
Flow-based Recurrent Belief State Learning for POMDPs

Xiaoyu Chen¹ Yao Mu² Ping Luo² Shengbo Eben Li³ Jianyu Chen^{1,4}

Abstract

Partially Observable Markov Decision Process (POMDP) provides a principled and generic framework to model real world sequential decision making processes but yet remains unsolved, especially for high dimensional continuous space and unknown models. The main challenge lies in how to accurately obtain the belief state, which is the probability distribution over the unobservable environment states given historical information. Accurately calculating this belief state is a precondition for obtaining an optimal policy of POMDPs. Recent advances in deep learning techniques show great potential to learn good belief states. However, existing methods can only learn approximated distribution with limited flexibility. In this paper, we introduce the **F**low-based **R**ecurrent **B**elief **S**tate model (FORBES), which incorporates normalizing flows into the variational inference to learn general continuous belief states for POMDPs. Furthermore, we show that the learned belief states can be plugged into downstream RL algorithms to improve performance. In experiments, we show that our methods successfully capture the complex belief states that enable multi-modal predictions as well as high quality reconstructions, and results on challenging visual-motor control tasks show that our method achieves superior performance and sample efficiency.

1. Introduction

Partially Observable Markov Decision Process (POMDP) (Åström, 1965) provides a principled and generic framework to model real world sequential decision making processes.

¹Institute for Interdisciplinary Information Sciences, Tsinghua University ²Department of Computer Science, The University of Hong Kong ³School of Vehicle and Mobility, Tsinghua University ⁴Shanghai Qizhi Institute. Correspondence to: Jianyu Chen <jianyuchen@tsinghua.edu.cn>.

Unlike Markov Decision Process (MDP), the observations of a POMDP are generally non-Markovian. Therefore, to make optimal decisions, the agent needs to consider all historical information, which is usually intractable. One effective solution is to obtain the belief state. The belief state is defined as the probability distribution of the unobservable environment state conditioned on the past observations and actions (Kaelbling et al., 1998). Such belief state accurately summarizes the history. Traditional methods of calculating belief states (Smallwood & Sondik, 1973; Sondik, 1971; Kaelbling et al., 1998) assume finite discrete space with a known model. In many real world problems, however, the underlying model remains unknown, and the state space is large and even continuous. To track belief states in POMDPs with continuous state and action spaces, another line of works (Thrun, 1999; Silver & Veness, 2010) uses Monte Carlo algorithms like particle filters to estimate belief states. Nishiyama et al. (2012) proposes to solve the POMDP based on models defined in appropriate reproducing kernel Hilbert spaces (RKHSs). However, this requires access to samples from hidden states during training.

With the recent advances of deep learning technologies, recent works mainly focus on POMDPs with unknown models and continuous state spaces. To capture belief states, a branch of works including Hausknecht & Stone (2015); Gregor et al. (2019b) uses vector-based representations, namely scalars, to represent belief states. However, vector-based belief states may fall short in making predictions for multiple future trajectories (as discussed in Appendix D). Another line of works proposes to learn belief states by approximating belief state distributions. The current state-of-the-art performance on many visual-motor control tasks is also achieved in this manner by sequentially maximizing the observation probability at each timestep using the variational inference (Hafner et al., 2019a; Zhu et al., 2020; Okada et al., 2020; Ma et al., 2020a; Mu et al., 2021). They approximate the belief states with distributions like diagonal Gaussians (Krishnan et al., 2015; Han et al., 2019; Gregor et al., 2019a; Hafner et al., 2019b;a; Lee et al., 2020), Gaussian mixture (Tschitschek et al., 2018), categorical distribution (Hafner et al., 2021), or particle filters (Ma et al., 2020b; Igl et al., 2018). However, they still cannot capture general belief states due to the intractability of complex distributions in high-dimensional continuous space. They either suffer from

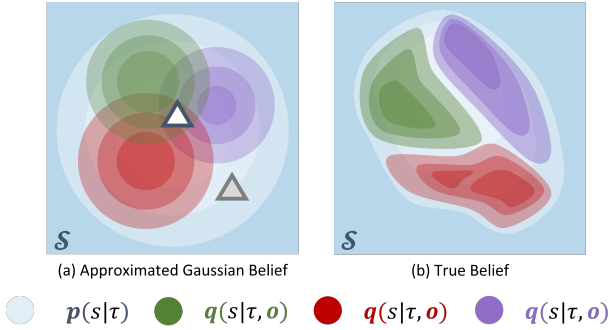


Figure 1: Difference between (a) spherical Gaussian belief states and (b) true belief states (better viewed in color). The spherical Gaussian belief states (a) approximate the true belief states in (b) using Gaussian assumptions, which may result in intersection (the white triangle) or vacancy (the gray triangle) points in the state space.

the curse of dimensionality, or instead make some assumptions and learn only the approximated distributions. Those approximation imposes strong restrictions and is problematic.

Taking the Gaussian belief assumption as an example: as shown in Figure 1, the blue area denotes the unobservable state space of the POMDP. Given the past information τ , the agent maintains a prior distribution of the state s , denoted as $p(s|\tau)$ (the distribution in white). Each colored distribution corresponds to the belief state after receiving a different new observation o , denoted as the posterior distribution $q(s|\tau, o)$. Consider an example of the true beliefs as shown in Figure 1(b), with their Gaussian approximations shown in Figure 1(a). The approximation error of Gaussian distributions will easily result in problems of intersecting belief which leads to a mixed-up state (e.g., the white triangle), and empty belief, which leads to a meaningless state (e.g., the grey triangle). This also explains the poor reconstruction problems in interactive environments observed by Okada & Taniguchi (2021). Furthermore, as mentioned in Hafner et al. (2021), the Gaussian approximation of belief states also makes it difficult to predict multi-modal future behaviours. Therefore, it is preferable to relax the Gaussian assumptions and use a more flexible family of distributions to learn accurate belief states as shown in Figure 1(b). For a more detailed discussion of the related works, please check Section 5 and Appendix D.

