
γ -Models: Generative Temporal Difference Learning for Infinite-Horizon Prediction

Michael Janner¹ Igor Mordatch² Sergey Levine^{1,2}
¹UC Berkeley ²Google Brain
{janner, svlevine}@eecs.berkeley.edu imordatch@google.com

Abstract

We introduce the γ -model, a predictive model of environment dynamics with an infinite probabilistic horizon. Replacing standard single-step models with γ -models leads to generalizations of the procedures central to model-based control, including the model rollout and model-based value estimation. The γ -model, trained with a generative reinterpretation of temporal difference learning, is a natural continuous analogue of the successor representation and a hybrid between model-free and model-based mechanisms. Like a value function, it contains information about the long-term future; like a standard predictive model, it is independent of task reward. We instantiate the γ -model as both a generative adversarial network and normalizing flow, discuss how its training reflects an inescapable tradeoff between training-time and testing-time compounding errors, and empirically investigate its utility for prediction and control.

1 Introduction

The common ingredient in all of model-based reinforcement learning is the dynamics model: a function used for predicting future states. The choice of the model’s prediction horizon constitutes a delicate trade-off. Shorter horizons make the prediction problem easier, as the near-term future increasingly begins to look like the present, but may not provide sufficient information for decision-making. Longer horizons carry more information, but present a more difficult prediction problem, as errors accumulate rapidly when a model is applied to its own previous outputs (Janner et al., 2019).

Can we avoid choosing a prediction horizon altogether? Value functions already do so by modeling the cumulative return over a discounted long-term future instead of an immediate reward, circumventing the need to commit to any single finite horizon. However, value prediction folds two problems into one by entangling environment dynamics with reward structure, making value functions less readily adaptable to new tasks in known settings than their model-based counterparts.

In this work, we propose a model that predicts over an infinite horizon with a geometrically-distributed timestep weighting (Figure 1). This γ -model, named for the dependence of its probabilistic horizon on a discount factor γ , is trained with a generative analogue of temporal difference learning suitable for continuous spaces. The γ -model bridges the gap between canonical model-based and model-free mechanisms. Like a value function, it is policy-conditioned and contains information about the distant future; like a conventional dynamics model, it is independent of reward and may be reused for new tasks within the same environment. The γ -model may be instantiated as both a generative adversarial network (Goodfellow et al., 2014) and a normalizing flow (Rezende & Mohamed, 2015).

The shift from standard single-step models to infinite-horizon γ -models carries several advantages:

Constant-time prediction Single-step models must perform an $\mathcal{O}(n)$ operation to predict n steps into the future; γ -models amortize the work of predicting over extended horizons during training such that long-horizon prediction occurs with a single feedforward pass of the model.




Figure 1: Conventional predictive models trained via maximum likelihood have a horizon of one. By interpreting temporal difference learning as a training algorithm for generative models, it is possible to predict with a probabilistic horizon governed by a geometric distribution. In the spirit of infinite-horizon control in model-free reinforcement learning, we refer to this formulation as infinite-horizon prediction.

Generalized rollouts and value estimation Probabilistic prediction horizons lead to generalizations of the core procedures of model-based reinforcement learning. For example, generalized rollouts allow for fine-tuned interpolation between training-time and testing-time compounding error. Similarly, terminal value functions appended to truncated γ -model rollouts allow for a gradual transition between model-based and model-free value estimation.

Omission of unnecessary information The predictions of a γ -model do not come paired with an associated timestep. While on the surface a limitation, we show why knowing precisely *when* a state will be encountered is not necessary for decision-making. Infinite-horizon γ -model prediction selectively discards the unnecessary information from a standard model-based rollout.

2 Related Work

The complementary strengths and weaknesses of model-based and model-free reinforcement learning have led to a number of works that attempt to combine these approaches. Common strategies include initializing a model-free algorithm with the solution found by a model-based planner (Levine & Koltun, 2013; Farshidian et al., 2014; Nagabandi et al., 2018), feeding model-generated data into an otherwise model-free optimizer (Sutton, 1990; Silver et al., 2008; Lampe & Riedmiller, 2014; Kalweit & Boedecker, 2017; Luo et al., 2019), using model predictions to improve the quality of target values for temporal difference learning (Buckman et al., 2018; Feinberg et al., 2018), leveraging model gradients for backpropagation (Nguyen & Widrow, 1990; Jordan & Rumelhart, 1992; Heess et al., 2015), and incorporating model-based planning without explicitly predicting future observations (Tamar et al., 2016; Silver et al., 2017; Oh et al., 2017; Kahn et al., 2018; Amos et al., 2018; Schrittwieser et al., 2019). In contrast to combining independent model-free and model-based components, we describe a framework for training a new class of predictive model with a generative, model-based reinterpretation of model-free tools.

Temporal difference models (TDMs) Pong et al. (2018) provide an alternative method of training models with what are normally considered to be model-free algorithms. TDMs interpret models as a special case of goal-conditioned value functions (Kaelbling, 1993; Foster & Dayan, 2002; Schaul et al., 2015; Andrychowicz et al., 2017), though the TDM is constrained to predict at a fixed horizon and is limited to tasks for which the reward depends only on the last state. In contrast, the γ -model predicts over a discounted infinite-horizon future and accommodates arbitrary rewards.

The most closely related prior work to γ -models is the successor representation (Dayan, 1993), a formulation of long-horizon prediction that has been influential in both cognitive science (Momennejad et al., 2017; Gershman, 2018) and machine learning (Kulkarni et al., 2016; Ma et al., 2018). In its original form, the successor representation is tractable only in tabular domains. Prior continuous variants have focused on policy evaluation based on expected state featurizations (Barreto et al., 2017, 2018; Hansen et al., 2020), forgoing an interpretation as a probabilistic model suitable for state prediction. Converting the tabular successor representation into a continuous generative model is non-trivial because the successor representation implicitly assumes the ability to normalize over a finite state space for interpretation as a predictive model.

Because of the discounted state occupancy’s central role in reinforcement learning, its approximation by Bellman equations has been the focus of multiple lines of work. Generalizations include β -models

(Sutton, 1995), allowing for arbitrary mixture distributions over time, and option models (Sutton et al., 1999), allowing for state-dependent termination conditions. While our focus is on generative models featuring the state-independent geometric timestep weighting of the successor representation, we are hopeful that the tools developed in this paper could also be applicable in the design of continuous analogues of these generalizations.

