

---

# Implicit Generative Modeling for Efficient Exploration

---

Neale Ratzlaff<sup>\*1</sup> Qinxun Bai<sup>2</sup> Li Fuxin<sup>1</sup> Wei Xu<sup>2</sup>

## Abstract

Efficient exploration remains a challenging problem in reinforcement learning, especially for those tasks where rewards from environments are sparse. In this work, we introduce an exploration approach based on a novel implicit generative modeling algorithm to estimate a Bayesian uncertainty of the agent’s belief of the environment dynamics. Each random draw from our generative model is a neural network that instantiates the dynamic function, hence multiple draws would approximate the posterior, and the variance in the predictions based on this posterior is used as an intrinsic reward for exploration. We design a training algorithm for our generative model based on the amortized Stein Variational Gradient Descent. In experiments, we demonstrate the effectiveness of this exploration algorithm in both pure exploration tasks and a downstream task, comparing with state-of-the-art intrinsic reward-based exploration approaches, including two recent approaches based on an ensemble of dynamic models. In challenging exploration tasks, our implicit generative model consistently outperforms competing approaches regarding data efficiency in exploration.

## 1. Introduction

Deep Reinforcement Learning (RL) has enjoyed recent success in a variety of applications, including super-human performance in Atari games (Mnih et al., 2013), robotic control (Lillicrap et al., 2015), image-based control tasks (Hafner et al., 2019), and playing the game of Go (Silver et al., 2016). Despite these achievements, many recent deep RL techniques still suffer from poor sample efficiency. Agents are often trained for millions, or even bil-

ions of simulation steps before achieving reasonable performance (Burda et al., 2018a). This lack of statistical efficiency makes it difficult to apply deep RL to real-world tasks, as the cost of acting in the real world is far greater than in a simulator. It is then a problem of utmost importance to design agents that make efficient use of collected data. In this work, we focus on efficient exploration which is widely considered to be one of the three key aspects in building a data-efficient agent (Sutton & Barto, 2018).

In particular, we focus on those challenging environments with sparse external rewards. In those environments, it is important for an effective agent to methodically explore a significant portion of the state space, since there may not be enough signals to indicate where the reward might be. Previous work usually utilize some sort of intrinsic reward driven by the uncertainty in an agent’s belief of the environment state (Osband et al., 2018). Intuitively, agents should explore more around states where they are not certain whether there could exist a previously unknown consequence – which could be an unexpected extrinsic reward. However, uncertainty modeling from a deep network has proven to be difficult with no approach (Snoek et al., 2019) that is proven to be universally applicable.

In this work, we introduce a new framework of Bayesian uncertainty modeling for intrinsic reward-based exploration in deep RL. The main component of our framework is a network generator, each draw of which is a neural network that serves as the dynamic function for the environment. Multiple draws approximate a posterior of the dynamic model, and the variance in the future state predictions based on this posterior is used as an intrinsic reward for exploration. In doing so, our framework characterizes the uncertainty of the agent’s belief of the environment dynamics in a non-parametric manner, avoiding restrictive distributional assumptions on the posterior, and explore a significantly larger model space than previous approaches. Recently, it has been shown (Ratzlaff & Fuxin, 2019) that training these kinds of generators can be done in classification problems and the resulting network samples can represent a rich distribution of diverse networks that perform approximately equally well on the classification task.

For our goal of training this generator for the dynamic function, we propose a new algorithm to optimize the KL di-

---

<sup>1</sup>Department of Electrical Engineering and Computer Science, Oregon State University, Corvallis, Oregon, USA <sup>2</sup>Horizon Robotics, Cupertino, California, USA. Correspondence to: Neale Ratzlaff <ratzlafn@oregonstate.edu>.

vergence between the implicit distribution (represented by draws from the generator) and the true posterior of the dynamic model (given the agent’s experience) via amortized Stein Variational Gradient Descent (SVGD) (Liu & Wang, 2016; Feng et al., 2017). Amortized SVGD allows direct minimization of the KL divergence between the implicit posterior and true posterior without parametric assumptions or Evidence Lower Bound (ELBO) approximations, and projects to a finite-dimensional parameter update.

Comparing with recent work (Pathak et al., 2019; Shyam et al., 2019) that maintain an ensemble of dynamic models and use the divergence or disagreement among them as an intrinsic reward for exploration, our implicit modeling of the posterior has two major advantages: First, it is a more flexible framework for approximating the model posterior compared to an ensemble-based approximation. After one training episode, it can provide an unlimited amount of draws whereas for an ensemble each draw would require independent training. Second, amortized SVGD (Feng et al., 2017) allows direct nonparametric minimization of the KL divergence, in contrast with existing ensemble-based methods that rely on the random initialization and/or bootstrapped experience sampling, which does not necessarily approximate the posterior.

In our experiments, we compare our approach with several state-of-the-art intrinsic reward-based exploration approaches, including two recent approaches that also leverage the uncertainty in dynamic models. Experiments show that our implementation consistently outperforms competing methods regarding data efficiency in exploration.

In summary, our contributions are:

- We propose a generative framework leveraging amortized SVGD to implicitly approximate the posterior of network parameters. Applying this framework to generate dynamic models of the environment, the uncertainty from the approximate posterior is used as an intrinsic reward for efficient exploration in deep RL.
- We evaluate on three challenging exploration tasks and compare with three state-of-the-art intrinsic reward-based methods, two of which are also based on uncertainty in dynamic models. The superior performance of our method shows the effectiveness of the proposed framework in estimating the Bayesian uncertainty in the dynamic model for efficient exploration. We also evaluate in a dense reward setting to show its potential for improving downstream tasks.

