robotic learning. *ArXiv*, abs/1910.11670, 2019.
- A. V. Nair, V. Pong, M. Dalal, S. Bahl, S. Lin, and S. Levine. Visual reinforcement learning with imagined goals. In *Advances in Neural Information Processing Systems*, pages 9191–9200, 2018.
- S. Nair and C. Finn. Hierarchical foresight: Self-supervised learning of long-horizon tasks via visual subgoal generation. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=H1gzR2VKDH>.
- OpenAI. Openai five. <https://blog.openai.com/openai-five/>, 2018.
- OpenAI, M. Andrychowicz, B. Baker, M. Chociej, R. Józefowicz, B. McGrew, J. W. Pachocki, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, J. Schneider, S. Sidor, J. Tobin, P. Welinder, L. Weng, and

- W. Zaremba. Learning dexterous in-hand manipulation. *CoRR*, abs/1808.00177, 2018. URL <http://arxiv.org/abs/1808.00177>.
- D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell. Curiosity-driven exploration by self-supervised prediction. *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 488–489, 2017.
- C. Paxton, Y. Barnoy, K. D. Katyal, R. Arora, and G. D. Hager. Visual robot task planning. *CoRR*, abs/1804.00062, 2018.
- L. Pinto and A. Gupta. Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 3406–3413. IEEE, 2016.
- S. Racanière, T. Weber, D. P. Reichert, L. Buesing, A. Guez, D. J. Rezende, A. P. Badia, O. Vinyals, N. M. O. Heess, Y. Li, R. Pascanu, P. W. Battaglia, D. Hassabis, D. Silver, and D. Wierstra. Imagination-augmented agents for deep reinforcement learning. *ArXiv*, abs/1707.06203, 2017.
- R. Rubinstein and D. Kroese. *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization, Monte-Carlo Simulation and Machine Learning*. 01 2004.
- O. Rybkin, K. Pertsch, F. Ebert, D. Jayaraman, C. Finn, and S. Levine. Goal-conditioned video prediction, 2020. URL <https://openreview.net/forum?id=Blg79grKPr>.
- T. Schaul, D. Horgan, K. Gregor, and D. Silver. Universal value function approximators. In *International Conference on Machine Learning*, 2015.
- J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel, T. P. Lillicrap, and D. Silver. Mastering atari, go, chess and shogi by planning with a learned model. *ArXiv*, abs/1911.08265, 2019.
- 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, T. Graepel, and D. Hassabis. Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–503, 2016. URL <http://www.nature.com/nature/journal/v529/n7587/full/nature16961.html>.
- A. Srinivas, A. Jabri, P. Abbeel, S. Levine, and C. Finn. Universal planning networks. *ArXiv*, abs/1804.00645, 2018.
- R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, second edition, 2018. URL <http://incompleteideas.net/book/the-book-2nd.html>.
- B. Thananjeyan, A. Balakrishna, U. Rosolia, F. Li, R. McAllister, J. E. Gonzalez, S. Levine, F. Borrelli, and K. Goldberg. Safety augmented value estimation from demonstrations (saved): Safe deep model-based rl for sparse cost robotic tasks. 2019.
- R. Veerapaneni, J. D. Co-Reyes, M. Chang, M. Janner, C. Finn, J. Wu, J. B. Tenenbaum, and S. Levine. Entity abstraction in visual model-based reinforcement learning. *ArXiv*, abs/1910.12827, 2019.
- R. Villegas, A. Pathak, H. Kannan, D. Erhan, Q. Le, and H. Lee. High fidelity video prediction with large stochastic recurrent neural networks. 11 2019.
- A. Wang, T. Kurutach, K. Liu, P. Abbeel, and A. Tamar. Learning robotic manipulation through visual planning and acting. *CoRR*, abs/1905.04411, 2019.
- T. Wang and J. Ba. Exploring model-based planning with policy networks. *ArXiv*, abs/1906.08649, 2019.
- M. Watter, J. T. Springenberg, J. Boedecker, and M. A. Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. *CoRR*, abs/1506.07365, 2015.
- A. Xie, F. Ebert, S. Levine, and C. Finn. Improvisation through physical understanding: Using novel objects as tools with visual foresight. *CoRR*, abs/1904.05538, 2019.
- T. Yu, D. Quillen, Z. He, R. R. Julian, K. Hausman, C. Finn, and S. Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. *ArXiv*, abs/1910.10897, 2019a.
- T. Yu, G. Shevchuk, D. Sadigh, and C. Finn. Unsupervised visuomotor control through distributional planning networks. *CoRR*, abs/1902.05542, 2019b.
- A. Zeng, S. Song, S. Welker, J. Lee, A. Rodriguez, and T. A. Funkhouser. Learning synergies between pushing and grasping with self-supervised deep reinforcement learning. *CoRR*, abs/1803.09956, 2018.
- M. Zhang, S. Vikram, L. Smith, P. Abbeel, M. J. Johnson, and S. Levine. Solar: Deep structured latent representations for model-based reinforcement learning. *ArXiv*, abs/1808.09105, 2018.
- Łukasz Kaiser, M. Babaeizadeh, P. Miłoś, B. Osiński, 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 *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=S1xCPJHtDB>.