

Organizing Experience: A Deeper Look at Replay Mechanisms for Sample-based Planning in Continuous State Domains

Yangchen Pan¹, Muhammad Zaheer¹, Adam White^{1,2}, Andrew Patterson¹, Martha White¹

¹University of Alberta

²DeepMind

pan6@ualberta.ca, mzaheer@ualberta.ca, adamwhite@google.com, andnpatt@iu.edu

Abstract

Model-based strategies for control are critical to obtain sample efficient learning. Dyna is a planning paradigm that naturally interleaves learning and planning, by simulating one-step experience to update the action-value function. This elegant planning strategy has been mostly explored in the tabular setting. The aim of this paper is to revisit sample-based planning, in stochastic and continuous domains with learned models. We first highlight the flexibility afforded by a model over Experience Replay (ER). Replay-based methods can be seen as stochastic planning methods that repeatedly sample from a buffer of recent agent-environment interactions and perform updates to improve data efficiency. We show that a model, as opposed to a replay buffer, is particularly useful for specifying which states to sample from during planning, such as predecessor states that propagate information in reverse from a state more quickly. We introduce a semi-parametric model learning approach, called Reweighted Experience Models (REMs), that makes it simple to sample next states or predecessors. We demonstrate that REM-Dyna exhibits similar advantages over replay-based methods in learning in continuous state problems, and that the performance gap grows when moving to stochastic domains, of increasing size.

1 Introduction

Experience replay has become nearly ubiquitous in modern large-scale, deep reinforcement learning systems [Schaul *et al.*, 2016]. The basic idea is to store an incomplete history of previous agent-environment interactions in a *transition buffer*. During planning, the agent selects a transition from the buffer and updates the value function as if the samples were generated online—the agent *replays* the transition. There are many potential benefits of this approach, including stabilizing potentially divergent non-linear Q-learning updates, and mimicking the effect of multi-step updates as in eligibility traces.

Experience replay (ER) is like a model-based RL system, where the transition buffer acts as a model of the world [Lin, 1992]. Using the data as your model avoids model errors that can cause bias in the updates (c.f. [Bagnell and Schneider,

2001]). One of ER’s most distinctive attributes as a model-based planning method, is that it does not perform multi-step rollouts of hypothetical trajectories according to a model; rather previous agent-environment transitions are replayed randomly or with priority from the transition buffer. Trajectory sampling approaches such as PILCO [Deisenroth and Rasmussen, 2011], Hallucinated-Dagger [Talvitie, 2017], and CPSRs [Hamilton *et al.*, 2014], unlike ER, can rollout unlikely trajectories ending up in hypothetical states that do not match any real state in the world when the model is wrong [Talvitie, 2017]. ER’s stochastic one-step planning approach was later adopted by Sutton’s Dyna architecture [Sutton, 1991].

Despite the similarities between Dyna and ER, there have been no comprehensive, direct empirical comparisons comparing the two and their underlying design-decisions. ER maintains a buffer of transitions for replay, and Dyna a search-control queue composed of stored states and actions from which to sample. There are many possibilities for how to add, remove and select samples from either ER’s transition buffer or Dyna’s search-control queue. It is not hard to imagine situations where a Dyna-style approach could be better than ER. For example, because Dyna models the environment, states leading into high priority states—predecessors—can be added to the queue, unlike ER. Additionally, Dyna can choose to simulate on-policy samples, whereas ER can only replay (likely off-policy) samples previously stored. In non-stationary problems, small changes can be quickly recognized and corrected in the model. On the other hand, these small changes might result in wholesale changes to the policy, potentially invalidating many transitions in ER’s buffer. It remains to be seen if these differences manifest empirically, or if the additional complexity of Dyna is worthwhile.

In this paper, we develop a novel semi-parametric Dyna algorithm, called REM-Dyna, that provides some of the benefits of both Dyna-style planning and ER. We highlight criteria for learned models used within Dyna, and propose Reweighted Experience Models (REMs) that are data-efficient, efficient to sample and can be learned incrementally. We investigate the properties of both ER and REM-Dyna, and highlight cases where ER can fail, but REM-Dyna is robust. Specifically, this paper contributes both (1) a new method extending Dyna to continuous-state domains—significantly outperforming previous attempts [Sutton *et al.*, 2008], and (2) a comprehensive investigation of the design decisions critical to the performance

of one-step, sample-based planning methods for reinforcement learning with function approximation. An **Appendix is publicly available on arXiv**, with theorem proof and additional algorithm and experimental details.

2 Background

We formalize an agent’s interaction with its environment as a discrete time Markov Decision Process (MDP). On each time step t , the agent observes the state of the MDP $S_t \in \mathcal{S}$, and selects an action $A_t \in \mathcal{A}$, causing a transition to a new state $S_{t+1} \in \mathcal{S}$ and producing a scalar reward on the transition $R_{t+1} \in \mathbb{R}$. The agent’s objective is to find an optimal policy $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$, which maximizes the expected return $Q_\pi(s, a)$ for all s, a , where $G_t \stackrel{\text{def}}{=} R_{t+1} + \gamma(S_t, A_t, S_{t+1})G_{t+1}$, $\gamma : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \in [0, 1]$, and $Q_\pi(s, a) = \mathbb{E}[G_t | S_t = s, A_t = a; A_{t+1:\infty} \sim \pi]$, with future states and rewards are sampled according to the one-step dynamics of the MDP. The generalization to a discount function γ allows for a unified specification of episodic and continuing tasks [White, 2017], both of which are considered in this work.