## 2. Problem Setup and Background

Consider a Markov Decision Process (MDP) represented as  $(\mathcal{S}, \mathcal{A}, P, r, \rho_0)$ , where  $\mathcal{S}$  is the state space,  $\mathcal{A}$  is the action space.  $P : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$  is the unknown dynam-

ics model, specifying the probability of transitioning to the next state  $s'$  from the current state  $s$  by taking the action  $a$ , as  $P(s'|s, a)$ .  $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$  is the reward function,  $\rho_0 : \mathcal{S} \rightarrow [0, 1]$  is the distribution of initial states. A policy is a function  $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ , which outputs a distribution over the action space for a given state  $s$ .

### 2.1. Exploration in Reinforcement Learning

In online decision-making problems, such as multi-arm bandits and reinforcement learning, a fundamental dilemma in an agent’s choice is exploitation versus exploration. Exploitation refers to making the best decision given current information, while exploration refers to gathering more information about the environment. In the standard reinforcement learning setting where the agent receives an external reward for each transition step, common recipes for exploration/exploitation trade-off include naive methods such as  $\epsilon$ -greedy (Sutton & Barto, 2018) and optimistic initialization (Lai & Robbins, 1985), posterior guided methods such as upper confidence bounds (Auer, 2002; Dani et al., 2008) and Thompson sampling (Thompson, 1933). We focus on the situation where external rewards are sparse or disregarded, here the above trade-off narrows down to the pure exploration problem of efficiently accumulating information about the environment. The common approach is to explore in a task-agnostic manner under some “intrinsic” reward. An exploration policy can then be trained with standard RL. Existing methods construct intrinsic rewards from visitation frequency of the state (Bellemare et al., 2016), prediction error of the dynamic model as “curiosity” (Pathak et al., 2017), diversity of visited states (Eysenbach et al., 2018), etc.

### 2.2. Dynamic Model Uncertainty as Intrinsic Reward

In order to model Bayesian uncertainty in online decision-making, two recent methods (Pathak et al., 2019; Shyam et al., 2019) train an ensemble of dynamic models and use the variation/information gain as an intrinsic reward for exploration. In this work, we follow the similar idea of exploiting the uncertainty in the dynamic model, but emphasize the implicit posterior modeling in contrast with directly training an ensemble of dynamic models.

Let  $f : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$  denote a model of the environment dynamics (represented by a neural network) we want to learn based on the agent experience  $\mathcal{D}$ . We design a generator module  $\mathcal{G}$  which takes a random draw from the standard normal distribution and outputs a sample vector of parameters  $\theta$  that determines  $f$  (denoted as  $f_\theta$ ). If samples from  $\mathcal{G}$  represent the posterior distribution  $p(f_\theta|D)$ , then given  $(s_t, a_t)$ , the uncertainty in the output of the dynamics model can be computed by the following variance among a set of samples  $\{\theta_i\}_{i=1}^m$  from  $\mathcal{G}$ , and used as an intrinsic




Figure 1. Architecture of the layer-wise generator of the dynamic model. A single shared noise sample  $z \in \mathbb{R}^d$  is drawn from a standard Gaussian with diagonal covariance, and input to layer-wise generators  $\{G_1, \dots, G_N\}$ . Each generator  $G_j$  outputs parameters  $\theta^j$  for the corresponding  $j$ -th layer of the neural network representing the dynamic model.

reward  $r^{in}$  for learning an exploration policy,

$$r_t^{in} = \frac{1}{m} \sum_{i=1}^m \left\| f_{\theta_i}(s_t, a_t) - \frac{1}{m} \sum_{\ell=1}^m f_{\theta_\ell}(s_t, a_t) \right\|^2. \quad (1)$$

When training the exploration policy, this intrinsic reward can be computed with rollouts in the environment, or simulated rollouts generated by the estimated dynamic model.

### 3. Posterior Approximation via Amortized SVGD

In this section, we introduce the core component of our exploration agent, the dynamic model generator  $\mathcal{G}$ . In the following subsections, we first introduce the design of this generator and then describe its training algorithm in detail. A summary of our algorithm is given in the last subsection.

#### 3.1. Implicit Posterior Generator

As shown in Fig. 1, the dynamic model is defined as an  $N$ -layer neural network function  $f_{\theta}(s, a)$ , with input (state, action) pair  $(s, a)$  and model parameters  $\theta = (\theta^1, \dots, \theta^N)$ , where  $\theta^j$  represents network parameters of the  $j$ -th layer. The generator module  $\mathcal{G}$  consists of exactly  $N$  layer-wise generators,  $\{G_1, \dots, G_N\}$ , where each  $G_j$  takes the random noise vector  $z \in \mathbb{R}^d$  as input, and outputs the corresponding parameter vector  $\theta^j = G_j(z; \eta^j)$ , where  $\eta^j$  are the parameters of  $G_j$ . Note that  $z$  is sampled from a  $d$ -dimensional standard normal distribution, and is shared across all generators to capture correlations between the generated parameters. As mentioned in Sec. 1, this framework has advantages in flexibility and efficiency, comparing with ensemble-based methods (Shyam et al., 2019; Pathak et al., 2019), since it maintains only parameters of the  $N$  generators, i.e.,  $\eta = (\eta^1, \dots, \eta^N)$ , and enables drawing an arbitrary number of sample networks to

approximate the posterior of the dynamic model.