In this paper, we propose a new method called **FI**ow-based **RE**cursive **BE**lief State model (FORBES) that is able to learn general continuous belief states for POMDPs. FORBES incorporates Normalizing Flows (Tabak & Turner, 2013; Rezende & Mohamed, 2015; Dinh et al., 2017) into the variational inference step to construct flexible belief states.

In experiments, we show that FORBES allows the agent to maintain flexible belief states, which result in multi-modal and precise predictions as well as higher quality reconstructions. We also demonstrate the results combining FORBES with downstream RL algorithms on challenging visual-motor control tasks (DeepMind Control Suite, (Tassa et al., 2018)). The results show the efficacy of FORBES in terms of improving both performance and sample efficiency.

Our contributions can be summarized as follows:

- We propose FORBES, the first flow-based belief state learning algorithm that is capable of learning general continuous belief states for POMDPs.
- We incorporate FORBES into a POMDP RL framework for visual-motor control tasks that can fully exploit the benefits brought by FORBES.
- Empirically, we show that FORBES allows the agent to learn flexible belief states that enable multi-modal predictions as well as high quality reconstructions and help improve both performance and sample efficiency for challenging visual-motor control tasks.

2. Preliminaries

2.1. Partially Observable Markov Decision Process




Figure 2: The PGM of POMDP. The grey circle represents the unobservable hidden states s , while the observations o , rewards r are observable, and the actions a are determined by the agent.

Formally, a Partially Observable Markov Decision Process (POMDP) is a 7-tuple $(\mathcal{S}, \mathcal{A}, T, R, \Omega, O, \gamma)$, where \mathcal{S} is a set of states, \mathcal{A} is a set of actions, T is a set of conditional transition probabilities between states, R is the reward function, Ω is a set of observations, O is a set of conditional observation probabilities, and γ is the discount factor.

At each timestep $t - 1$, the state of the environment is $s_{t-1} \in \mathcal{S}$. The agent takes an action $a_{t-1} \in \mathcal{A}$, which causes the environment to transit to state s_t with probability $T(s_t | s_{t-1}, a_{t-1})$. The agent then receives an observation $o_t \in \Omega$ which depends on the new state of the environment s_t with probability $O(o_t | s_t)$. Finally, the agent receives a reward r_{t-1} equal to $R(s_{t-1})$. The agent’s goal is

to maximize the the expected sum of discounted rewards $\mathbb{E}[\sum_{t=0}^{\infty} \gamma^t r_t]$. Such a POMDP model can also be described as a probabilistic graphical model (PGM) as shown in Figure 2. After having taken action a_{t-1} and observing o_t , an agent needs to update its belief state, which is defined as the probability distribution of the environment state conditioned on all historical information:

$$b(s_t) = p(s_t | \tau_t, o_t) \quad (1)$$

where $\tau_t = \{o_1, a_1, \dots, o_{t-1}, a_{t-1}\}$.

2.2. Normalizing Flow

Instead of using the Gaussian family to approximate the prior and posterior belief distributions, we believe it is more desirable to use a family of distributions that is highly flexible and preferably flexible enough to describe all possible true belief states. Therefore, we use Normalizing Flows (Tabak & Turner, 2013; Rezende & Mohamed, 2015) to parameterize those distributions.

Rather than directly parameterizing statistics of the distribution itself, Normalizing Flows model the transformations, or the ‘‘flow’’ progress, needed to derive such a distribution. More specifically, it describes a sequence of invertible mappings that gradually transform a relatively simple probability density to a more flexible and complex one.

Let $f_\theta : \mathbb{R}^D \rightarrow \mathbb{R}^D$ to be an invertible and differentiable mapping in state space parameterized by θ . Given a random variable $\mathbf{x} \in \mathbb{R}^D$ with probability distribution $p(\mathbf{x})$, we can derive the probability of the transformed random variable $\mathbf{z} = f_\theta(\mathbf{x})$ by applying the change of variable formula:

$$p(\mathbf{z}) = p(\mathbf{x}) \left| \det \frac{\partial f_\theta^{-1}}{\partial \mathbf{z}} \right| \quad (2)$$

$$\log p(\mathbf{z}) = \log p(\mathbf{x}) - \log \left| \det \frac{\partial f_\theta}{\partial \mathbf{z}} \right| \quad (3)$$

To construct a highly flexible family of distributions, we can propagate the random variable at beginning \mathbf{z}_0 through a sequence of K mappings and get $\mathbf{z}_K = f_{\theta_K} \circ f_{\theta_{K-1}} \circ \dots \circ f_{\theta_1}(\mathbf{z}_0)$ with the probability

$$\log p_K(\mathbf{z}_K) = \log p(\mathbf{z}_0) - \sum_{k=1}^K \log \left| \det \frac{\partial f_{\theta_k}}{\partial \mathbf{z}_{k-1}} \right| \quad (4)$$

Given a relatively simple distribution of \mathbf{z}_0 , say, Gaussian distribution, by iteratively applying the transformations, the flow is capable of representing a highly complex distribution with the probability that remains tractable. The parameters $\theta_1, \dots, \theta_K$ determine the transformations of the flow.