3 Preliminaries

We consider an infinite-horizon Markov decision process (MDP) defined by the tuple $(\mathcal{S}, \mathcal{A}, p, r, \gamma)$, with state space \mathcal{S} and action space \mathcal{A} . The transition distribution and reward function are given by $p : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}^+$ and $r : \mathcal{S} \rightarrow \mathbb{R}$, respectively. The discount is denoted by $\gamma \in (0, 1)$. A policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}^+$ induces a conditional occupancy $\mu(\mathbf{s} \mid \mathbf{s}_t, \mathbf{a}_t)$ over future states:

$$\mu(\mathbf{s} \mid \mathbf{s}_t, \mathbf{a}_t) = (1 - \gamma) \sum_{\Delta t=1}^{\infty} \gamma^{\Delta t-1} p(\mathbf{s}_{t+\Delta t} = \mathbf{s} \mid \mathbf{s}_t, \mathbf{a}_t, \pi). \quad (1)$$

We denote parametric approximations of p (μ) as p_θ (μ_θ), in which the subscripts denote model parameters. Standard model-based reinforcement learning algorithms employ the single-step model p_θ for long-horizon decision-making by performing multi-step model-based rollouts.

4 Generative Temporal Difference Learning

Our goal is to make long-horizon predictions without the need to repeatedly apply a single-step model. Instead of modeling states at a particular instant in time by approximating the environment transition distribution $p(\mathbf{s}_{t+1} \mid \mathbf{s}_t, \mathbf{a}_t)$, we aim to predict a weighted distribution over all possible future states according to $\mu(\mathbf{s} \mid \mathbf{s}_t, \mathbf{a}_t)$. In principle, this can be posed as a conventional maximum likelihood problem:

$$\max_{\theta} \mathbb{E}_{\mathbf{s}_t, \mathbf{a}_t, \mathbf{s} \sim \mu(\cdot \mid \mathbf{s}_t, \mathbf{a}_t)} [\log \mu_\theta(\mathbf{s} \mid \mathbf{s}_t, \mathbf{a}_t)].$$

However, doing so would require collecting samples from the occupancy μ independently for each policy of interest. Forgoing the ability to re-use data from multiple policies when training dynamics models would sacrifice the sample efficiency that often makes model usage compelling in the first place, so we instead aim to design an off-policy algorithm for training μ_θ . We accomplish this by reinterpreting temporal difference learning as a method for training generative models.

Instead of collecting only on-policy samples from $\mu(\mathbf{s} \mid \mathbf{s}_t, \mathbf{a}_t)$, we observe that μ admits a convenient recursive form. Consider a modified MDP in which there is a $1 - \gamma$ probability of terminating at each timestep. The distribution over the state at termination, denoted as the exit state \mathbf{s}_e , corresponds to first sampling from a termination timestep $\Delta t \sim \text{Geom}(1 - \gamma)$ and then sampling from the per-timestep distribution $p(\mathbf{s}_{t+\Delta t} \mid \mathbf{s}_t, \mathbf{a}_t, \pi)$. The distribution over \mathbf{s}_e corresponds exactly to that in the definition of the occupancy μ in Equation 1, but also lends itself to an interpretation as a mixture over only two components: the distribution at the immediate next timestep, in the event of termination, and that over all subsequent timesteps, in the event of non-termination. This mixture yields the following target distribution:

$$p_{\text{targ}}(\mathbf{s}_e \mid \mathbf{s}_t, \mathbf{a}_t) = \underbrace{(1 - \gamma)p(\mathbf{s}_e \mid \mathbf{s}_t, \mathbf{a}_t)}_{\text{single-step distribution}} + \underbrace{\gamma \mathbb{E}_{\mathbf{s}_{t+1} \sim p(\cdot \mid \mathbf{s}_t, \mathbf{a}_t)} [\mu_\theta(\mathbf{s}_e \mid \mathbf{s}_{t+1})]}_{\text{model bootstrap}}. \quad (2)$$

We use the shorthand $\mu_\theta(\mathbf{s}_e \mid \mathbf{s}_{t+1}) = \mathbb{E}_{\mathbf{a}_{t+1} \sim \pi(\cdot \mid \mathbf{s}_{t+1})} [\mu_\theta(\mathbf{s}_e \mid \mathbf{s}_{t+1}, \mathbf{a}_{t+1})]$. The target distribution p_{targ} is reminiscent of a temporal difference target value: the state-action conditioned occupancy $\mu_\theta(\mathbf{s}_e \mid \mathbf{s}_t, \mathbf{a}_t)$ acts as a Q -function, the state-conditioned occupancy $\mu_\theta(\mathbf{s}_e \mid \mathbf{s}_{t+1})$ acts as a value function, and the single-step distribution $p(\mathbf{s}_{t+1} \mid \mathbf{s}_t, \mathbf{a}_t)$ acts as a reward function. However, instead of representing a scalar target value, p_{targ} is a distribution from which we may sample future states \mathbf{s}_e . We can use this target distribution in place of samples from the true discounted occupancy μ :

$$\max_{\theta} \mathbb{E}_{\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_e \sim (1-\gamma)p(\cdot \mid \mathbf{s}_t, \mathbf{a}_t) + \gamma \mathbb{E}[\mu_\theta(\cdot \mid \mathbf{s}_{t+1})]} [\log \mu_\theta(\mathbf{s}_e \mid \mathbf{s}_t, \mathbf{a}_t)].$$

This formulation differs from a standard maximum likelihood learning problem in that the target distribution depends on the current model. By bootstrapping the target distribution in this manner, we

are able to use only empirical $(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1})$ transitions from one policy in order to train an infinite-horizon predictive model μ_θ for any other policy. Because the horizon is governed by the discount γ , we refer to such a model as a γ -model.