#### 3.2. Training with Amortized SVGD

We now introduce the training algorithm of the generator module  $\mathcal{G}$ . Assuming that the true posterior of the dynamic model given the agent’s experience  $\mathcal{D}$  is  $p(f|\mathcal{D})$ , and the implicit distribution captured by  $\mathcal{G}$  is  $q(f_{\theta})$ , where  $f_{\theta}$  denotes the function values obtained by evaluating  $f_{\theta}$  on  $\mathcal{D}$ . We want  $q(f_{\theta}|\mathcal{D})$  to be as close as possible to  $p(f|\mathcal{D})$ , such closeness is commonly measured by the KL divergence  $\mathbf{D}_{\text{KL}}[q(f_{\theta}|\mathcal{D})||p(f|\mathcal{D})]$ . The traditional approach for finding  $q$  that minimizes  $\mathbf{D}_{\text{KL}}[q(f_{\theta}|\mathcal{D})||p(f|\mathcal{D})]$  is variational inference (VI), by maximizing the ELBO (Blei et al., 2017). But standard VI necessitates restricting the parametric form of the target posterior. Recently, a non-parametric VI framework, Stein Variational Gradient Descent (SVGD) (Liu & Wang, 2016), was proposed, that represents  $q$  with a set of particles rather than making any parametric assumptions, and approximates the functional gradient descent w.r.t.  $\mathbf{D}_{\text{KL}}[q(f_{\theta}|\mathcal{D})||p(f|\mathcal{D})]$  by iterative particle evolution. We apply SVGD to our sampled network functions, and follow the idea of amortized SVGD (Feng et al., 2017) to project the functional gradients to the parameter space of  $\eta$  by back-propagation through the generators.

Given a set of dynamic functions  $\{f_{\theta_i}\}_{i=1}^m$  sampled from  $\mathcal{G}$ , SVGD updates each function by

$$f_{\theta_i} \leftarrow f_{\theta_i} + \epsilon \phi^*(f_{\theta_i}), \quad i = 1, \dots, m,$$

where  $\epsilon$  is step size, and  $\phi^*$  is the function in the unit ball of a reproducing kernel Hilbert space (RKHS)  $\mathcal{H}$  that maximally decreases the KL divergence between the distribution  $q$  represented by  $\{f_{\theta_i}\}_{i=1}^m$  and the target posterior  $p$ .

Let  $q_{[\epsilon\phi]}$  refer to the distribution of updated particles. The optimal choice of  $\phi$  can be found by solving the following optimization problem:

$$\phi^* = \max_{\phi \in \mathcal{H}} \left\{ -\frac{d}{d\epsilon} \mathbf{D}_{\text{KL}}(q_{[\epsilon\phi]}||p), \quad s.t. \|\phi\|_{\mathcal{H}} \leq 1 \right\}. \quad (2)$$

It was shown in (Liu & Wang, 2016) that the KL-divergence can be expressed as a linear functional of  $\phi$ ,

$$-\frac{d}{d\epsilon} \mathbf{D}_{\text{KL}}(q_{[\epsilon\phi]}||p)|_{\epsilon=0} = \mathbb{E}_{f_{\theta} \sim q} [\mathcal{S}_p \phi(f_{\theta})] \quad (3)$$

where  $\mathcal{S}_p$  is the Stein operator (Stein et al., 2004):

$$\mathcal{S}_p \phi(f_{\theta}) = \nabla_{f_{\theta}} \log p(f_{\theta})^T \phi(f_{\theta}) + \nabla_{f_{\theta}}^T \phi(f_{\theta})$$

Hence, eq. (2) has a closed form solution,

$$\phi^*(f_\theta) = \mathbb{E}_{f_\theta \sim q} [\nabla_{f_\theta} \log p(f_\theta) k(f_\theta, f_{\theta_i}) + \nabla_{f_\theta} k(f_\theta, f_{\theta_i})], \quad (4)$$

where  $k(\cdot, \cdot)$  is the positive definite kernel associated with the RKHS. The log-likelihood term for  $f_\theta$  corresponds to the negation of the regression loss of future state prediction for all transitions in  $\mathcal{D}$ , i.e.,  $\log p(f_\theta) = -\sum_{(s,a,s') \in \mathcal{D}} L(f_\theta(s,a), s')$ . Given that each  $\theta_i$  is generated by  $\mathcal{G}(z; \eta)$ , the update rule for  $\eta$  can be obtained by the chain rule,

$$\eta \leftarrow \eta + \epsilon \sum_{i=1}^m \nabla_{\eta} \phi^*(f_{\theta_i})|_{\theta_i = \mathcal{G}(z_i; \eta)} \quad (5)$$

where  $\phi^*(\mathcal{G}(z_i; \eta))$  can be computed by (4) using empirical expectation from sampled batch  $\{\theta_i\}_{i=1}^m$ ,

$$\phi^*(f_{\theta_i}) = \frac{1}{m} \sum_{\ell=1}^m \left\{ - \left[ \sum_{(s,a,s') \in \mathcal{D}} \nabla_{f_{\theta_\ell}} L(f_{\theta_\ell}(s,a), s') \right] \cdot k(f_{\theta_\ell}(s,a), f_{\theta_i}(s,a)) + \nabla_{f_{\theta_\ell}} k(f_{\theta_\ell}(s,a), f_{\theta_i}(s,a)) \right\}, \quad (6)$$

where  $k(\cdot, \cdot)$  is the Gaussian kernel evaluated at function outputs, which is in the state space.

### 3.3. Summary of the Exploration Algorithm

To condense what we have proposed so far, we summarize in Algorithm 1 the procedure used to train the generator of dynamic models and the exploration policies.