An effective transformation that is widely accepted is affine coupling layer (Dinh et al., 2017; Kingma & Dhariwal, 2018; Kingma et al., 2017). Given the input $\mathbf{x} \in \mathbb{R}^D$, let s and

t stand for scale and translation functions which are usually parameterized by neural networks, where $s, t : \mathbb{R}^k \rightarrow \mathbb{R}^{D-k}$, $k < D$. The output, \mathbf{y} , can be viewed as a concatenation of its first k dimensions $\mathbf{y}_{1:k}$ and the remaining part $\mathbf{y}_{k+1:D}$:

$$\begin{aligned} \mathbf{y}_{1:k} &= \mathbf{x}_{1:k}, \\ \mathbf{y}_{k+1:D} &= \mathbf{x}_{k+1:D} \odot \exp(s(\mathbf{x}_{1:k})) + t(\mathbf{x}_{1:k}) \end{aligned} \quad (5)$$

where \odot denotes the element-wise product (see details about affine coupling layer in Appendix A).

3. Flow-based Recurrent Belief State Learning

3.1. Flow-based Recurrent Belief State model

We propose the FLOW-based RECURRENT BELIEF STATE model (FORBES) which learns general continuous belief states via normalizing flows under the variational inference framework. Specifically, the FORBES model consists of components needed to construct the PGM of POMDP as shown in Figure 2:

$$\begin{aligned} \text{State transition model : } & p(s_t | s_{t-1}, a_{t-1}) \\ \text{Observation model : } & p(o_t | s_t) \\ \text{Reward model : } & p(r_t | s_t) \end{aligned} \quad (6)$$

In addition, we have a belief inference model $q(s_t | \tau_t, o_t)$ to approximate the true posterior distribution $p(s_t | \tau_t, o_t)$ as defined in Equation 1, where $\tau_t = \{o_1, a_1, \dots, o_{t-1}, a_{t-1}\}$ is the past information. The above components of FORBES can be optimized jointly by maximizing the Evidence Lower Bound (ELBO) (Jordan et al., 1999) or more generally the variational information bottleneck (Tishby et al., 2000; Alemi et al., 2016):

$$\begin{aligned} & \log p(o_{1:T} | a_{1:T}) \\ & \geq \sum_{t=1}^T \left(\mathbb{E}_{q(s_t | o_{\leq t}, a_{< t})} [\log p(o_t | s_t) + \log p(r_t | s_t)] \right. \\ & \quad \left. - \mathbb{E} [D_{\text{KL}}(q(s_t | \tau_t, o_t) || p(s_t | s_{t-1}, a_{t-1}))] \right) \doteq \mathcal{J}_{\text{Model}} \end{aligned} \quad (7)$$

Detailed derivations can be found in Appendix J. In practice, the state transition model, observation model, reward model, and belief inference model can be represented by stochastic deep neural networks parameterized by ψ :

$$p_\psi(s_t | s_{t-1}, a_{t-1}), p_\psi(o_t | s_t), p_\psi(r_t | s_t), q_\psi(s_t | \tau_t, o_t)$$

where their outputs usually follow simple distributions such as diagonal Gaussians. The parameterized belief inference model $q_\psi(s_t | \tau_t, o_t)$ acts as an encoder that encodes the historical information using a combination of convolutional neural networks and recurrent neural networks.

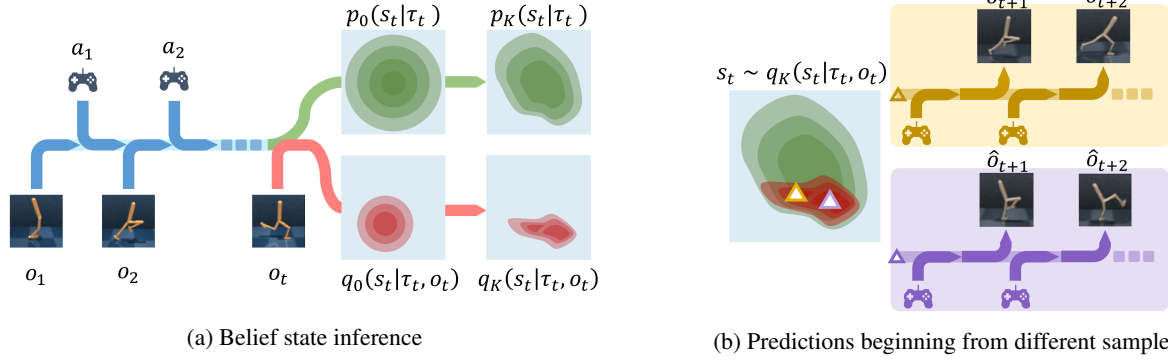


Figure 3: The algorithm framework of FORBES. Figure 3a shows how to calculate prior and posterior belief distribution given previous information. The blue arrows bring in historical observations and actions, and the green path shows the evolution of prior belief distribution. The red path takes an additional o_t and shows the evolution of posterior belief distribution. Figure 3b shows the predictions of future trajectories starting from different samples (yellow and purple triangles) given the future actions.