In this paper we are concerned with model-based approaches to finding optimal policies. In all approaches we consider here the agent forms an estimate of the value function from data: $\hat{q}_\pi(S_t, A_t, \theta) \approx \mathbb{E}[G_t | S_t = s, A_t = a]$. The value function is parameterized by $\theta \in \mathbb{R}^n$ allowing both linear and non-linear approximations. We consider sample models, that given an input state and action need only output one possible next state and reward, sampled according to the one-step dynamics of the MDP: $M : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S} \times \mathbb{R}$.

In this paper, we focus on stochastic one-step planning methods, where one-step transitions are sampled from a model to update an action-value function. The agent interacts with the environment on each time step, selecting actions according to its current policy (e.g., ϵ -greedy with respect to \hat{q}_π), observing next states and rewards, and updating \hat{q}_π . Additionally, the agent also updates a model with these *observed sample* transitions $\langle S_t, A_t, S_{t+1}, R_t \rangle$ on each time step. After updating the value function and the model, the agent executes m steps of planning. On each planning step, the agent samples a start state S and action A in some way (called *search control*), then uses the model to *simulate* the next state and reward. Using this *hypothetical transition* the agent updates \hat{q}_π in the usual way. In this generic framework, the agent can interleave learning, planning, and acting—all in realtime. Two well-known implementations of this framework are ER [Lin, 1992], and the Dyna architecture [Sutton, 1991].

3 One-step Sample-based Planning Choices

There are subtle design choices in the construction of stochastic, one-step, sample-based planning methods that can significantly impact performance. These include how to add states and actions to the search-control queue for Dyna, how to select states and actions from the queue, and how to sample next states. These choices influence the design of our REM algorithm, and so we discuss them in this section.

One important choice for Dyna-style methods is whether to sample a next state, or compute an expected update over all possible transitions. A sample-based planner samples s', r, γ ,

given s, a , and stochastically updates $\hat{q}_\pi(s, a, \theta)$. An alternative is to approximate full dynamic programming updates, to give an expected update, as done by stochastic factorization approaches [Barreto *et al.*, 2011; Kveton and Theodorou, 2012; Barreto *et al.*, 2014; Yao *et al.*, 2014; Barreto *et al.*, 2016; Pires and Szepesvári, 2016], kernel-based RL (KBRL) [Ormoneit and Sen, 2002], and kernel mean embeddings (KME) for RL [Grunewalder *et al.*, 2012; Van Hoof *et al.*, 2015; Lever *et al.*, 2016]. Linear Dyna [Sutton *et al.*, 2008] computes an expected next reward and expected next feature vector for the update, which corresponds to an expected update when \hat{q}_π is a linear function of features. We advocate for a sampled update, because approximate dynamic programming updates, such as KME and KBRL, are typically too expensive, couple the model and value function parameterization and are designed for a batch setting. Computation can be more effectively used by sampling transitions.

There are many possible refinements to the search-control mechanism, including prioritization and backwards-search. For tabular domains, it is feasible to simply store all possible states and actions, from which to simulate. In continuous domains, however, care must be taken to order and delete stored samples. A basic strategy is to simply store recent transitions (s, a, s', r, γ) for the transition buffer in ER, or state and actions (s, a) for the search-control queue in Dyna. This, however, provides little information about which samples would be most beneficial for learning. Prioritizing how samples are drawn, based on absolute TD-error $|R + \gamma \max_{a'} \hat{q}_\pi(s', a') - \hat{q}_\pi(s, a)|$, has been shown to be useful for both tabular Dyna [Sutton and Barto, 1998], and ER with function approximation [Schaul *et al.*, 2016]. When the buffer or search-control queue gets too large, one then must also decide whether to delete transitions based on recency or priority. In the experiments, we explore this question about the efficacy of recency versus priorities for adding and deleting.

ER is limited in using alternative criteria for search-control, such as backward search. A model allows more flexibility in obtaining useful states and action to add to the search-control queue. For example, a model can be learned to simulate *predecessor states*—states leading into (a high-priority) s for a given action a . Predecessor states can be added to the search-control queue during planning, facilitating a type of backward search. The idea of backward search and prioritization were introduced together for tabular Dyna [Peng and Williams, 1993; Moore and Atkeson, 1993]. Backward search can only be applied in ER in a limited way because its buffer is unlikely to contain transitions from multiple predecessor states to the current state in planning. [Schaul *et al.*, 2016] proposed a simple heuristic to approximate prioritization with predecessors, by updating the priority of the most recent transition on the transition buffer to be at least as large as the transition that came directly after it. This heuristic, however, does not allow a systematic backward-search.

A final possibility we consider is using the current policy to select the actions during search control. Conventionally, Dyna draws the action from the search-control queue using the same mechanism used to sample the state. Alternatively, we can sample the state via priority or recency, and then query the model using the action the learned policy would select in the

current state: $s', R, \gamma \sim \hat{P}(s, \pi(s), \cdot, \cdot, \cdot)$. This approach has the advantage that planning focuses on actions that the agent currently estimates to be the best. In the tabular setting, this *on-policy sampling* can result in dramatic efficiency improvements for Dyna [Sutton and Barto, 1998], while [Gu *et al.*, 2016] report improvement from on-policy sample of transitions, in a setting with multi-step rollouts. ER cannot emulate on-policy search control because it replays full transitions (s, a, s', r, γ) , and cannot query for an alternative transition if a different action than a is taken.