---

#### Algorithm 1 Exploration with an Implicit Distribution

---

**Initialize** Generator  $\mathcal{G}_\eta$ , parameters  $T, m$

**Initialize** Policy  $\pi$ , Experience buffer  $\mathcal{D}$

**while** True **do**

**while** episode not done: **do**

$f_\Theta \leftarrow \mathcal{G}(z; \eta), z \sim \mathcal{N}(0, I^d)$

$\eta \leftarrow$  evaluate (5), (6) on  $\mathcal{D}$

$\tilde{\mathcal{D}} \sim \text{MDP}(f_\Theta)$

$\mathcal{D}_\pi \leftarrow \mathcal{D} \cup \tilde{\mathcal{D}}$ ,

$R_\pi \leftarrow r^{in}(f_\theta, s, a | (s, a) \sim \mathcal{D}_\pi)$  by (1)

$\pi \leftarrow$  update policy on  $(\mathcal{D}_\pi, R_\pi)$

$\mathcal{D}_T \leftarrow$  rollout  $\pi$  for  $T$  steps

$\mathcal{D} \leftarrow \mathcal{D} \cup \mathcal{D}_T$

**end**

**end**

---

Our algorithm starts with a buffer  $\mathcal{D}$  of random transitions and explores for some fixed number of episodes. For each episode, our algorithm samples a set of dynamic models

$f_\Theta = \{f_{\theta_i}\}$  from the generator  $\mathcal{G}$ , and updates the generator parameters  $\eta$  using amortized SVGD (5) and (6). For the policy update, the intrinsic reward (1) is evaluated on the actual experience  $\mathcal{D}$  and the simulated experience  $\tilde{\mathcal{D}}$  generated by  $f_{\theta_i}$ . The exploration policy is then updated using a model-free RL algorithm on the collected experience  $\mathcal{D}_\pi$  and intrinsic rewards  $R_\pi$ . The updated exploration policy is then used to rollout in the environment for  $T$  steps so that new transitions are collected and added to the buffer  $\mathcal{D}$ . The process is repeated until the end of the episode.

## 4. Related Work

Efficient exploration remains a major challenge in deep reinforcement learning (Fortunato et al., 2017; Burda et al., 2018b; Eysenbach et al., 2018; Burda et al., 2018a), and there is no consensus on the *correct* way to explore an environment. One practical guiding principle for efficient exploration is the reduction of the agent’s epistemic uncertainty of the environment (Chaloner & Verdinelli, 1995; Osband et al., 2017). Osband et al. (2016) uses a bootstrap ensemble of DQNs, where the predictions of the ensemble are used as an estimate of the agent’s uncertainty over the value function. Osband et al. (2018) proposed to augment the predictions of a DQN agent by adding the contribution from a prior to the value estimate. In contrast to our method, these approaches seek to estimate the uncertainty in the *value function*, while we focus on exploration with intrinsic reward by estimating the uncertainty of the *dynamic model*. Fortunato et al. (2017) add parameterized noise to the agent’s weights, to induce state-dependant exploration beyond  $\epsilon$ -greedy or entropy bonus.

Methods for constructing intrinsic rewards for exploration have become the subject of increased study. One well-known approach is to use the prediction error of an inverse dynamics model as an intrinsic reward (Pathak et al., 2017; Schmidhuber, 1991). Schmidhuber (1991) and Sun et al. (2011) proposed using the learning progress of the agent as an intrinsic reward. Count based methods (Bellemare et al., 2016; Ostrovski et al., 2017) give a reward proportional to the visitation count of a state. Houthoofd et al. (2016) formulate exploration as a variational inference problem, and use Bayesian neural networks (BNN) to maintain the agent’s belief over the transition dynamics. The BNN predictions are used to estimate a form of Bayesian information gain called compression improvement. The variational approach is also explored in Mohamed & Rezende (2015); Gregor et al. (2016); Salge et al. (2014), who proposed using intrinsic rewards based on a variational lower bound on empowerment: the mutual information between an action and the induced next state. This reward is used to learn a set of discriminative low-level skills. The most closely-related work to ours are two recent methods (Pathak et al., 2019;

Shyam et al., 2019) that compute intrinsic rewards from an ensemble of dynamic models. Disagreement among the ensemble members in next-state predictions is computed as an intrinsic reward. Shyam et al. (2019) also uses *active exploration* (Schmidhuber, 2003; Chua et al., 2018), in which the agent is trained in a surrogate MDP, to maximize intrinsic reward before acting in the real environment. Our method follows the similar idea of exploiting the uncertainty in the dynamic model, but instead suggests an implicit generative modeling of the posterior of the dynamic function, which enables a more flexible approximation of the posterior uncertainty with better sample efficiency.

There has been a wealth of research on nonparametric particle-based variational inference methods (Liu & Wang, 2016; Dai et al., 2016; Ambrogioni et al., 2018), where particles are maintained to represent the variational distribution, and updated by solving an optimization problem within an RKHS. Notably, we use amortized SVGD (Feng et al., 2017) to optimize our generator for approximately sampling from the posterior of the dynamic model. In addition to amortized SVGD, other nonparametric methods for training implicit samplers with particle-based variational inference have been proposed, such as Li & Turner (2018).

## 5. Experiments

In this section we conduct experiments to compare our approach to the existing state-of-the-art in efficient exploration with intrinsic rewards to illustrate the following:

- An agent with an implicit posterior over dynamic models explores more effectively and efficiently than agents using a single model or a static ensemble.
- Agents seeking external reward find better policies when initialized from powerful exploration policies. Our ablation studies shows that the better the exploration policy as an initialization, the better the downstream task policy can learn.