In FORBES we provide special treatments for the belief inference model and the state transition model to represent more complex and flexible posterior and prior distributions. As shown in Figure 3(a), the input images $o_{1:t}$ and actions $a_{1:t-1}$ are encoded with $q_\psi(s_t|\tau_t, o_t)$ (the blue and the red path). Then our final inferred belief is obtained by propagating $q_\psi(s_t|\tau_t, o_t)$ through a set of normalizing flow mappings denoted $f_{\theta_K} \circ \dots \circ f_{\theta_1}$ to get a representative posterior distribution $q_{\psi, \theta}(s_t|\tau_t, o_t)$. For convenience, we denote $q_0 = q_\psi$ and $q_K = q_{\psi, \theta}$. On the other hand, $o_{1:t-1}$ and $a_{1:t-2}$ are encoded with $q_\psi(s_{t-1}|\tau_{t-1}, o_{t-1})$ (the blue path), then the state transition model is used to obtain the prior guess of the state $p_\psi(s_t|\tau_t) = \mathbb{E}_{q_\psi(s_{t-1}|\tau_{t-1}, o_{t-1})} [p_\psi(s_t|s_{t-1}, a_{t-1})]$ (the green path). Then our final prior is obtained by propagating $p_\psi(s_t|\tau_t)$ through another set of normalizing flow mappings denoted $f_{\omega_K} \circ \dots \circ f_{\omega_1}$ to get a representative prior distribution $p_{\psi, \omega}(s_t|\tau_t)$. For convenience, we denote $p_0 = p_\psi$ and $p_K = p_{\psi, \omega}$. Then as shown in Figure 3(b), we can sample the initial state s_t (the yellow and purple triangles) from the belief states $q_K(s_t|\tau_t, o_t)$. For each sampled initial state, we can use the state transition model to predict the future states \hat{s}_{t+h} given the future actions $a_{t:t+h-1}$, and then use the observation model to reconstruct the observations \hat{o}_{t+h} , where h is the prediction horizon.

With the above settings, we can substitute the density probability inside the KL-divergence term in Equation 7 with Normalizing Flow:

$$\begin{aligned} \log q_K(s_t|\tau_t, o_t) &= \log q_0(s_t|\tau_t, o_t) - \sum_{k=1}^K \log \left| \det \frac{\partial f_{\theta_k}}{\partial s_{t, k-1}} \right| \\ \log p_K(s_t|\tau_t) &= \log p_0(s_t|\tau_t) - \sum_{k=1}^K \log \left| \det \frac{\partial f_{\omega_k}}{\partial s_{t, k-1}} \right| \end{aligned} \quad (8)$$

where $p_K(s_t|s_{t-1}, a_{t-1}) = p_K(s_t|\tau_t)$ given the sampled s_{t-1} from $q_K(s_{1:t}|\tau_t, o_t)$. $s_{t,k}$ is the state variable s_t transformed by k layers of normalizing flows, and $s_{t,0} = s_t$.

To further demonstrate the properties of FORBES, we provide the following theorems.

Theorem 3.1. *The approximation error of the log-likelihood when maximizing the $\mathcal{J}_{\text{Model}}$ (the derived ELBO) defined in Equation 7 is:*

$$\begin{aligned} &\log p(o_{1:T}, r_{1:T}|a_{1:T}) - \mathcal{J}_{\text{Model}} \\ &= \mathbb{E}_{q_K(s_{1:T}|o_{1:T}, a_{1:T-1})} \left[\sum_{t=1}^T D_{\text{KL}}(q(s_t|\tau_t, o_t) \| p(s_t|\tau_t, o_t)) \right] \end{aligned} \quad (9)$$

where $p(s_t|\tau_t, o_t)$ denotes the true belief states.

Detailed proofs can be found in Appendix.K. Theorem 3.1 suggests that, when the learning algorithm maximizes the $\mathcal{J}_{\text{Model}}$ (the derived ELBO), then the D_{KL} terms in the right-hand side are minimized, which indicate the KL-divergence between the learned belief states $q(s_t|\tau_t, o_t)$ and the true belief states $p(s_t|\tau_t, o_t)$. Clearly, if $p(s_t|\tau_t, o_t)$ is a complex distribution and $q(s_t|\tau_t, o_t)$ is chosen from a restricted distribution class such as diagonal Gaussian, then when the algorithm maximizes the $\mathcal{J}_{\text{Model}}$ (the derived ELBO), there will still be a potentially large KL-divergence between the learned and the true belief states.

Therefore, naturally there raises the problem that is normalizing flow a universal distributional approximator that is capable of accurately representing arbitrarily complex belief states, so the KL-divergence terms in the right-hand side of Equation (9) can be minimized to approach zero? The answer is yes for a wide range of normalizing flows. To be specific, Teshima et al. (2020) provides theoretical results for the family of the flow used in FORBES.