This bootstrapped model training may be incorporated into a number of different generative modeling frameworks. We discuss two cases here. (1) When the model μ_θ permits only sampling, we may train μ_θ by minimizing an f -divergence from samples:

$$\mathcal{L}_1(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}) = D_f(\mu_\theta(\cdot | \mathbf{s}_t, \mathbf{a}_t) || (1 - \gamma)p(\cdot | \mathbf{s}_t, \mathbf{a}_t) + \gamma\mu_\theta(\cdot | \mathbf{s}_{t+1})). \quad (3)$$

This objective leads naturally to an adversarially-trained γ -model. (2) When the model μ_θ permits density evaluation, we may minimize an error defined on log-densities directly:

$$\mathcal{L}_2(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}) = \mathbb{E}_{\mathbf{s}_e} \left[\left\| \log \mu_\theta(\mathbf{s}_e | \mathbf{s}_t, \mathbf{a}_t) - \log \left((1 - \gamma)p(\mathbf{s}_e | \mathbf{s}_t, \mathbf{a}_t) + \gamma\mu_\theta(\mathbf{s}_e | \mathbf{s}_{t+1}) \right) \right\|_2^2 \right]. \quad (4)$$

This objective is suitable for γ -models instantiated as normalizing flows. Due to the approximation of a target log-density $\log \left((1 - \gamma)p(\cdot | \mathbf{s}_t, \mathbf{a}_t) + \gamma\mathbb{E}_{\mathbf{s}_{t+1}} [\mu_\theta(\cdot | \mathbf{s}_{t+1})] \right)$ using a single next state \mathbf{s}_{t+1} , \mathcal{L}_2 is unbiased for deterministic dynamics and a bound in the case of stochastic dynamics. We provide complete algorithmic descriptions of both variants and highlight practical training considerations in Section 6.

5 Analysis and Applications of γ -Models

Using the γ -model for prediction and control requires us to generalize procedures common in model-based reinforcement learning. In this section, we derive the γ -model rollout and show how it can be incorporated into a reinforcement learning procedure that hybridizes model-based and model-free value estimation. First, however, we show that the γ -model is a continuous, generative counterpart to another type of long-horizon model: the successor representation.

5.1 γ -Models as a Continuous Successor Representation

The successor representation M is a prediction of expected visitation counts (Dayan, 1993). It has a recurrence relation making it amenable to tabular temporal difference algorithms:

$$M(\mathbf{s}_e | \mathbf{s}_t, \mathbf{a}_t) = \mathbb{E}_{\mathbf{s}_{t+1} \sim p(\cdot | \mathbf{s}_t, \mathbf{a}_t)} [\mathbb{1}[\mathbf{s}_e = \mathbf{s}_{t+1}] + \gamma M(\mathbf{s}_e | \mathbf{s}_{t+1})]. \quad (5)$$

Adapting the successor representation to continuous state spaces in a way that retains an interpretation as a probabilistic model has proven challenging. However, variants that forego the ability to sample in favor of estimating expected state features have been developed (Barreto et al., 2017).

The form of the successor recurrence relation bears a striking resemblance to that of the target distribution in Equation 2, suggesting a connection between the generative, continuous γ -model and the discriminative, tabular successor representation. We now make this connection precise.

Proposition 1. *The global minimum of both \mathcal{L}_1 and \mathcal{L}_2 is achieved if and only if the resulting γ -model produces samples according to the normalized successor representation:*

$$\mu_\theta(\mathbf{s}_e | \mathbf{s}_t, \mathbf{a}_t) = (1 - \gamma)M(\mathbf{s}_e | \mathbf{s}_t, \mathbf{a}_t).$$

Proof. In the case of either objective, the global minimum is achieved only when

$$\mu_\theta(\mathbf{s}_e | \mathbf{s}_t, \mathbf{a}_t) = (1 - \gamma)p(\mathbf{s}_e | \mathbf{s}_t, \mathbf{a}_t) + \gamma\mathbb{E}_{\mathbf{s}_{t+1} \sim p(\cdot | \mathbf{s}_t, \mathbf{a}_t)} [\mu_\theta(\mathbf{s}_e | \mathbf{s}_{t+1})]$$

for all $\mathbf{s}_t, \mathbf{a}_t$. We recognize this optimality condition exactly as the recurrence defining the successor representation M (Equation 5), scaled by $(1 - \gamma)$ such that μ_θ integrates to 1 over \mathbf{s}_e . \square

5.2 γ -Model Rollouts

Standard single-step models, which correspond to γ -models with $\gamma = 0$, can predict multiple steps into the future by making iterated autoregressive predictions, conditioning each step on their own output from the previous step. These sequential rollouts form the foundation of most model-based

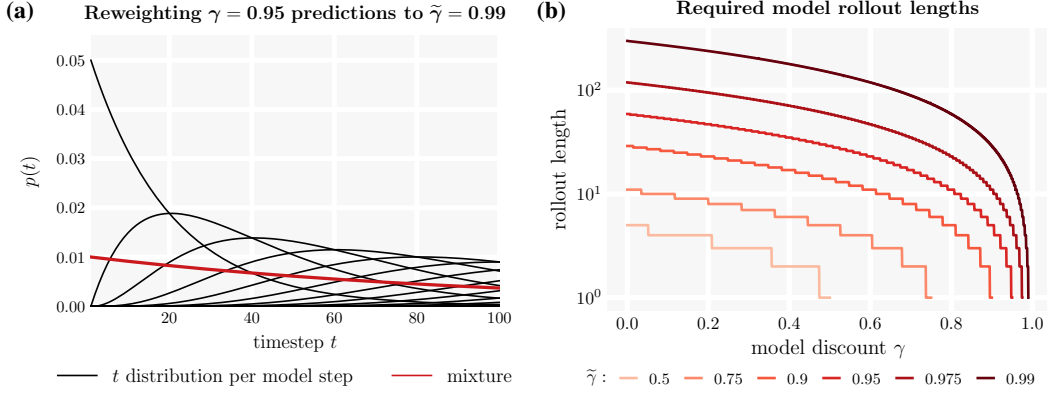


Figure 2: **(a)** The first step from a γ -model samples states at timesteps distributed according to a geometric distribution with parameter $1 - \gamma$; all subsequent steps have a negative binomial timestep distribution stemming from the sum of independent geometric random variables. When these steps are reweighted according to Theorem 1, the resulting distribution follows a geometric distribution with smaller parameter (corresponding to a larger discount value $\tilde{\gamma}$). **(b)** The number of steps needed to recover 95% of the probability mass from distributions induced by various target discounts $\tilde{\gamma}$ for all valid model discounts γ . When using a standard single-step model, corresponding to the case of $\gamma = 0$, a 299-step model rollout is required to reweight to a discount of $\tilde{\gamma} = 0.99$.

reinforcement learning algorithms. We now generalize these rollouts to γ -models for $\gamma > 0$, allowing us to decouple the discount used during model training from the desired horizon in control. When working with multiple discount factors, we explicitly condition an occupancy on its discount as $\mu(\mathbf{s}_e | \mathbf{s}_t; \gamma)$. In the results below, we omit the model parameterization θ whenever a statement applies to both a discounted occupancy μ and a parametric γ -model μ_θ .