To evaluate the proposed method in terms of exploration efficiency, we first consider exploration tasks agnostic of any external reward. In this setting, the agent explores the environment irrespective of any downstream task. Then, to further investigate the potential of our exploration policies, we consider transferring the learned exploration policy to downstream task policies where a dense external reward is provided. Note that both cases are important for understanding and applying exploration policies. In sparse reward settings, such as a maze, the reward could occur at any location, without informative hints accessible at other locations. Therefore an effective agent must be able to efficiently explore the entire state space in order to consistently find rewards under different task settings. In dense reward settings, the trade-off between exploration and exploitation

plays a central role in efficient policy learning. Our experiments show that even for a state-of-the-art model-free algorithm like Soft Actor-Critic (SAC) (Haarnoja et al., 2018), that already incorporates a strong exploration mechanism, spending some initial rollouts to learn a powerful exploration policy as an initialization of the task policy still considerably improves the learning efficiency.

### 5.1. Pure Exploration Results

For pure exploration, we consider three challenging continuous control tasks in which efficient exploration is known to be difficult. In each environment, the dynamics are non-linear and cannot be solved with tabular approaches. As explained in the beginning of Section 5, the agent does not receive any external reward and is motivated purely by the uncertainty in its belief of the environment.

**Experimental setup** To validate the effectiveness of our method, we compare with several state-of-the-art formulations of intrinsic reward. Specifically, we conduct experiments comparing the following methods:

- (*Ours*) The proposed intrinsic reward, using the estimated variance from an implicit distribution of the dynamic model.
- (*Random*) Random exploration as a naive baseline.
- (*ICM*) Error between predicted next state and observed next state (Pathak et al., 2017).
- (*Disagreement*) Variance of predictions from an ensemble of dynamic models (Pathak et al., 2019).
- (*MAX*) Jensen-Renyi information gain of the dynamic function (Shyam et al., 2019).

### Implementation details

Since our goal is to compare the performance across different intrinsic rewards, we fix the model architecture, training pipeline, and hyper-parameters across all methods,<sup>1</sup> shared hyper-parameters follow the *MAX* default settings. For the purpose of computing the information gain, dynamic models for *MAX* predict both mean and variance of the next state, while for other methods, dynamic models predict only the mean. Since our method trains a generator of dynamic models instead of a fixed-size ensemble, we fix the number of models we sample from the generator at  $m = 32$ , which equals the ensemble size for *MAX*, and *Disagreement*. For all experiments except for the Chain environment, we use SAC v1 (Haarnoja et al., 2018) as the model-free RL algorithm used to train the exploration policies.

<sup>1</sup>We use the codebase of *MAX* as a basis and implement *Ours*, *ICM*, and *Disagreement* intrinsic rewards under the same framework. The full *Disagreement* method includes an additional differentiable reward function that we compare with separately in the supplementary material.




Figure 2. The NChain environment.

### 5.1.1. TOY TASK: NCHAIN

As a sanity check, we first follow *MAX* (Shyam et al., 2019) by evaluating our method on a stochastic version of the toy environment NChain. As shown in Fig. 2, the chain is a finite sequence of  $N$  states. Each episode starts from state 1 and lasts for  $N + 9$  steps. For each step, the agent can move forward to the next state in the chain or backward to the previous state. Attempting to move off the edge of the chain results in the agent staying still. Reward is only afforded to the agent at the edge states: 0.01 for reaching state 0, and 1.0 for reaching state  $N - 1$ . In addition, there is uncertainty built into the environment: each state is designated as a *flip-state* with probability 0.5. When acting from a flip-state, the agent’s actions are reversed, i.e., moving forward will result in movement backward, and vice-versa. Given the (initially) random dynamics and a sufficiently long chain, we expect an agent using an  $\epsilon$ -greedy exploration strategy to exploit only the small reward of state 0. In contrast, agents with exploration policies which actively reduce uncertainty can efficiently discover all states in the chain. Fig. 3 shows that our agent navigates the chain in less than 15 episodes, while the  $\epsilon$ -greedy agent (double DQN) does not make meaningful progress. We also evaluate each of the methods enumerated in section 5.1.




Figure 3. Results on the 40-link chain environment. Each line is the mean of three runs, with the shaded regions corresponding to  $\pm 1$  standard deviation. Our method and *MAX* actively reduce uncertainty in the chain, and are able to quickly explore to the end of the chain.  $\epsilon$ -greedy *DDQN* fails to explore more than 40% of the chain. Both *ICM* and *Disagreement* perform better than *DDQN* but explore less efficiently compared to *MAX* and our method

We find that actively reducing uncertainty is critical to exploring the chain. We believe that because *ICM* explores using the prediction error of the dynamic model, a chain initialized with simple dynamics (few flip states) may lead to poor exploration. Though *Disagreement* uses a similar intrinsic reward as *Ours*, we suspect the use of a static en-

semble leads to a lack of predictive diversity, as the ensemble can easily overfit to the dynamics of the chain, limiting exploration. *MAX* may avoid overfitting to the chain due to using stochastic neural networks. Our method however, directly promotes model diversity using amortized SVGD, and uses the uncertainty in our dynamic model to explore new states. We provide additional details of the NChain experiments in the supplementary material.

### 5.1.2. ACROBOT CONTROL

The first continuous control environment that we evaluate is a modified version of the Acrobot. As shown in figure 4, the Acrobot environment begins with a hanging down pendulum which consists of two links connected by an actuated joint. Normally, a discrete action  $a \in \{-1, 0, 1\}$  either applies a unit force on the joint in the left or right direction ( $a = \pm 1$ ), or not ( $a = 0$ ). We modify the environment such that a continuous action  $a \in [-1, 1]$  applies a force  $F = |a|$  in the corresponding direction.