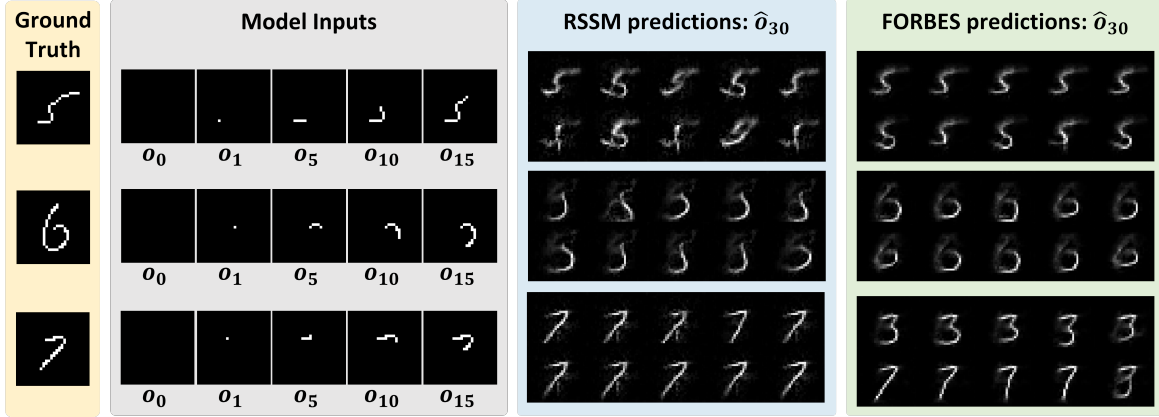


Figure 4: Predictions on sequential MNIST of two models. This is a digit writing task. The fully written digits are shown in the leftmost column. We use incomplete writing processes (the first 15 frames, partially shown in the grey column) as the inputs and let the models predict the complete digit (as shown in the blue/green columns). The results show that FORBES can make precise yet diverse predictions with less blur and no mode mixup.

Besides the aforementioned affine coupling flow, many works show the distributional universality of other flows (Kong & Chaudhuri, 2020; Huang et al., 2018). Ideally, the universal approximation property of the flow model $q_K(s_t | \tau_t, o_t)$ allows us to approximate the true posterior $p(s_t | \tau_t, o_t)$ with arbitrary accuracy. Thus, compared to previous methods, FORBES helps close the gap between the log-likelihood and the ELBO to obtain a more accurate belief state. Though we usually cannot achieve the ideal zero KL-divergence in practice, our method can get a smaller approximation error, equally a higher ELBO than previous works. We verify this statement in section 4.1.

3.2. POMDP RL framework based on FORBES

To show the advantage of the belief states inferred by the FORBES model compared to the existing belief inference method in visual-motor control tasks, we incorporate FORBES into a flow-based belief reinforcement learning algorithm for learning the optimal policy in POMDPs. Inspired by Hafner et al. (2019a), the algorithm follows an actor-critic framework but is slightly modified to exploit better the flexible nature of FORBES: The critic estimates the accumulated future rewards, and the actor chooses actions to maximize the estimated cumulated rewards. Instead of using only one sample, both the actor and critic operate on top of the samples of belief states learned by FORBES. They thus benefit from the accurate representations learned by the FORBES model. Note that this is an approximation of the true value on belief, which avoids the intractable integration through observation model.

The critic $v_\xi(s_\tau)$ aims to predict the discounted sum of future rewards that the actor can achieve given an initial state s_t , known as the state value $\mathbb{E}\left(\sum_{\tau=t}^{t+\infty} \gamma^{\tau-t} r_\tau\right)$, where

ξ denote the parameters of the critic network and H is the prediction horizon. We leverage temporal-difference to learn this value, where the critic is trained towards a value target that is constructed from the intermediate reward and the critic output for the next step’s state. In order to trade-off the bias and the variance of the state value estimation, we use the more general TD(λ) target (Sutton & Barto, 2018), which is a weighted average of n-step returns for different horizons and is defined as follows:

$$V_\tau^\lambda \doteq \hat{r}_\tau + \hat{\gamma}_\tau \begin{cases} (1 - \lambda)v_\xi(s_{\tau+1}) + \lambda V_{\tau+1}^\lambda & \text{if } \tau < t + H, \\ v_\xi(s_{t+H}) & \text{if } \tau = t + H. \end{cases} \quad (10)$$

To better utilize the flexibility belief states from FORBES, we run the sampling method multiple times to capture the diverse predictions. Specifically, we sample N states from the belief state given by FORBES and then rollout trajectories of future states and rewards using the state transition model and the reward model. Finally, we train the critic to regress the TD(λ) target return using a mean squared error loss:

$$\mathcal{J}_{\text{Critic}}(\xi) = \mathbb{E}\left[\sum_{i=1}^N \sum_{\tau=t}^{t+H} \frac{1}{2} (v_\xi(s_{i,\tau}) - \text{sg}(V_{i,\tau}^\lambda))^2\right]. \quad (11)$$

where $\text{sg}(\cdot)$ is the stop gradient operation. The actor $a_\tau \sim q_\phi(a_\tau | s_\tau)$ aims to output actions that maximize the prediction of long-term future rewards made by the critic and is trained directly by backpropagating the value gradients through the sequence of sampled states and actions, i.e., maximize:

$$\mathcal{J}_{\text{Actor}}(\phi) = \mathbb{E}\left(\sum_{i=1}^N \sum_{\tau=t}^{t+H} V_{i,\tau}^\lambda\right) \quad (12)$$

We jointly optimize the model loss $\mathcal{J}_{\text{Model}}$ with respect to the model parameters ψ , θ and ω , the critic loss $\mathcal{J}_{\text{Critic}}$

Algorithm 1 FORBES Algorithm

Input: buffer \mathcal{B} , imagination horizon H , interacting step T , batch size B , batch length L , number of trajectories N . Initialize buffer \mathcal{B} with S random seed episodes.

while not converged **do**

for $c = 1, \dots, C$ **do**

 Draw B data sequences $\{(o_t, a_t, r_t)\}_{t=k}^{k+L}$ from \mathcal{B}

 Infer belief state $q_K(s_t|s_{t-1}, a_{t-1}, o_t)$.

for $i = 1, \dots, N$ **do**

 Rollout imaginary trajectories $\{(s_{i,\tau}, a_{i,\tau})\}_{\tau=t}^{t+H}$ with belief transition model.

end for

 For each $s_{i,\tau}$, predict rewards $p_\psi(r_{i,\tau}|s_{i,\tau})$ and values $v_\phi(s_{i,\tau})$ *{Calculate returns}*

 Update $\theta, \omega, \xi, \phi, \psi$ using Equation (7), (8), (11), (12) and (13) *{Optimize parameters}*

end for

 Reset environment and get o_1 .

for $t = 1, \dots, T$ **do**

 Compute $s_t \sim q_K(s_t|s_{t-1}, a_{t-1}, o_t)$ from history.