4 Reweighted Experience Models for Dyna

In this section, we highlight criteria for selecting amongst the variety of available sampling models, and then propose a semi-parametric model—called Reweighted Experience Models—as one suitable model that satisfies these criteria.

4.1 Generative Models for Dyna

Generative models are a fundamental tool in machine learning, providing a wealth of possible model choices. We begin by specifying our desiderata for online sample-based planning and acting. First, the model learning should be *incremental and adaptive*, because the agent incrementally interleaves learning and planning. Second, the models should be *data-efficient*, in order to achieve the primary goal of improving data-efficiency of learning value functions. Third, due to policy non-stationarity, the models need to be *robust to forgetting*: if the agent stays in a part of the world for quite some time, the learning algorithm should not overwrite—or forget—the model in other parts of the world. Fourth, the models need to be able to be queried as *conditional models*. Fifth, *sampling should be computationally efficient*, since a slow sampler will reduce the feasible number of planning steps.

Density models are typically learned as a mixture of simpler functions or distributions. In the most basic case, a simple distributional form can be used, such as a Gaussian distribution for continuous random variables, or a categorical distribution for discrete random variables. For conditional distributions, $p(s'|s, a)$, the parameters to these distributions, like the mean and variance of s' , can be learned as a (complex) function of s, a . More general distributions can be learned using mixtures, such as mixture models or belief networks. A Conditional Gaussian Mixture Model, for example, could represent $p(s'|s, a) = \sum_{i=1}^b \alpha_i(s, a) \mathcal{N}(s' | \mu(s, a), \Sigma(s, a))$, where α_i, μ and Σ are (learned) functions of s, a . In belief networks—such as Boltzmann distributions—the distribution is similarly represented as a sum over hidden variables, but for more general functional forms over the random variables—such as energy functions. To condition on s, a , those variables in the network are fixed both for learning and sampling.

Kernel density estimators (KDE) are similar to mixture models, but are non-parametric: means in the mixture are the training data, with a uniform weighting: $\alpha_i = 1/T$ for T samples. KDE and conditional KDE is consistent [Holmes *et al.*, 2007]—since the model is a weighting over observed data—providing low model-bias. Further, it is data-efficient, easily enables conditional distributions, and is well-understood theoretically and empirically. Unfortunately, it scales linearly

in the data, which is not compatible with online reinforcement learning problems. Mixture models, on the other hand, learn a compact mixture and could scale, but are expensive to train incrementally and have issues with local minima.

Neural network models are another option, such as Generative Adversarial Networks [Goodfellow *et al.*, 2014] and Stochastic Neural Networks [Sohn *et al.*, 2015; Alain *et al.*, 2016]. Many of the energy-based models, however, such as Boltzmann distributions, require computationally expensive sampling strategies [Alain *et al.*, 2016]. Other networks—such as Variational Auto-encoders—sample inputs from a given distribution, to enable the network to sample outputs. These neural network models, however, have issues with forgetting [McCloskey and Cohen, 1989; French, 1999; Goodfellow *et al.*, 2013], and require more intensive training strategies—often requiring experience replay themselves.

4.2 Reweighted Experience Models

We propose a semi-parametric model to take advantage of the properties of KDE and still scale with increasing experience. The key properties of REM models are that 1) it is straightforward to specify and sample both forward and reverse models for predecessors— $p(s'|s, a)$ and $p(s|s', a)$ —using essentially the same model (the same prototypes); 2) they are data-efficient, requiring few parameters to be learned; and 3) they can provide sufficient model complexity, by allowing for a variety of kernels or metrics defining similarity.

REM models consist of a subset of prototype transitions $\{(s_i, a_i, s'_i, r_i, \gamma_i)\}_{i=1}^b$, chosen from all T transitions experienced by the agent, and their corresponding weights $\{c_i\}_{i=1}^b$. These prototypes are chosen to be representative of the transitions, based on a similarity given by a *product kernel* k

$$p(s, a, s', r, \gamma | s_i, a_i, s'_i, r_i, \gamma_i) \stackrel{\text{def}}{=} k((s, a, s', r, \gamma), (s_i, a_i, s'_i, r_i, \gamma_i)) \stackrel{\text{def}}{=} k_s(s, s_i) k_a(a, a_i) k_{s', r, \gamma}((s', r, \gamma), (s'_i, r_i, \gamma_i)). \quad (1)$$

A product kernel is a product of separate kernels. It is still a valid kernel, but simplifies dependences and simplifies computing conditional densities, which are key for Dyna, both for forward and predecessor models. They are also key for obtaining a consistent estimate of the $\{c_i\}_{i=1}^b$, described below.

We first consider Gaussian kernels for simplicity. For states, $k_s(s, s_i) = (2\pi)^{-d/2} |\mathbf{H}_s|^{-1/2} \exp(-(\mathbf{s} - \mathbf{s}_i)^\top \mathbf{H}_s^{-1} (\mathbf{s} - \mathbf{s}_i))$ with covariance \mathbf{H}_s . For discrete actions, the similarity is an indicator $k_a(a, a_i) = 1$ if $a = a_i$ and otherwise 0. For next state, reward and discount, a Gaussian kernel is used for $k_{s', r, \gamma}$ with covariance $\mathbf{H}_{s', r, \gamma}$. We set the covariance matrix $\mathbf{H}_s = b^{-1} \Sigma_s$, where Σ_s is a sample covariance, and use a conditional covariance for (s, r, γ) .