Theorem 1. Let $\mu_n(\mathbf{s}_e | \mathbf{s}_t; \gamma)$ denote the distribution over states at the n^{th} sequential step of a γ -model rollout beginning from state \mathbf{s}_t . For any desired discount $\tilde{\gamma} \in [\gamma, 1)$, we may reweight the samples from these model rollouts according to the weights

$$\alpha_n = \frac{(1 - \tilde{\gamma})(\tilde{\gamma} - \gamma)^{n-1}}{(1 - \gamma)^n}$$

to obtain the state distribution drawn from $\mu_1(\mathbf{s}_e | \mathbf{s}_t; \tilde{\gamma}) = \mu(\mathbf{s}_e | \mathbf{s}_t; \tilde{\gamma})$. That is, we may reweight the steps of a γ -model rollout so as to match the distribution of a $\tilde{\gamma}$ -model with larger discount:

$$\mu(\mathbf{s}_e | \mathbf{s}_t; \tilde{\gamma}) = \sum_{n=1}^{\infty} \alpha_n \mu_n(\mathbf{s}_e | \mathbf{s}_t; \gamma).$$

Proof. See Appendix A. □

This reweighting scheme has two special cases of interest. A standard single-step model, with $\gamma = 0$, yields $\alpha_n = (1 - \tilde{\gamma})\tilde{\gamma}^{n-1}$. These weights are familiar from the definition of the discounted state occupancy in terms of a per-timestep mixture (Equation 1). Setting $\gamma = \tilde{\gamma}$ yields $\alpha_n = 0^{n-1}$, or a weight of 1 on the first step and 0 on all subsequent steps.¹ This result is also expected: when the model discount matches the target discount, only a single forward pass of the model is required.

Figure 2 visually depicts the reweighting scheme and the number of steps required for truncated model rollouts to approximate the distribution induced by a larger discount. There is a natural tradeoff with γ -models: the higher γ is, the fewer model steps are needed to make long-horizon predictions, reducing model-based compounding prediction errors (Asadi et al., 2019; Janner et al., 2019). However, increasing γ transforms what would normally be a standard maximum likelihood problem (in the case of single-step models) into one resembling approximate dynamic programming (with a model bootstrap), leading to model-free bootstrap error accumulation (Kumar et al., 2019). The primary distinction is whether this accumulation occurs during training, when the work of sampling from the occupancy μ is being amortized, or during “testing”, when the model is being used for rollouts. While this horizon-based error compounding cannot be eliminated entirely, γ -models allow for a continuous interpolation between the two extremes.

¹We define 0^0 as $\lim_{x \rightarrow 0} x^x = 1$.

5.3 γ -Model-Based Value Expansion

We now turn our attention from prediction with γ -models to value estimation for control. In tabular domains, the state-action value function can be decomposed as the inner product between the successor representation M and the vector of per-state rewards (Gershman, 2018). Taking care to account for the normalization from the equivalence in Proposition 1, we can similarly estimate the Q function as the expectation of reward under states sampled from the γ -model:

$$\begin{aligned} Q(\mathbf{s}_t, \mathbf{a}_t; \gamma) &= \sum_{\Delta t=1}^{\infty} \gamma^{\Delta t-1} \int_{\mathcal{S}} r(\mathbf{s}_e) p(\mathbf{s}_{t+\Delta t} = \mathbf{s}_e \mid \mathbf{s}_t, \mathbf{a}_t, \pi) d\mathbf{s}_e \\ &= \int_{\mathcal{S}} r(\mathbf{s}_e) \sum_{\Delta t=1}^{\infty} \gamma^{\Delta t-1} p(\mathbf{s}_{t+\Delta t} = \mathbf{s}_e \mid \mathbf{s}_t, \mathbf{a}_t, \pi) d\mathbf{s}_e \\ &= \frac{1}{1-\gamma} \mathbb{E}_{\mathbf{s}_e \sim \mu(\cdot \mid \mathbf{s}_t, \mathbf{a}_t; \gamma)} [r(\mathbf{s}_e)] \end{aligned} \quad (6)$$

This relation suggests a model-based reinforcement learning algorithm in which Q -values are estimated by a γ -model without the need for sequential model-based rollouts. However, in some cases it may be practically difficult to train a generative γ -model with discount as large as that of a discriminative Q -function. While one option is to chain together γ -model steps as in Section 5.2, an alternative solution often effective with single-step models is to combine short-term value estimates from a truncated model rollout with a terminal model-free value prediction:

$$V_{\text{MVE}}(\mathbf{s}_t; \tilde{\gamma}) = \sum_{n=1}^H \tilde{\gamma}^{n-1} r(\mathbf{s}_{t+n}) + \tilde{\gamma}^H V(\mathbf{s}_{t+H}; \tilde{\gamma}).$$

This hybrid estimator is referred to as a model-based value expansion (MVE; Feinberg et al. 2018). There is a hard transition between the model-based and model-free value estimation in MVE, occurring at the model horizon H . We may replace the single-step model with a γ -model for a similar estimator in which there is a probabilistic prediction horizon, and as a result a gradual transition:

$$V_{\gamma\text{-MVE}}(\mathbf{s}_t; \tilde{\gamma}) = \frac{1}{1-\tilde{\gamma}} \sum_{n=1}^H \alpha_n \mathbb{E}_{\mathbf{s}_e \sim \mu_n(\cdot \mid \mathbf{s}_t; \gamma)} [r(\mathbf{s}_e)] + \left(\frac{\tilde{\gamma} - \gamma}{1-\gamma} \right)^H \mathbb{E}_{\mathbf{s}_e \sim \mu_H(\cdot \mid \mathbf{s}_t; \gamma)} [V(\mathbf{s}_e; \tilde{\gamma})].$$

The γ -MVE estimator allows us to perform γ -model-based rollouts with horizon H , reweight the samples from this rollout by solving for weights α_n given a desired discount $\tilde{\gamma} > \gamma$, and correct for the truncation error stemming from the finite rollout length using a terminal value function with discount $\tilde{\gamma}$. As expected, MVE is a special case of γ -MVE, as can be verified by considering the weights corresponding to $\gamma = 0$ described in Section 5.2. This estimator, along with the simpler value estimation in Equation 6, highlights the fact that it is not necessary to have timesteps associated with states in order to use predictions for decision-making. We provide a more thorough treatment of γ -MVE, complete with pseudocode for a corresponding actor-critic algorithm, in Appendix B.