 Compute $a_t \sim \pi(a_t|s_t)$ with action model.

 Add exploration noise to action.

 Execute a_t and get o_{t+1}, r_t .

end for

 Add experience to buffer $\mathcal{B} = \mathcal{B} \cup \{(o_t, a_t, r_t)_{t=1}^T\}$

end while

with respect to the critic parameters ξ and the actor \mathcal{J}_{Actor} loss with respect to the actor parameters ϕ using the Adam optimizer with different learning rates:

$$\min_{\psi, \xi, \phi, \theta, \omega} \alpha_0 \mathcal{J}_{Critic}(\xi) - \alpha_1 \mathcal{J}_{Actor}(\phi) - \alpha_2 \mathcal{J}_{Model}(\psi, \theta, \omega) \quad (13)$$

where $\alpha_0, \alpha_1, \alpha_2$ are coefficients for different components, and we summarize the whole framework of optimizing in Algorithm 1.

4. Experiments

Our experiments evaluate FORBES on two image-based tasks. We first demonstrate the belief learning capacity on a digit writing task in Section 4.1, and show that FORBES captures beliefs that allow for multi-modal yet precise long-term predictions as well as higher ELBO. For large-scale experiments, we test the proposed POMDP RL framework based on FORBES in Section 4.2. The results of multiple challenging visual-motor control tasks from DeepMind Control Suite (Tassa et al., 2018) show that FORBES outperforms baselines in terms of performance and sample efficiency. In Section 4.3, we further provide ablation studies of the multiple imagined trajectories technique used in our method.

4.1. Digit Writing Tasks

In this experiment, we validate the capacity of FORBES by modelling the partially observable sequence with visual inputs. We adopt the MNIST Sequence Dataset (D. De Jong, 2016) that consists of sequences of handwriting MNIST digit strokes. This problem can be viewed as a special case of POMDP, whose action space is \emptyset and rewards remain 0. Such a problem setting separates the belief learning and policy optimizing problem and allows us to concentrate on the former one in this section. We convert the digit stroke to a sequence of images of size 28×28 to simulate the writing process. At time step t , the agent can observe o_t that has already written t pixels, and we train the agent maximizing \mathcal{J}_{Model} in Equation 7 except for the reward reconstruction term. As shown in Figure 4, we randomly select three digits




Figure 5: ELBO on digit writing.

as examples (see Appendix L for more results) and show the inputs as well as the prediction outputs of our model and the RSSM (Hafner et al., 2019b) baseline, which is the previous state-of-the-art method for learning continuous belief states of POMDPs. The leftmost column is the ground truth of the fully written digits. During the testing, we feed the initial 15 frames $\{o_1, o_2, \dots, o_{15}\}$ to the model, and the columns in grey exhibit a part of the inputs. Then we sample several states from the inferred belief state and rollout via the learned state transition model (Equation (6)) for 15 steps and show the reconstruction results of the predictions. As shown in the blue and green columns on the right of Figure 4, though RSSM can also predict the future strokes in general, the reconstructions are relatively blurred and mix different digits up. It also fails to give diverse predictions. However, FORBES can make precise yet diverse predictions. Each prediction is clear and distinct from other digits. Given the beginning of the digit 7, FORBES successfully predicts both 7 and 3 since they have a similar beginning. The results can be partially explained via the mixed-up belief and the empty belief as shown in Figure 1, which support the claim that FORBES can better capture the complex belief states.

We also provide the quantitative results in Figure 5, which is the ELBO on test digits sequence set that is never seen during training. The results show that FORBES can achieve

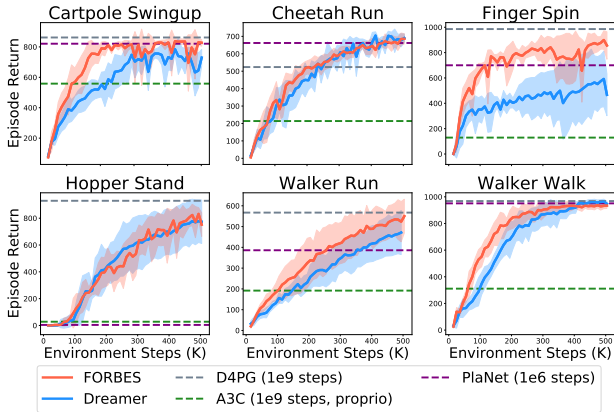


Figure 6: Performance on DeepMind Control Suite. The shaded areas show the standard deviation across 5 seeds. FORBES achieves better performance and sample efficiency in various challenging tasks.

a tighter ELBO, which verifies the theoretical results in 3.1. The details of the implementation can be found in Appendix B.