To focus on efficient exploration, we test the ability of each exploration method to sweep the entire lower hemisphere: positioning the acrobot completely horizontal towards both (left and right) directions. Given this is a relatively simple task and can be solved by random exploration, as shown in Figure 4, all four intrinsic reward methods solve it within just hundreds of steps and our method is the most efficient one. The takeaway here is that in relatively simple environments where there might be little room for improvement over state-of-the-art, our method still achieves a better performance due to its flexibility and efficiency in approximating the model posterior. As we will see in subsequent experiments, this observation scales well with the increasing difficulty of the environments.

### 5.1.3. ANT MAZE NAVIGATION

Next, we evaluate on the Ant Maze environment. In the Ant control task, the agent provides torques to each of the 8 joints of the ant. The provided observation contains the pose of the torso as well as the angles and velocities of each joint. For the purpose of exploration, we place the Ant in a U-shaped maze (shown in figure 6(a)), where the goal is to reach the end of the maze, discovering all the states. The agent’s performance is measured by the percentage of the maze explored during evaluation. Figure 5(a) shows the result of each method over 5 seeds. Our agent consistently navigates to the end of the maze faster than the other competing methods. While *MAX* (Shyam et al., 2019) also navigates the maze, the implicit uncertainty modeling scheme in our method allows our agent to better estimate the state novelty, which leads to a considerably faster exploration. To see that our agent fully explores the maze, and does not only trace out a single trajectory, we include state visita-





Figure 4. Performance of each method on the Acrobot environment (average of five seeds), with error bars representing  $\pm 1$  standard deviation. The length of each horizontal bar indicates the number of environment steps each agent/method takes to swing the acrobot to fully horizontal on both (left and right) directions.

tion diagrams in figures 6(b)-6(e). We see that the agent explores many paths through the maze, and has not left any large portion of the maze unexplored.


To provide a more intuitive understanding of the effect of an intrinsic reward and how it might correlate to the performance, we also plot in Figure 5(b) the intrinsic reward observed by our agent at each exploration step, compared with that observed by the *MAX* agent. For fair comparison we plot the intrinsic reward from eq.(1) for both methods. We can see that after step 2K, predictions from the *MAX* ensemble start to become increasingly similar, leading to a decline in intrinsic reward (Fig. 5(b)) as well as a slow-down in exploration speed (Fig. 5(a)). We hypothesize this is because in a regular ensemble, all members are updating their gradients on the same experiences without an explicit term to match the real posterior, leading to all agents eventually converging to the same representation. In contrast, our intrinsic reward keeps increasing around step 2K and remains high as we continue to quickly explore new states in the maze, only starting to decline once we have solved the maze at approximately step 5,500.

#### 5.1.4. ROBOTIC MANIPULATION

The final task is an exploration task in a robotic manipulation environment, HandManipulateBlock. As shown in Figure 7(a), a robotic hand is given a palm-sized block for



(a) Ant Navigation Task Results



(b) Ant Intrinsic Rewards

Figure 5. Figure (a) shows the performance of each method with mean and  $\pm 1$  standard deviation (shaded region) over five seeds.  $x$ -axis is the number of steps the ant has moved,  $y$ -axis is the percentage of the U-shaped maze that has been explored. Figure (b) shows the proposed intrinsic reward magnitude for each step in the environment, calculated for both our method and *MAX*.

manipulation. The agent has actuation control of the 20 joints that make up the hand, and its exploration performance is measured by the percentage of possible rotations of the cube that the agent performs. This is different from the original goal of this environment since we want to evaluate task-agnostic exploration rather than goal-based policies. In particular, the state of the cube is represented by Cartesian coordinates along with a quaternion to represent the rotation. We transform the quaternion to Euler angles and discretize the resulting state space by 45 degree intervals. The agent is evaluated based on how many of the 512 total states are visited.





Figure 6. Figure (a) displays U-shaped ant maze. Figures (b-e) show the behavior of the agent at different stages of training, over 5 seeds. Points are color-coded with blue points occurring at the beginning of the episode, and red points at the end.



(a) Robotic Hand



(b) Manipulation Task Results

Figure 7. (a) The Robotic Hand task in motion. (b) Performance of each method with mean and  $\pm 1$  standard deviation (shaded region) over five seeds.  $x$ -axis is the number of manipulation steps,  $y$ -axis is the number of rotation states of the block that has been explored. Our method (red) explores clearly faster than all other methods.

This task is far more challenging than previous tasks, having a larger state space and action space. Additionally, states are more difficult to reach than the Ant Maze environment: requiring manipulation of 20 joints instead of 8. In order to explore in this environment, an agent must also learn how to rotate the block without dropping it. Figure 7(b) shows the performance of each method over 5 seeds. This environment proved very challenging for all methods: none succeeded in exploring more than half of the state space. Still, our method performs the best by a clear margin.

## 5.2. Policy Transfer Experiments

So far, we have demonstrated that the proposed implicit generative modeling of the posterior over dynamic models leads to more effective and efficient pure exploration policies. While the efficiency of pure exploration is important under sparse reward settings, a natural follow-up question is whether a strong pure exploration policy would also be beneficial for downstream tasks where dense rewards are available. We give an answer to this question by performing the following experiments in the widely-used HalfCheetah environment.