6 Practical Training of γ -Models

Because γ -model training differs from standard dynamics modeling primarily in the bootstrapped target distribution and not in the model parameterization, γ -models are in principle compatible with any generative modeling framework. We focus on two representative scenarios, differing in whether the generative model class used to instantiate the γ -model allows for tractable density evaluation.

Training without density evaluation When the γ -model parameterization does not allow for tractable density evaluation, we minimize a bootstrapped f -divergence according to \mathcal{L}_1 (Equation 3) using only samples from the model. The generative adversarial framework provides a convenient way to train a parametric generator by minimizing an f -divergence of choice given only samples from a target distribution p_{targ} and the ability to sample from the generator (Goodfellow et al., 2014; Nowozin et al., 2016). In the case of bootstrapped maximum likelihood problems, our target distribution is induced by the model itself (alongside a single-step transition distribution), meaning that we only need sample access to our γ -model in order to train μ_θ as a generative adversarial network (GAN).

Introducing an auxiliary discriminator D_ϕ and selecting the Jensen-Shannon divergence as our f -divergence, we can reformulate minimization of the original objective \mathcal{L}_1 as a saddle-point

Algorithm 1 γ -model training **without density evaluation**

- 1: **Input** \mathcal{D} : dataset of transitions, π : policy, λ : step size, τ : delay parameter
 - 2: Initialize parameter vectors $\theta, \bar{\theta}, \phi$
 - 3: **while** not converged **do**
 - 4: Sample transitions $(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1})$ from \mathcal{D} and actions $\mathbf{a}_{t+1} \sim \pi(\cdot | \mathbf{s}_{t+1})$
 - 5: **Sample from bootstrapped target** $\mathbf{s}_e^+ \sim (1 - \gamma)\delta_{\mathbf{s}_{t+1}} + \gamma\mu_{\bar{\theta}}(\cdot | \mathbf{s}_{t+1}, \mathbf{a}_{t+1})$
 - 6: **Sample from current model** $\mathbf{s}_e^- \sim \mu_{\theta}(\cdot | \mathbf{s}_t, \mathbf{a}_t)$
 - 7: **Evaluate objective** $\mathcal{L} = \log D_{\phi}(\mathbf{s}_e^+ | \mathbf{s}_t, \mathbf{a}_t) + \log(1 - D_{\phi}(\mathbf{s}_e^- | \mathbf{s}_t, \mathbf{a}_t))$
 - 8: Update model parameters $\theta \leftarrow \theta - \lambda \nabla_{\theta} \mathcal{L}$; $\phi \leftarrow \phi + \lambda \nabla_{\phi} \mathcal{L}$
 - 9: Update target parameters $\bar{\theta} \leftarrow \tau\theta + (1 - \tau)\bar{\theta}$
 - 10: **end while**
-

Algorithm 2 γ -model training **with density evaluation**

- 1: **Input** \mathcal{D} : dataset of transitions, π : policy, λ : step size, τ : delay parameter, σ^2 : variance
 - 2: Initialize parameter vectors $\theta, \bar{\theta}$; **let** f **denote the Gaussian pdf**
 - 3: **while** not converged **do**
 - 4: Sample transitions $(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1})$ from \mathcal{D} and actions $\mathbf{a}_{t+1} \sim \pi(\cdot | \mathbf{s}_{t+1})$
 - 5: **Sample from bootstrapped target** $\mathbf{s}_e \sim (1 - \gamma)\mathcal{N}(\mathbf{s}_{t+1}, \sigma^2) + \gamma\mu_{\bar{\theta}}(\cdot | \mathbf{s}_{t+1}, \mathbf{a}_{t+1})$
 - 6: **Construct target values** $T = \log((1 - \gamma)f(\mathbf{s}_e | \mathbf{s}_{t+1}, \sigma^2) + \gamma\mu_{\bar{\theta}}(\mathbf{s}_e | \mathbf{s}_{t+1}, \mathbf{a}_{t+1}))$
 - 7: **Evaluate objective** $\mathcal{L} = \|\log \mu_{\theta}(\mathbf{s}_e | \mathbf{s}_t, \mathbf{a}_t) - T\|_2^2$
 - 8: Update model parameters $\theta \leftarrow \theta - \lambda \nabla_{\theta} \mathcal{L}$
 - 9: Update target parameters $\bar{\theta} \leftarrow \tau\theta + (1 - \tau)\bar{\theta}$
 - 10: **end while**
-

optimization over the following objective:

$$\hat{\mathcal{L}}_1(\mathbf{s}_t, \mathbf{a}_t) = \mathbb{E}_{\mathbf{s}_e^+ \sim p_{\text{targ}}(\cdot | \mathbf{s}_t, \mathbf{a}_t)} [\log D_{\phi}(\mathbf{s}_e^+ | \mathbf{s}_t, \mathbf{a}_t)] + \mathbb{E}_{\mathbf{s}_e^- \sim \mu_{\theta}(\cdot | \mathbf{s}_t, \mathbf{a}_t)} [\log(1 - D_{\phi}(\mathbf{s}_e^- | \mathbf{s}_t, \mathbf{a}_t))],$$

which is minimized over μ_{θ} and maximized over D_{ϕ} . As in \mathcal{L}_1 , p_{targ} refers to the bootstrapped target distribution in Equation 2. In this formulation, μ_{θ} produces samples by virtue of a deterministic mapping of a random input vector $\mathbf{z} \sim \mathcal{N}(0, I)$ and conditioning information $(\mathbf{s}_t, \mathbf{a}_t)$. Other choices of f -divergence may be instantiated by different choices of activation function (Nowozin et al., 2016).