4.2. Visual-motor control tasks

We experimentally evaluate the performance of FORBES on Reinforcement Learning on a variety of visual-motor control tasks from the DeepMind Control Suite (Tassa et al., 2018), illustrated in Figure 6. Across all the tasks, the observations are $64 \times 64 \times 3$ images. These environments provide different challenges. The Cartpole-Swingup task requires a long planning horizon and memorizing the state of the cart when it is out of view; Finger-Spinning includes contact dynamics between the finger and the object; Cheetah-Run exhibits high-dimensional state and action spaces; the Walker-Walk and Walker-Run are challenging because the robot has to learn to first stand up and then walk; Hopper Stand is based on a single-legged robot, which is sensitive to the reaction force on the ground and thus needs more accurate control. As for baselines, we include the scores for A3C Mnih et al. (2016) with state inputs (1e9 steps), D4PG Barth-Maron et al. (2018) (1e9 steps), PlaNet (Hafner et al., 2019b) (1e6 steps) and Dreamer Hafner et al. (2019a) with pixel inputs. All the scores of baselines are aligned with the ones reported in Hafner et al. (2019a) (see details in Appendix C). We use $N = 4$ trajectories. The details of the implementations and hyperparameters can be found in Appendix B.

Our experiments empirically show that FORBES achieves superior performance and sample efficiency on challenging visual-motor control tasks. As illustrated in Figure 6, FORBES achieves higher scores than Dreamer (Hafner et al., 2019a) in most of the tasks and achieves better performance than PlaNet (Hafner et al., 2019b) within much fewer environment steps. See Appendix G for more results. We provide some insights into the results. As shown in Section

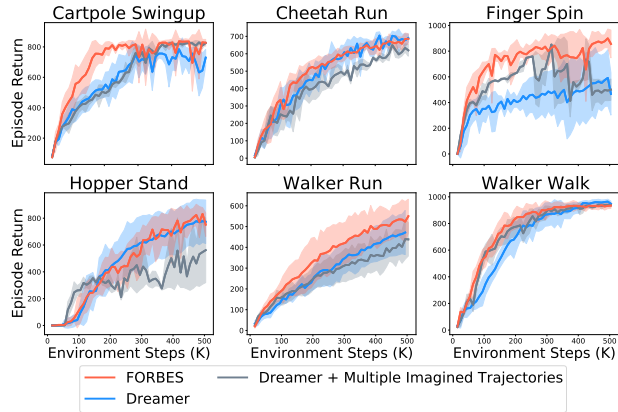


Figure 7: Comparison of the performance between FORBES and Dreamer with multiple imagined trajectories.

4.1, baselines with Gaussian assumptions may suffer from the mixed-up belief and empty belief issues, while FORBES can better capture the general belief states. Furthermore, multiple imagined trajectories can better utilize the diversity in the rollout. Therefore, the inner coherency within the model components allows the agent a better performance. See Appendix F for results on more environments. We further discuss the role of multiple imagined trajectories and other components in the next section.

4.3. Ablation Study

In order to verify that the outperformance of FORBES is not simply due to increasing the number of imagined trajectories, we conducted an ablation study in this section. We compare FORBES with the “Dreamer + multiple imagined trajectories” baseline by increasing the number of imagined trajectories in Dreamer to the same as in FORBES ($N = 4$). As shown in Figure 7, no consistent and obvious gain can be observed after increasing the number of trajectories to Dreamer. The agent gains slight improvements in two environments and suffers from slight performance loss on other tasks. This result indicates that increasing the number of imagined trajectories may only be effective when the agent can make diverse predictions as in FORBES. The Gaussian assumptions lead to the lack of trajectory diversity, so that increasing the number of imagined trajectories will not effectively help. Besides, Appendix E compare different N to illustrate the effect of multiple imagined trajectories. Appendix H add parameters to baselines to illustrate the performance gain is not due to more parameters.

4.4. Visual-motor control tasks with multimodal stochasticity

To better exhibit the advantages of FORBES on challenging tasks, we further make our environment more stochastic by extending DMC to multimodal DMC, aiming to solve the




Figure 8: Comparison of the performance between FORBES and Dreamer on multimodal DMC with different c .

multimodal beliefs problems.

With the observations left unchanged, we add a random binary state s_m . s_m is sampled from $\{+1, -1\}$ with equal probability at the beginning of each episode and then stays fixed. For each t in this episode, we add $s_m \times c$ to a_t , where c is a hyperparameter. This is equal to a piecewise-defined dynamic function. Since s_m is unobservable, this introduces clear, distinct multimodality to the belief. In this case, the agent needs to accurately capture the belief distribution in order to act optimally.

As illustrated in Figure 8, the results show that FORBES consistently gains significant improvements, which indicates that FORBES can capture the belief more accurately.

5. Related Work

POMDP: POMDP solving approaches can be divided into two categories based on whether their state, action and observation spaces are discrete or continuous. Discrete space POMDP solvers, in general, either approximate the value function using point-based methods (Kurniawati et al., 2008; Shani et al., 2013) or using Monte-Carlo sampling in the belief space (Thrun, 1999; Andrieu & Doucet, 2002; Silver & Veness, 2010; Kurniawati & Yadav, 2016) to make the POMDP problem tractable. Monte Carlo algorithms like particle filters make it possible to handle POMDPs with continuous state space by maintaining sets of samples drawn from the belief states. Other continuous space POMDP solvers often approximate the belief states as a distribution with few parameters (typically Gaussian) and solve the problem analytically either using gradients (Van Den Berg et al., 2012; Indelman et al., 2015) or using random sampling in the belief space (Agha-Mohammadi et al., 2014; Hollinger & Sukhatme, 2014). However, most of the classical POMDP methods mentioned above are based on an accurately known dynamic model, which is a restricted assumption in many real world tasks. More recently, Nishiyama et al. (2012) proposes to solve the POMDP based on models defined in appropriate RKHSs, which represent probability distributions as embeddings in RKHSs. However, the embeddings are learned from training samples, and therefore this method

requires access to samples from hidden states during training.