We first train a task-agnostic exploration policy following Alg. 1 for 10K environment steps. The trained policy is then used to warm up the (downstream) task policy for an additional 10K environment steps. This warm-up stage is followed by standard training of the task policy using external rewards. Warm-up periods are often used as an initial exploration stage, to collect enough data for meaningful off-policy updates. The baseline SAC is SAC v1 (Haarnoja et al., 2018), where the warm-up stage consists of taking uniformly random actions for the first 10K steps, before performing any parameter updates. Given that our

trained exploration policies explore much more efficiently than acting randomly, we can examine if warming up SAC with our exploration policy offers a benefit over the standard uniform warm-up strategy.

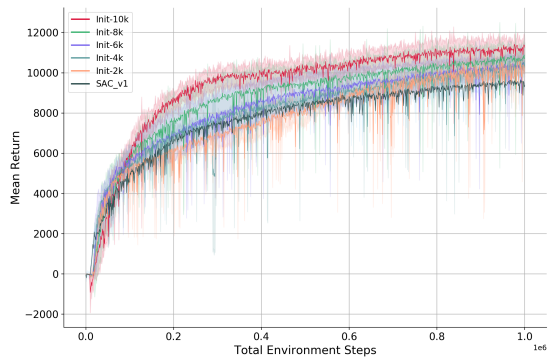
In particular, we first train a pure exploration policy for 10K steps, for each method on HalfCheetah. We then freeze the parameters of the pure exploration policy, and use them to initialize a new agent in a HalfCheetah environment where the external reward is known. We then warm up the agent, taking actions and collecting data according to the newly initialized policy. After 10K steps of warm-up we begin training as normal, with respect to the external reward. We evaluate this procedure that we call “policy transfer”, by comparing the performance of SAC at 1M steps, after the task policy has been warmed up using exploration policies trained by *MAX*, *ICM*, *Disagreement*, and *Ours* respectively. We also include SAC v1 (with a uniform warm-up strategy) as a baseline. For the training of the task policy, we follow the recommended settings for HalfCheetah given in the original SAC v1 method. In the supplementary material, we detail our hyper-parameter choices, as well as show that our choice of hyper-parameters does not unfairly favor our method.

Figure 8(a) shows the performance of all compared methods on HalfCheetah. We can see that the comparatively small number of initial steps spent on pure exploration pays off when the agent switches to the downstream task. Even though SAC is widely regarded as a strong baseline with a maximum entropy-based exploration mechanism, all intrinsic reward methods are able to improve the baseline more or less, by introducing a pure exploration stage before standard training of SAC. We also observe that the stronger the pure exploration policy is, the more it can improve the training efficiency of the downstream task. Task policies initialized with our exploration policy (*Ours*) still perform





(a) Policy Transfer Performance with Warm-up



(b) Policy Transfer with Varying Exploration Time

Figure 8. Policy transfer results. Figure (a) shows results with policy warm-up on the HalfCheetah environment. We show results for the SAC baseline with uniform warm-up, and SAC with 10K-step pure exploration initialization and warm-up using different intrinsic reward methods: ICM, Disagreement, MAX, and our proposed method (Ours) respectively. To make the comparison fair, agents initialized with pure exploration policies perform 10k less steps during task training than the baseline SAC. Figure (b) shows the performance of downstream SAC policy training warmed up with our proposed exploration policy under different numbers of pure exploration steps. Init- $N$ k refers to the SAC agent initialized from our exploration policy which has been trained for  $N$ -thousand exploration steps.

the best with a clear margin.

We also conduct an ablation study to better understand the relationship between the number of steps used to train the exploration policy, and the improvement it brings to downstream task training. We compare multiple variants of *Ours* in Figure 8(a), with different numbers of pure exploration steps: 2K, 4K, 6K, 8K, and 10K steps. As shown in Figure 8(b), with only 2K or more steps of initial pure exploration, our approach improves upon the SAC baseline in the downstream task. The longer our exploration policy is trained, the more beneficial it is to the training of the downstream task. We note that by using an exploration policy trained for just 4K steps, our agent performs strictly better on the downstream task than the SAC baseline.

Our next study shows that even in a setting without a warm-up stage, initializing the task policy with a pure exploration policy still benefit the downstream task learning. In figure 9 we show an evaluation of SAC on the HalfCheetah environment without the warm-up stage. We report performance of the SAC baseline, as well as SAC initialized with exploration policies trained by *MAX*, *ICM*, *Disagreement*, and *Ours*. We can see that without initial random exploration, the performance of SAC suffers dramatically. Policies initialized with pure exploration policies outperform the baseline following the same trend as in the setting with warm-up. *Ours* still performs the best with a clear margin.

## 6. Conclusion and Future Work

In this work, we introduced a new method for representing the agent’s uncertainty of the environment dynamics. Uti-




Figure 9. Policy transfer results without warm-up stage, on the HalfCheetah environment. We show our results for the SAC baseline without uniform warm-up, as well as an SAC initialized with exploration policies trained by *MAX*, *ICM*, *Disagreement*, and *Ours*, respectively

lizing amortized SVGD, we learned an approximate posterior over dynamics models. We use this approximate posterior to formulate an intrinsic reward based on the uncertainty estimated from samples of this distribution, enabling efficient exploration in difficult environments, Future work includes investigating the efficacy of learning an approximate posterior of the agent’s value or policy model, as well as more efficient sampling techniques to reduce the computational cost inherent to many model-based algorithms. We would also investigate principled methods of combining intrinsic and external rewards, and how different exploration policies influence the downstream task performance.

## Acknowledgements

Neale Ratzlaff and Li Fuxin were partially supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. N66001-17-12-4030, HR001120C0011 and HR001120C0022. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of DARPA.