Training with density evaluation When the γ -model permits density evaluation, we may bypass saddle point approximations to an f -divergence and directly regress to target density values, as in objective \mathcal{L}_2 (Equation 4). This is a natural choice when the γ -model is instantiated as a conditional normalizing flow (Rezende & Mohamed, 2015). Evaluating target values of the form

$$T(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}, \mathbf{s}_e) = \log((1 - \gamma)p(\mathbf{s}_e | \mathbf{s}_t, \mathbf{a}_t) + \gamma\mu_{\theta}(\mathbf{s}_e | \mathbf{s}_{t+1}))$$

requires density evaluation of not only our γ -model, but also the single-step transition distribution. There are two options for estimating the single-step densities: (1) a single-step model p_{θ} may be trained alongside the γ -model μ_{θ} for the purposes of constructing targets $T(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1}, \mathbf{s}_e)$, or (2) a simple approximate model may be constructed on the fly from $(\mathbf{s}_t, \mathbf{a}_t, \mathbf{s}_{t+1})$ transitions. We found $p_{\theta} = \mathcal{N}(\mathbf{s}_{t+1}, \sigma^2)$, with σ^2 a constant hyperparameter, to be sufficient.

Stability considerations To alleviate the instability caused by bootstrapping, we appeal to the standard solution employed in model-free reinforcement learning: decoupling the regression targets from the current model by way of a “delayed” target network (Mnih et al., 2015). In particular, we use a delayed γ -model $\mu_{\bar{\theta}}$ in the bootstrapped target distribution p_{targ} , with the parameters $\bar{\theta}$ given by an exponentially-moving average of previous parameters θ .

We summarize the above scenarios in Algorithms 1 and 2. We isolate model training from data collection and focus on a setting in which a static dataset is provided, but this algorithm may also be used in a data-collection loop for policy improvement. Further implementation details, including all hyperparameter settings and network architectures, are included in Appendix C.

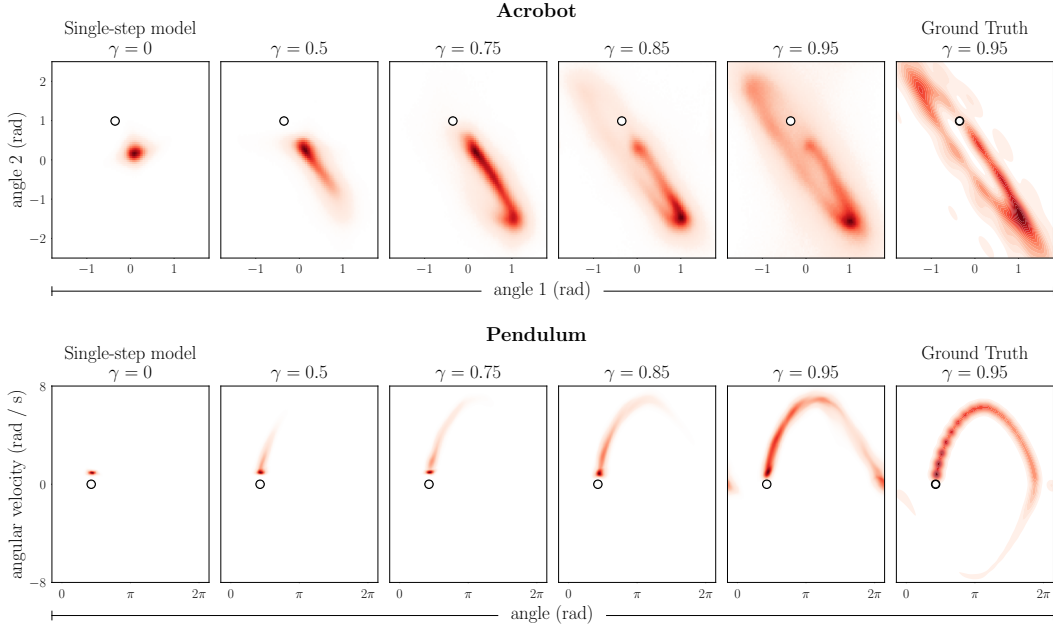


Figure 3: Visualization of the predicted distribution from a **single** feedforward pass of normalizing flow γ -models trained with varying discounts γ . The conditioning state s_t is denoted by \circ . The leftmost plots, with $\gamma = 0$, correspond to a single-step model. For comparison, the rightmost plots show a Monte Carlo estimation of the discounted occupancy from 100 environment trajectories.

7 Experiments

Our experimental evaluation is designed to study the viability of γ -models as a replacement of conventional single-step models for long-horizon state prediction and model-based control.

7.1 Prediction

We investigate γ -model predictions as a function of discount in continuous-action versions of two benchmark environments suitable for visualization: acrobot (Sutton, 1996) and pendulum. The training data come from a mixture distribution over all intermediate policies of 200 epochs of optimization with soft-actor critic (SAC; Haarnoja et al. 2018). The final converged policy is used for γ -model training. We refer to Appendix C for implementation and experiment details.

Figure 3 shows the predictions of a γ -model trained as a normalizing flow according to Algorithm 2 for five different discounts, ranging from $\gamma = 0$ (a single-step model) to $\gamma = 0.95$. The rightmost column shows the ground truth discounted occupancy corresponding to $\gamma = 0.95$, estimated with Monte Carlo rollouts of the policy. Increasing the discount γ during training has the expected effect of qualitatively increasing the predictive lookahead of a single feedforward pass of the γ -model. We found flow-based γ -models to be more reliable than GAN parameterizations, especially at higher discounts. Corresponding GAN γ -model visualizations can be found in Appendix E for comparison.

Equation 6 expresses values as an expectation over a single feedforward pass of a γ -model. We visualize this relation in Figure 4, which depicts γ -model predictions on the pendulum environment for a discount of $\gamma = 0.99$ and the resulting value map estimated by taking expectations over these predicted state distributions. In comparison, value estimation for the same discount using a single-step model would require 299-step rollouts in order to recover 95% of the probability mass (see Figure 2).