First consider a KDE model, for comparison, where all experience is used to define the distribution

$$p_k(s, a, s', r, \gamma) = \frac{1}{T} \sum_{i=1}^T k((s, a, s', r, \gamma), (s_i, a_i, s'_i, r_i, \gamma_i))$$

This estimator puts higher density around more frequently observed transitions. A conditional estimator is similarly intuitive, and also a consistent estimator [Holmes *et al.*, 2007],

$$N_k(s, a) = \frac{1}{T} \sum_{i=1}^T k_s(s, s_i) k_a(a, a_i)$$

$$p_k(s', r, \gamma | s, a) = \frac{1}{N_k(s, a)} \sum_{i=1}^T k_s(s, s_i) k_a(a, a_i) k_{s', r, \gamma}((s', r, \gamma), (s'_i, r_i, \gamma_i))$$

The experience (\mathbf{s}_i, a_i) similar to (\mathbf{s}, a) has higher weight in the conditional estimator: distributions centered at $(\mathbf{s}'_i, r_i, \gamma_i)$ contribute more to specifying $p(\mathbf{s}', r, \gamma | \mathbf{s}, a)$. Similarly, it is straightforward to specify the conditional density $p(\mathbf{s} | \mathbf{s}', a)$.

When only prototype transitions are stored, joint and conditional densities can be similarly specified, but prototypes must be weighted to reflect the density in that area. We therefore need a method to select prototypes and to compute weightings. Selecting representative prototypes or centers is a very active area of research, and we simply use a recent incremental and efficient algorithm designed to select prototypes [Schlegel *et al.*, 2017]. For the reweighting, however, we can design a more effective weighting exploiting the fact that we will only query the model using conditional distributions.

Reweighting approach. We develop a reweighting scheme that takes advantage of the fact that Dyna only requires conditional models. Because $p(\mathbf{s}', r, \gamma | \mathbf{s}, a) = p(\mathbf{s}, a, \mathbf{s}', r, \gamma) / p(\mathbf{s}, a)$, a simple KDE strategy is to estimate coefficients p_i on the entire transition $(\mathbf{s}_i, a_i, \mathbf{s}'_i, r_i, \gamma_i)$ and q_i on (\mathbf{s}_i, a_i) , to obtain accurate densities $p(\mathbf{s}, a, \mathbf{s}', r, \gamma)$ and $p(\mathbf{s}, a)$. However, there are several disadvantages to this approach. The p_i and q_i need to constantly adjust, because the policy is changing. Further, when adding and removing prototypes incrementally, the other p_i and q_i need to be adjusted. Finally, p_i and q_i can get very small, depending on visitation frequency to a part of the environment, even if p_i/q_i is not small. Rather, by directly estimating the conditional coefficients $c_i = p(\mathbf{s}'_i, r_i, \gamma_i | \mathbf{s}_i, a_i)$, we avoid these problems. The distribution $p(\mathbf{s}', r, \gamma | \mathbf{s}, a)$ is stationary even with a changing policy; each c_i can converge even during policy improvement and can be estimated independently of the other c_i .

We can directly estimate c_i , because of the conditional independence assumption made by product kernels. To see why, for prototype i in the product kernel in Equation (1),

$$p(\mathbf{s}, a, \mathbf{s}', r, \gamma | \mathbf{s}_i, a_i, \mathbf{s}'_i, r_i, \gamma_i) = p(\mathbf{s}, a | \mathbf{s}_i, a_i) p(\mathbf{s}', r, \gamma | \mathbf{s}'_i, r_i, \gamma_i)$$

Rewriting $p(\mathbf{s}_i, a_i, \mathbf{s}'_i, r_i, \gamma_i) = c_i p(\mathbf{s}_i, a_i)$ and because $p(\mathbf{s}, a | \mathbf{s}_i, a_i) p(\mathbf{s}_i, a_i) = p(\mathbf{s}_i, a_i | \mathbf{s}, a) p(\mathbf{s}, a)$, we can rewrite the probability as

$$\begin{aligned} p(\mathbf{s}, a, \mathbf{s}', r, \gamma) &= \sum_{i=1}^b [c_i p(\mathbf{s}_i, a_i)] p(\mathbf{s}, a | \mathbf{s}_i, a_i) p(\mathbf{s}', r, \gamma | \mathbf{s}'_i, r_i, \gamma_i) \\ &= \sum_{i=1}^b c_i p(\mathbf{s}', r, \gamma | \mathbf{s}'_i, r_i, \gamma_i) p(\mathbf{s}_i, a_i | \mathbf{s}, a) p(\mathbf{s}, a) \end{aligned}$$

$$\begin{aligned} \text{giving } p(\mathbf{s}', r, \gamma | \mathbf{s}, a) &= \frac{1}{p(\mathbf{s}, a)} \sum_{i=1}^b c_i p(\mathbf{s}', r, \gamma | \mathbf{s}'_i, r_i, \gamma_i) p(\mathbf{s}_i, a_i | \mathbf{s}, a) p(\mathbf{s}, a) \\ &= \sum_{i=1}^b c_i k_{\mathbf{s}', r, \gamma}((\mathbf{s}', r, \gamma), (\mathbf{s}'_i, r_i, \gamma_i)) k_{\mathbf{s}}(\mathbf{s}_i, \mathbf{s}) k_a(a_i, a). \end{aligned}$$

Now we simply need to estimate c_i . Again using the conditional independence property, we can prove the following.