## References

- Luca Ambrogioni, Umut Guclu, Yagmur Gucluturk, and Marcel van Gerven. Wasserstein variational gradient descent: From semi-discrete optimal transport to ensemble variational inference. *arXiv preprint arXiv:1811.02827*, 2018.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying count-based exploration and intrinsic motivation. In *Advances in Neural Information Processing Systems*, pp. 1471–1479, 2016.
- David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877, 2017.
- Yuri Burda, Harri Edwards, Deepak Pathak, Amos Storkey, Trevor Darrell, and Alexei A Efros. Large-scale study of curiosity-driven learning. *arXiv preprint arXiv:1808.04355*, 2018a.
- Yuri Burda, Harrison Edwards, Amos Storkey, and Oleg Klimov. Exploration by random network distillation. *arXiv preprint arXiv:1810.12894*, 2018b.
- Kathryn Chaloner and Isabella Verdinelli. Bayesian experimental design: A review. *Statistical Science*, pp. 273–304, 1995.
- Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems*, pp. 4754–4765, 2018.
- Bo Dai, Niao He, Hanjun Dai, and Le Song. Provable bayesian inference via particle mirror descent. In *Artificial Intelligence and Statistics*, pp. 985–994, 2016.
- Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. 2008.
- Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you need: Learning skills without a reward function. *arXiv preprint arXiv:1802.06070*, 2018.
- Yihao Feng, Dilin Wang, and Qiang Liu. Learning to draw samples with amortized stein variational gradient descent. *arXiv preprint arXiv:1707.06626*, 2017.
- Meire Fortunato, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Ian Osband, Alex Graves, Vlad Mnih, Remi Munos, Demis Hassabis, Olivier Pietquin, et al. Noisy networks for exploration. *arXiv preprint arXiv:1706.10295*, 2017.
- Karol Gregor, Danilo Jimenez Rezende, and Daan Wierstra. Variational intrinsic control. *arXiv preprint arXiv:1611.07507*, 2016.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *arXiv preprint arXiv:1801.01290*, 2018.
- Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. In *International Conference on Machine Learning*, pp. 2555–2565, 2019.
- Rein Houthoofd, Xi Chen, Xi Chen, Yan Duan, John Schulman, Filip De Turck, and Pieter Abbeel. Vime: Variational information maximizing exploration. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 29*, pp. 1109–1117. Curran Associates, Inc., 2016.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- Yingzhen Li and Richard E. Turner. Gradient estimators for implicit models. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=SJi9WOeRb>.
- Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm. In *Advances in neural information processing systems*, pp. 2378–2386, 2016.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin

- Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.
- Shakir Mohamed and Danilo Jimenez Rezende. Variational information maximisation for intrinsically motivated reinforcement learning. In *Advances in neural information processing systems*, pp. 2125–2133, 2015.
- Ian Osband, Charles Blundell, Alexander Pritzel, and Benjamin Van Roy. Deep exploration via bootstrapped dqn. In *Advances in neural information processing systems*, pp. 4026–4034, 2016.
- Ian Osband, Benjamin Van Roy, Daniel Russo, and Zheng Wen. Deep exploration via randomized value functions. *arXiv preprint arXiv:1703.07608*, 2017.
- Ian Osband, John Aslanides, and Albin Cassirer. Randomized prior functions for deep reinforcement learning. In *Advances in Neural Information Processing Systems*, pp. 8617–8629, 2018.
- Georg Ostrovski, Marc G Bellemare, Aäron van den Oord, and Rémi Munos. Count-based exploration with neural density models. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2721–2730. JMLR. org, 2017.
- Deepak Pathak, Pulkit Agrawal, Alexei A Efros, and Trevor Darrell. Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 16–17, 2017.
- Deepak Pathak, Dhiraj Gandhi, and Abhinav Gupta. Self-supervised exploration via disagreement. In *International Conference on Machine Learning*, pp. 5062–5071, 2019.
- Neale Ratzlaff and Li Fuxin. Hypergan: A generative model for diverse, performant neural networks. *arXiv preprint arXiv:1901.11058*, 2019.
- Christoph Salge, Cornelius Glackin, and Daniel Polani. Empowerment—an introduction. In *Guided Self-Organization: Inception*, pp. 67–114. Springer, 2014.
- Juergen Schmidhuber. Exploring the predictable. In *Advances in evolutionary computing*, pp. 579–612. Springer, 2003.
- Jürgen Schmidhuber. Curious model-building control systems. In *Proc. international joint conference on neural networks*, pp. 1458–1463, 1991.
- Pranav Shyam, Wojciech Jaśkowski, and Faustino Gomez. Model-based active exploration. In *International Conference on Machine Learning*, pp. 5779–5788, 2019.
- David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484, 2016.
- Jasper Snoek, Yaniv Ovadia, Emily Fertig, Balaji Lakshminarayanan, Sebastian Nowozin, D Sculley, Joshua Dillon, Jie Ren, and Zachary Nado. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. In *Advances in Neural Information Processing Systems*, pp. 13969–13980, 2019.
- Charles Stein, Persi Diaconis, Susan Holmes, Gesine Reinert, et al. Use of exchangeable pairs in the analysis of simulations. In *Stein’s Method*, pp. 1–25. Institute of Mathematical Statistics, 2004.
- Yi Sun, Faustino Gomez, and Jürgen Schmidhuber. Planning to be surprised: Optimal bayesian exploration in dynamic environments. In *International Conference on Artificial General Intelligence*, pp. 41–51. Springer, 2011.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- William R Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.