7.2 Control

To study the utility of the γ -model for model-based reinforcement learning, we use the γ -MVE estimator from Section 5.3 as a drop-in replacement for value estimation in SAC. We compare this approach to the state-of-the-art in model-based and model-free methods, with representative algorithms consisting of SAC, PPO (Schulman et al., 2017), MBPO (Janner et al., 2019), and MVE (Feinberg et al., 2018). In γ -MVE, we use a model discount of $\gamma = 0.8$, a value discount of

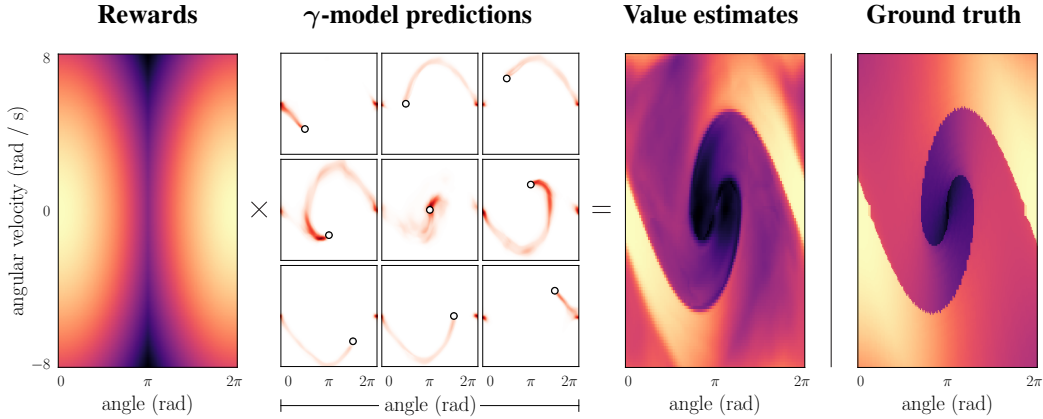


Figure 4: Values are expectations of reward over a single feedforward pass of a γ -model (Equation 6). We visualize γ -model predictions ($\gamma = 0.99$) from nine starting states, denoted by \circ , in the pendulum benchmark environment. Taking the expectation of reward over each of these predicted distributions yields a value estimate for the corresponding conditioning state. The rightmost plot depicts the value map produced by value iteration on a discretization of the same environment for reference.




Figure 5: Comparative performance of γ -MVE and four prior reinforcement learning algorithms on continuous control benchmark tasks. γ -MVE retains the asymptotic performance of SAC with sample-efficiency matching that of MBPO. Shaded regions depict standard deviation among 5 seeds.

$\tilde{\gamma} = 0.99$ and a single model step ($n = 1$). We use a model rollout length of 5 in MVE such that it has an effective horizon identical to that of γ -MVE. Other hyperparameter settings can once again be found in Appendix C; details regarding the evaluation environments can be found in Appendix D. Figure 5 shows learning curves for all methods. We find that γ -MVE converges faster than prior algorithms, twice as quickly as SAC, while retaining their asymptotic performance.

8 Discussion, Limitations, and Future Work

We have introduced a new class of predictive model, a γ -model, that is a hybrid between standard model-free and model-based mechanisms. It is policy-conditioned and infinite-horizon, like a value function, but independent of reward, like a standard single-step model. This new formulation of infinite-horizon prediction allows us to generalize the procedures integral to model-based control, yielding new variants of model rollouts and model-based value estimation.

Our experimental evaluation shows that, on tasks with low to moderate dimensionality, our method learns accurate long-horizon predictive distributions without sequential rollouts and can be incorporated into standard model-based reinforcement learning methods to produce results that are competitive with state-of-the-art algorithms. Scaling up our framework to more complex tasks, including high-dimensional continuous control problems and tasks with image observations, presents a number of additional challenges. We are optimistic that continued improvements in training techniques for generative models and increased stability in off-policy reinforcement learning will also carry benefits for γ -model training.

Broader Impact

γ -models provide a way to study infinite-horizon prediction in the same way that we can study the problem of control in infinite-horizon MDPs, possibly providing a path for more accurate long-term modeling without test-time compounding error. However, transforming a maximum likelihood problem into an approximate dynamic programming one is not without its drawbacks: maximum likelihood learning techniques are substantially better understood at this point in time and can be accompanied by rigorous guarantees even when incorporating function approximation. It will take much future work not just on γ -models, but on dynamic programming methods more broadly, for them to work reliably enough for deployment in real-world, safety-critical systems. While much work remains, we are optimistic for the long-term viability of temporal difference learning as an algorithm for training long-horizon dynamics models given its empirical success in long-horizon model-free control.

Acknowledgements

We thank Karthik Narasimhan, Pulkit Agrawal, Anirudh Goyal, Pedro Tsividis, Anusha Nagabandi, Aviral Kumar, and Michael Chang for formative discussions about model-based reinforcement learning and generative modeling. This work was partially supported by computational resource donations from Amazon. M.J. is supported by fellowships from the National Science Foundation and the Open Philanthropy Project.

References

- Amos, B., Rodriguez, I. D. J., Sacks, J., Boots, B., and Kolter, J. Z. Differentiable mpc for end-to-end planning and control. In *Advances in Neural Information Processing Systems*, 2018.
- Andrychowicz, M., Wolski, F., Ray, A., Schneider, J., Fong, R., Welinder, P., McGrew, B., Tobin, J., Abbeel, P., and Zaremba, W. Hindsight experience replay. In *Advances in Neural Information Processing Systems*. 2017.
- Asadi, K., Misra, D., Kim, S., and Littman, M. L. Combating the compounding-error problem with a multi-step model. *arXiv preprint arXiv:1905.13320*, 2019.
- Barreto, A., Dabney, W., Munos, R., Hunt, J. J., Schaul, T., van Hasselt, H. P., and Silver, D. Successor features for transfer in reinforcement learning. In *Advances in Neural Information Processing Systems 30*. 2017.
- Barreto, A., Borsa, D., Quan, J., Schaul, T., Silver, D., Hessel, M., Mankowitz, D., Zidek, A., and Munos, R. Transfer in deep reinforcement learning using successor features and generalised policy improvement. In *Proceedings of the International Conference on Machine Learning*, 2018.
- Buckman, J., Hafner, D., Tucker, G., Brevdo, E., and Lee, H. Sample-efficient reinforcement learning with stochastic ensemble value expansion. In *Advances in Neural Information Processing Systems*, 2018.
- Chua, K., Calandra, R., McAllister, R., and Levine, S. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*. 2018.
- Dayan, P. Improving generalization for temporal difference learning: The successor representation. *Neural Computation*, 5:613, 1993.
- Durkan, C., Bekasov, A., Murray, I., and Papamakarios, G. Neural spline flows. In *Advances in Neural Information Processing Systems*. 2019.
- Farshidian, F., Neunert, M., and Buchli, J. Learning of closed-loop motion control. In *International Conference on Intelligent Robots and Systems*, 2014.
- Feinberg, V., Wan, A., Stoica, I., Jordan, M. I., Gonzalez, J. E., and Levine, S. Model-based value estimation for efficient model-free reinforcement learning. In *International Conference on Machine Learning*, 2018.