Theorem 1. Let $\rho(t, i) \stackrel{\text{def}}{=} k_{\mathbf{s}, a}((\mathbf{s}_t, a_t), (\mathbf{s}_i, a_i))$ be the similarity of \mathbf{s}_t, a_t for sample t to \mathbf{s}_i, a_i for prototype i . Then

$$c_i = \underset{c_i \geq 0}{\operatorname{argmin}} \sum_{t=1}^T (c_i - k_{\mathbf{s}', r, \gamma}((\mathbf{s}'_t, r_t, \gamma_t), (\mathbf{s}'_i, r_i, \gamma_i)))^2 \rho(t, i)$$

is a consistent estimator of $p(\mathbf{s}'_i, r_i, \gamma_i | \mathbf{s}_i, a_i)$.

The proof for this theorem, and a figure demonstrating the difference between KDE and REM, are provided in the appendix.

Though there is a closed form solution to this objective, we use an incremental stochastic update to avoid storing additional variables and for the model to be more adaptive. For each transition, the c_i are updated for each prototype as

$$c_i \leftarrow (1 - \rho(t, i)) c_i + \rho(t, i) k_{\mathbf{s}', r, \gamma}((\mathbf{s}'_t, r_t, \gamma_t), (\mathbf{s}'_i, r_i, \gamma_i))$$

The resulting REM model is

$$\begin{aligned} \beta_i(\mathbf{s}, a) &\stackrel{\text{def}}{=} \frac{c_i}{N(\mathbf{s}, a)} k_{\mathbf{s}}(\mathbf{s}, \mathbf{s}_i) k_a(a, a_i) \\ \text{where } N(\mathbf{s}, a) &\stackrel{\text{def}}{=} \sum_{i=1}^b c_i k_{\mathbf{s}}(\mathbf{s}, \mathbf{s}_i) k_a(a, a_i) \\ p(\mathbf{s}', r, \gamma | \mathbf{s}, a) &\stackrel{\text{def}}{=} \sum_{i=1}^b \beta_i(\mathbf{s}, a) k_{\mathbf{s}', r, \gamma}((\mathbf{s}', r, \gamma), (\mathbf{s}'_i, r_i, \gamma_i)). \end{aligned}$$

To sample predecessor states, with $p(\mathbf{s} | \mathbf{s}', a)$, the same set of b prototypes can be used, with a separate set of conditional weightings estimated as $c'_i \leftarrow (1 - \rho'(t, i)) c'_i + \rho'(t, i) k_{\mathbf{s}}(\mathbf{s}, \mathbf{s}_i)$ for $\rho'(t, i) \stackrel{\text{def}}{=} k_{\mathbf{s}}(\mathbf{s}', \mathbf{s}'_i) k_a(a, a_i)$.

Sampling from REMs. Conveniently, to sample from the REM conditional distribution, the similarity across next states and rewards need not be computed. Rather, only the coefficients $\beta_i(\mathbf{s}, a)$ need to be computed. A prototype is sampled with probability $\beta_i(\mathbf{s}, a)$; if prototype j is sampled, then the density (Gaussian) centered around $(\mathbf{s}'_j, r_j, \gamma_j)$ is sampled.

In the implementation, the terms $(2\pi)^{-d/2} |\mathbf{H}_{\mathbf{s}}|^{-1/2}$ in the Gaussian kernels are omitted, because as fixed constants they can be normalized out. All kernel values then are in $[0, 1]$, providing improved numerical stability and the straightforward initialization $c_i = 1$ for new prototypes. REMs are linear in the number of prototypes, for learning and sampling, with complexity per-step independent of the number of samples.

Addressing issues with scaling with input dimension. In general, any nonnegative kernel $k_{\mathbf{s}}(\cdot, \mathbf{s})$ that integrates to one is possible. There are realistic low-dimensional physical systems for which Gaussian kernels have been shown to be highly effective, such as in robotics [Deisenroth and Rasmussen, 2011]. Kernel-based approaches can, however, extend to high-dimensional problems with specialized kernels. For example, convolutional kernels for images have been shown to be competitive with neural networks [Mairal *et al.*, 2014]. Further, learned similarity metrics or embeddings enable data-driven models—such as neural networks—to improve performance, by replacing the Euclidean distance. This combination of probabilistic structure from REMs and data-driven similarities for neural networks is a promising next step.

5 Experiments

We first empirically investigate the design choices for ER's buffer and Dyna's search-control queue in the tabular setting. Subsequently, we examine the utility of REM-Dyna, our proposed model-learning technique, by comparing it with ER and other model learning techniques in the function approximation setting. Maintaining the buffer or queue involves determining how to add and remove samples, and how to prioritize samples. All methods delete the oldest samples. Our experiments (not shown here), showed that deleting samples of lowest priority—computed from TD error—is not effective in the problems we studied. We investigate three different settings:

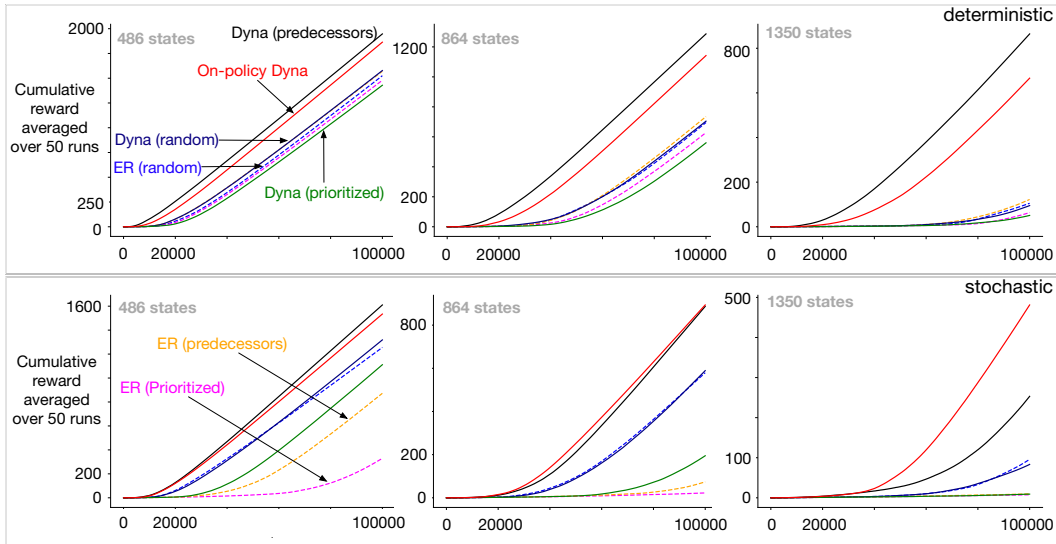


Figure 1: Learning curves for a varying number of states, for the deterministic (upper) and stochastic (lower) gridworlds. The x-axis shows the number of interactions with the environment

- 1) *Random*: samples are drawn randomly.
- 2) *Prioritized*: samples are drawn probabilistically according to the absolute TD error of the transitions [Schaul *et al.*, 2016, Equation 1] (exponent = 1).
- 3) *Predecessors*: same as *Prioritized*, and predecessors of the current state are also added to the buffer or queue.

We also test using *On-policy* transitions for Dyna, where only s is stored on the queue and actions simulated according to the current policy; the queue is maintained using priorities and predecessors. In Dyna, we use the learned model to sample predecessors s of the current s' , for all actions a , and add them to the queue. In ER, with no environment model, we use a simple heuristic which adds the priority of the current sample to the preceding sample in the buffer [Schaul *et al.*, 2016]. Note that [van Seijen and Sutton, 2015] relate Dyna and ER, but specifically for a theoretical equivalence in policy evaluation based on a non-standard form of replay related to true online methods, and thus we do not include it.

Experimental settings: All experiments are averaged over many independent runs, with the randomness controlled based on the run number. All learning algorithms use ϵ -greedy action selection ($\epsilon = 0.1$) and Q-learning to update the value function in both *learning* and *planning* phases. The step-sizes are swept in $[0.1, 1.0]$. The size of the search-control queue and buffer was fixed to 1024—large enough for the micro-worlds considered—and the number of planning steps was fixed to 5.

A natural question is if the conclusions from experiments in the below microworlds extend to larger environments. Microworlds are specifically designed to highlight phenomena in larger domains, such as creating difficult-to-reach, high-reward states in River Swim described below. The computation and model size are correspondingly scaled down, to reflect realistic limitations when moving to larger environments. The trends obtained when varying the size and stochasticity of these environments provides insights into making such changes in larger environments. Experiments, then, in microworlds enable a

more systematic issue-oriented investigation and suggest directions for further investigation for use in real domains.

Results in the Tabular Setting: To gain insight into the differences between Dyna and ER, we first consider them in the deterministic and stochastic variants of a simple gridworld with increasing state space size. ER has largely been explored in deterministic problems, and most work on Dyna has only considered the tabular setting. The gridworld is discounted with $\gamma = 0.95$, and episodic with obstacles and one goal, in which case the reward is +100. The agent can take four actions. In the stochastic variant each action takes the agent to the intended next state with probability 0.925, or one of the other three adjacent states with probability 0.025. In the deterministic setting, Dyna uses a table to store next state and reward for each state and action; in stochastic, it estimates the probabilities of each observed transition via transition counts.

Figure 1 shows the reward accumulated by each agent over 100,000 time-steps. We observe that: 1) Dyna with priorities and predecessors outperformed all variants of ER, and the performance gap increases with gridworld size. 2) TD-error based prioritization on Dyna’s search control queue improved performance only when combined with the addition of predecessors; otherwise, unprioritized variants outperformed prioritized variants. We hypothesize that this could be due to out-dated priorities, previously suggested to be problematic [Peng and Williams, 1993; Schaul *et al.*, 2016]. 3) ER with prioritization performs slightly worse than unprioritized ER variants for the deterministic setting, but its performance degrades considerably in the stochastic setting. 4) On-Policy Dyna with priorities and predecessors outperformed the regular variant in the stochastic domain with a larger state space. 5) Dyna with similar search-control strategies to ER, such as recency and priorities, does not outperform ER; only with the addition of improved search-control strategies is there an advantage. 6) Deleting samples from the queue or transitions from the buffer according to recency was always better than




Figure 2: (a) compares variants of ER and REM-Dyna. REM-Dyna with predecessor states and Random ER accumulate significantly more reward than all other variants, with REM-Dyna statistically significantly better (non-overlapping confidence intervals) than ER by the end of the run. (b) shows Dyna with different models. REM-Dyna is statistically significantly better than NN-Dyna and Linear-Dyna. For NNs and REMs, using predecessors is significantly better, unlike Linear-Dyna which learns inaccurate models. (c) Results on River Swim, with number of steps required to obtain a ratio of 80%, 85% and 90% between the cumulative reward for the agent relative to the cumulative reward of the optimal policy. If there is no entry, then the agent was unable to achieve that performance within the 20,000 learning steps.

deleting according to priority for both Dyna and ER.