- Foster, D. and Dayan, P. Structure in the space of value functions. *Machine Learning*, 49:325, 2002.
- Gershman, S. J. The successor representation: Its computational logic and neural substrates. *Journal of Neuroscience*, 2018.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in Neural Information Processing Systems*. 2014.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, 2018.
- Hansen, S., Dabney, W., Barreto, A., Warde-Farley, D., de Wiele, T. V., and Mnih, V. Fast task inference with variational intrinsic successor features. In *International Conference on Learning Representations*, 2020.
- Heess, N., Wayne, G., Silver, D., Lillicrap, T., Tassa, Y., and Erez, T. Learning continuous control policies by stochastic value gradients. In *Advances in Neural Information Processing Systems*, 2015.
- Janner, M., Fu, J., Zhang, M., and Levine, S. When to trust your model: Model-based policy optimization. In *Advances in Neural Information Processing Systems*, 2019.
- Jordan, M. I. and Rumelhart, D. E. Forward models: Supervised learning with a distal teacher. *Cognitive Science*, 16:307, 1992.
- Kaelbling, L. P. Learning to achieve goals. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1993.
- Kahn, G., Villafior, A., Ding, B., Abbeel, P., and Levine, S. Self-supervised deep reinforcement learning with generalized computation graphs for robot navigation. In *International Conference on Robotics and Automation*, 2018.
- Kalweit, G. and Boedecker, J. Uncertainty-driven imagination for continuous deep reinforcement learning. In *Conference on Robot Learning*, 2017.
- Kulkarni, T. D., Saeedi, A., Gautam, S., and Gershman, S. J. Deep successor reinforcement learning, 2016.
- Kumar, A., Fu, J., Tucker, G., and Levine, S. Stabilizing off-policy q-learning via bootstrapping error reduction. In *Advances in Neural Information Processing Systems*, 2019.
- Lampe, T. and Riedmiller, M. Approximate model-assisted neural fitted Q-iteration. In *International Joint Conference on Neural Networks*, 2014.
- Levine, S. and Koltun, V. Guided policy search. In *International Conference on Machine Learning*, 2013.
- Luo, Y., Xu, H., Li, Y., Tian, Y., Darrell, T., and Ma, T. Algorithmic framework for model-based deep reinforcement learning with theoretical guarantees. In *International Conference on Learning Representations*, 2019.
- Ma, C., Wen, J., and Bengio, Y. Universal successor representations for transfer reinforcement learning. *arXiv preprint arXiv:1804.03758*, 2018.
- Mao, X., Li, Q., Xie, H., Lau, R. Y. K., and Wang, Z. Least squares generative adversarial networks. *arXiv preprint arXiv:1611.04076*, 2016.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D. Human-level control through deep reinforcement learning. *Nature*, 2015.

- Momennejad, I., Russek, E. M., Cheong, J. H., Botvinick, M. M., Daw, N. D., and Gershman, S. J. The successor representation in human reinforcement learning. *Nature Human Behaviour*, 1(9): 680–692, 2017.
- Moore, A. W. *Efficient Memory-based Learning for Robot Control*. PhD thesis, University of Cambridge, 1990.
- Nagabandi, A., Kahn, G., S. Fearing, R., and Levine, S. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *International Conference on Robotics and Automation*, 2018.
- Nguyen, D. H. and Widrow, B. Neural networks for self-learning control systems. *IEEE Control Systems Magazine*, 1990.
- Nowozin, S., Cseke, B., and Tomioka, R. f-gan: Training generative neural samplers using variational divergence minimization. In *Advances in Neural Information Processing Systems*. 2016.
- Oh, J., Singh, S., and Lee, H. Value prediction network. In *Advances in Neural Information Processing Systems*, 2017.
- Pong, V., Gu, S., Dalal, M., and Levine, S. Temporal difference models: Model-free deep RL for model-based control. In *International Conference on Learning Representations*, 2018.
- Rezende, D. and Mohamed, S. Variational inference with normalizing flows. *Proceedings of Machine Learning Research*, 2015.
- Schaul, T., Horgan, D., Gregor, K., and Silver, D. Universal value function approximators. In *Proceedings of the International Conference on Machine Learning*, 2015.
- Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., Guez, A., Lockhart, E., Hassabis, D., Graepel, T., Lillicrap, T., and Silver, D. Mastering atari, go, chess and shogi by planning with a learned model. *arXiv preprint arXiv:1911.08265*, 2019.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- Silver, D., Sutton, R. S., and Müller, M. Sample-based learning and search with permanent and transient memories. In *Proceedings of the International Conference on Machine Learning*, 2008.
- Silver, D., van Hasselt, H., Hessel, M., Schaul, T., Guez, A., Harley, T., Dulac-Arnold, G., Reichert, D., Rabinowitz, N., Barreto, A., and Degris, T. The predictron: End-to-end learning and planning. In *International Conference on Machine Learning*, 2017.
- Sutton, R. TD models: Modeling the world at a mixture of time scales. In *Proceedings of the 12th International Conference on Machine Learning*, 1995.
- Sutton, R. S. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *International Conference on Machine Learning*, 1990.
- Sutton, R. S. Generalization in reinforcement learning: Successful examples using sparse coarse coding. In *Advances in Neural Information Processing Systems*. 1996.
- Sutton, R. S., Precup, D., and Singh, S. Between MDPs and semi-MDPs: A framework for temporal abstraction in reinforcement learning. *Artificial Intelligence*, 112(1):181 – 211, 1999.
- Tamar, A., Wu, Y., Thomas, G., Levine, S., and Abbeel, P. Value iteration networks. In *Advances in Neural Information Processing Systems*. 2016.