Results for Continuous States. We recreate the above experiments for continuous states, and additionally explore the utility of REMs for Dyna. We compare to using a Neural Network model—with two layers, trained with the Adam optimizer on a sliding buffer of 1000 transitions—and to a Linear model predicting features-to-expected next features rather than states, as in Linear Dyna. We improved upon the original Linear Dyna by learning a reverse model and sweeping different step-sizes for the models and updates to \hat{q}_π .

We conduct experiments in two tasks: a Continuous Gridworld and River Swim. Continuous Gridworld is a continuous variant of a domain introduced by [Peng and Williams, 1993], with $(x, y) \in [0, 1]^2$, a sparse reward of 1 at the goal, and a long wall with a small opening. Agents can choose to move 0.05 units up, down, left, right, which is executed successfully with probability 0.9 and otherwise the environment executes a random move. Each move has noise $\mathcal{N}(0, 0.01)$. River Swim is a difficult exploration domain, introduced as a tabular domain [Strehl and Littman, 2008], as a simple simulation of a fish swimming up a river. We modify it to have a continuous state space $[0, 1]$. On each step, the agent can go right or left, with the river pushing the agent towards the left. The right action succeeds with low probability depending on the position, and the left action always succeeds. There is a small reward 0.005 at the leftmost state (close to 0), and a relatively large reward 1.0 at the rightmost state (close to 1). The optimal policy is to constantly select right. Because exploration is difficult in this domain, instead of ϵ -greedy, we induced a bit of extra exploration by initializing the weights to 1.0. For both domains, we use a coarse tile-coding, similar to state-aggregation.

REM-Dyna obtains the best performance on both domains, in comparison to the ER variants and other model-based approaches. For search-control in the continuous state domains, the results in Figures 2 parallels the conclusions from the tabular case. For the alternative models, REMs outperform both Linear models and NN models. For Linear models, the model-accuracy was quite low and the step-size selection sen-

sitive. We hypothesize that this additional tuning inadvertently improved the Q-learning update, rather than gaining from Dyna-style planning; in River Swim, Linear Dyna did poorly. Dyna with NNs performs poorly because the NN model is not data-efficient; after several 1000s of more learning steps, however, the model does finally become accurate. This highlights the necessity for data-efficient models, for Dyna to be effective. In Riverswim, no variant of ER was within 85% of optimal, in 20,000 steps, whereas all variants of REM-Dyna were, once again particularly for REM-Dyna with Predecessors.

6 Conclusion

In this work, we developed a semi-parametric model learning approach, called Reweighted Experience Models (REMs), for use with Dyna for control in continuous state settings. We revisited a few key dimensions for maintaining the search-control queue for Dyna, to decide how to select states and actions from which to sample. These included understanding the importance of using recent samples, prioritizing samples (with absolute TD-error), generating predecessor states that lead into high-priority states, and generating on-policy transitions. We compared Dyna to the simpler alternative, Experience Replay (ER), and considered similar design decisions for its transition buffer. We highlighted several criteria for the model to be useful in Dyna, for one-step sampled transitions, namely being data-efficient, robust to forgetting, enabling conditional models and being efficient to sample. We developed a new semi-parametric model, REM, that uses similarities to a representative set of prototypes, and requires only a small set of coefficients to be learned. We provided a simple learning rule for these coefficients, taking advantage of a conditional independence assumption and that we only require conditional models. We thoroughly investigate the differences between Dyna and ER, in several microworlds for both tabular and continuous states, showing that Dyna can provide significant gains through the use of predecessors and on-policy transitions. We further highlight that REMs are an effective model for Dyna, compared to using a Linear model or a Neural Network model.

References

- [Alain *et al.*, 2016] Guillaume Alain, Yoshua Bengio, Li Yao, Jason Yosinski, Éric Thibodeau-Laufer, Saizheng Zhang, and Pascal Vincent. GSNs: generative stochastic networks. *Information and Inference: A Journal of the IMA*, 2016.
- [Bagnell and Schneider, 2001] J A Bagnell and J G Schneider. Autonomous helicopter control using reinforcement learning policy search methods. In *IEEE International Conference on Robotics and Automation*, 2001.
- [Barreto *et al.*, 2011] A Barreto, D Precup, and J Pineau. Reinforcement Learning using Kernel-Based Stochastic Factorization. In *Advances in Neural Information Processing Systems*, 2011.
- [Barreto *et al.*, 2014] A Barreto, J Pineau, and D Precup. Policy Iteration Based on Stochastic Factorization. *Journal of Artificial Intelligence Research*, 2014.
- [Barreto *et al.*, 2016] A Barreto, R Beirigo, J Pineau, and D Precup. Incremental Stochastic Factorization for Online Reinforcement Learning. In *AAAI Conference on Artificial Intelligence*, 2016.
- [Deisenroth and Rasmussen, 2011] M Deisenroth and C E Rasmussen. PILCO: A model-based and data-efficient approach to policy search. In *International Conference on Machine Learning*, 2011.
- [French, 1999] R M French. Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4):128–135, 1999.
- [Goodfellow *et al.*, 2013] I J Goodfellow, M Mirza, D Xiao, A Courville, and Y Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013.
- [Goodfellow *et al.*, 2014] I J Goodfellow, J Pouget-Abadie, MehMdi Mirza, B Xu, D Warde-Farley, S Ozair, A C Courville, and Y Bengio. Generative Adversarial Nets. In *Advances in Neural Information Processing Systems*, 2014.
- [Grunewalder *et al.*, 2012] Steffen Grunewalder, Guy Lever, Luca Baldassarre, Massi Pontil, and Arthur Gretton. Modelling transition dynamics in MDPs with RKHS embeddings. In *International Conference on Machine Learning*, 2012.
- [Gu *et al.*, 2016] Shixiang Gu, Timothy P Lillicrap, Ilya Sutskever, and Sergey Levine. Continuous Deep Q-Learning with Model-based Acceleration. In *International Conference on Machine Learning*, 2016.
- [Hamilton *et al.*, 2014] W L Hamilton, M M Fard, and J Pineau. Efficient learning and planning with compressed predictive states. *Journal of Machine Learning Research*, 2014.
- [Holmes *et al.*, 2007] Michael P Holmes, Alexander G Gray, and Charles L Isbell. Fast Nonparametric Conditional Density Estimation. *Uncertainty in AI*, 2007.
- [Kveton and Theodorou, 2012] B Kveton and G Theodorou. Kernel-Based Reinforcement Learning on Representative States. *AAAI Conference on Artificial Intelligence*, 2012.
- [Lever *et al.*, 2016] Guy Lever, John Shawe-Taylor, Ronnie Stafford, and Csaba Szepesvári. Compressed Conditional Mean Embeddings for Model-Based Reinforcement Learning. In *AAAI Conference on Artificial Intelligence*, 2016.
- [Lin, 1992] Long-Ji Lin. Self-Improving Reactive Agents Based On Reinforcement Learning, Planning and Teaching. *Machine Learning*, 1992.
- [Mairal *et al.*, 2014] Julien Mairal, Piotr Koniusz, Zaid Harchaoui, and Cordelia Schmid. Convolutional Kernel Networks. *Advances in Neural Information Processing Systems*, 2014.
- [McCloskey and Cohen, 1989] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of learning and motivation*, 24:109–165, 1989.
- [Moore and Atkeson, 1993] Andrew W Moore and Christopher G Atkeson. Prioritized sweeping: Reinforcement learning with less data and less time. *Machine learning*, 13(1):103–130, 1993.
- [Ormoneit and Sen, 2002] Dirk Ormoneit and Šaunak Sen. Kernel-Based Reinforcement Learning. *Machine Learning*, 2002.
- [Peng and Williams, 1993] Jing Peng and Ronald J Williams. Efficient Learning and Planning Within the Dyna Framework. *Adaptive behavior*, 1993.
- [Pires and Szepesvári, 2016] Bernardo Avila Pires and Csaba Szepesvári. Policy Error Bounds for Model-Based Reinforcement Learning with Factored Linear Models. In *Annual Conference on Learning Theory*, 2016.
- [Schaul *et al.*, 2016] Tom Schaul, John Quan, Ioannis Antonoglou, and David Silver. Prioritized Experience Replay. In *International Conference on Learning Representations*, 2016.
- [Schlegel *et al.*, 2017] Matthew Schlegel, Yangchen Pan, Jiecao Chen, and Martha White. Adapting Kernel Representations Online Using Submodular Maximization. In *International Conference on Machine Learning*, 2017.
- [Sohn *et al.*, 2015] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning Structured Output Representation using Deep Conditional Generative Models. In *Advances in Neural Information Processing Systems*, 2015.
- [Strehl and Littman, 2008] A. Strehl and M Littman. An analysis of model-based Interval Estimation for Markov Decision Processes. *Journal of Computer and System Sciences*, 2008.
- [Sutton and Barto, 1998] R.S. Sutton and A G Barto. *Reinforcement Learning: An Introduction*. MIT press, 1998.
- [Sutton *et al.*, 2008] R Sutton, C Szepesvári, A Geramifard, and M Bowling. Dyna-style planning with linear function approximation and prioritized sweeping. In *Conference on Uncertainty in Artificial Intelligence*, 2008.
- [Sutton, 1991] R.S. Sutton. Integrated modeling and control based on reinforcement learning and dynamic programming. In *Advances in Neural Information Processing Systems*, 1991.
- [Talvitie, 2017] Erik Talvitie. Self-Correcting Models for Model-Based Reinforcement Learning. In *AAAI Conference on Artificial Intelligence*, 2017.
- [Van Hoof *et al.*, 2015] H Van Hoof, J. Peters, and G Neumann. Learning of Non-Parametric Control Policies with High-Dimensional State Features. *AI and Statistics*, 2015.
- [van Seijen and Sutton, 2015] H van Seijen and R.S. Sutton. A deeper look at planning as learning from replay. In *International Conference on Machine Learning*, 2015.
- [White, 2017] Martha White. Unifying task specification in reinforcement learning. In *International Conference on Machine Learning*, 2017.
- [Yao *et al.*, 2014] Hengshuai Yao, Csaba Szepesvári, Bernardo Avila Pires, and Xinhua Zhang. Pseudo-MDPs and factored linear action models. In *ADPRL*, 2014.