MBRL for visual-motor control: Recent researches in model-based reinforcement learning (MBRL) for visual-motor control provides promising methods to solve POMDPs with high-dimensional continuous space and unknown models since visual-motor control tasks can be naturally modelled as POMDP problems. Learning effective latent dynamics models to solve challenging visual-motor control problems is becoming feasible through advances in deep generative modeling and latent variable models (Krishnan et al., 2015; Karl et al., 2016; Doerr et al., 2018; Buesing et al., 2018; Ha & Schmidhuber, 2018; Han et al., 2019; Hafner et al., 2019b;a). Among which, the recurrent state-space model (RSSM) based methods (Hafner et al., 2019b;a) provide a principled way to learn continuous latent belief states for POMDPs by variational inference and learns behaviours based on the belief states using model-based reinforcement learning, which achieves high performance on visual-motor control tasks. However, they assume the belief states obey diagonal Gaussian distributions. Such assumptions impose strong restrictions to belief inference and lead to limitations in practice, including mode collapse, posterior collapse and object vanishing in reconstruction (Bowman et al., 2016; Salimans et al., 2015; Okada & Taniguchi, 2020). In addition to the diagonal Gaussian distributions, (Tschitschek et al., 2018) uses a Gaussian mixture to approximate the belief states. More recently, (Hafner et al., 2021) proposes to approximate the belief states by assuming a discrete latent space and results in superior performance. However, our algorithm makes no assumption and has the capability to approach arbitrary continuous distribution according to the theoretical analysis. Other works like (Hausknecht & Stone, 2015; Gregor et al., 2019b) use a vector-based representation of belief states. However, this deterministic representation prohibits the agent from consistently forecasting the future since the results of the reconstructed observation contain multimodality, and one can hardly keep the samples stay in the same mode across time. Please check Appendix D for more details. A few works propose particle filter based methods that use samples to approximate the belief states (Ma et al., 2020b; Igl et al., 2018). However, particle filters are reported to experience the curse of dimensionality (Daum & Huang, 2003; Bengtsson et al., 2008) and therefore suffer from insufficient sample efficiency and performance (Lee et al. (2020)). For a more detailed discussion of the related works, please refer to Appendix D.

Normalizing Flows: Normalizing Flows (NF) are a family of generative models which produce tractable distributions with analytical density. For a transformation $f : \mathbf{R}^D \rightarrow \mathbf{R}^D$, the computational time cost of the log determinant is $\mathcal{O}(D^3)$. Thus most previous works choose to make the

computation more tractable. (Rezende & Mohamed, 2015; van den Berg et al., 2019) propose to use restricted functional form of f . Another choice is to force the Jacobian of f to be lower triangular by using an autoregressive model (Kingma et al., 2016; Papamakarios et al., 2018). These models usually excel at density estimation, but the inverse computation can be time-consuming. Dinh et al. (2015; 2017); Kingma & Dhariwal (2018) propose the coupling method to make the Jacobian triangular and ensure the forward and inverse can be computed with a single pass. The applications of NF include image generation (Ho et al., 2019; Kingma & Dhariwal, 2018), video generation (Kumar et al., 2019) and reinforcement learning (Mazouze et al., 2020; Ward et al., 2019; Touati et al., 2020).

6. Conclusion

General continuous belief states inference is a crucial yet challenging problem in high-dimensional Partially Observable Markov Decision Process (POMDP) problems. In this paper, we propose the FLOW-based RECURRENT BELIEF STATE model (FORBES) that can learn general continuous belief states by incorporating normalizing flows into the variational inference framework and then effectively utilize the learned belief states in downstream RL tasks. We show that theoretically, our method can accurately learn the true belief states and we verify the effectiveness of our method in terms of both the quality of learned belief states and the final performance of our extended POMDP RL framework on two visual input environments. The digit writing tasks demonstrate that our method can learn general belief states that enable precise and multi-modal predictions and high-quality reconstructions. General belief inference plays a vital role in solving the POMDP, and our method paves a way towards it. In the future, we will explore further approaches to improve the accuracy of belief states inference and information seeking, such as combining contrastive learning and using advanced network architectures such as transformers to build normalizing flows.

7. Acknowledgment

This work is supported by the Ministry of Science and Technology of the People’s Republic of China, the 2030 Innovation Megaprojects “Program on New Generation Artificial Intelligence” (Grant No. 2021AAA0150000), NSF China with U20A20334 and 52072213.

References

Agha-Mohammadi, A.-A., Chakravorty, S., and Amato, N. M. Firm: Sampling-based feedback motion-planning under motion uncertainty and imperfect measurements. *The International Journal of Robotics Research*, 33(2): 268–304, 2014.

Alemi, A. A., Fischer, I., Dillon, J. V., and Murphy, K. Deep variational information bottleneck. *arXiv preprint arXiv:1612.00410*, 2016.

Andrieu, C. and Doucet, A. Particle filtering for partially observed gaussian state space models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4):827–836, 2002.

Barth-Maron, G., Hoffman, M. W., Budden, D., Dabney, W., Horgan, D., TB, D., Muldal, A., Heess, N., and Lillicrap, T. Distributed distributional deterministic policy gradients, 2018.

Bengtsson, T., Bickel, P., and Li, B. Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems. *Probability and Statistics: Essays in Honor of David A. Freedman*, pp. 316–334, 2008. doi: 10.1214/193940307000000518. URL <http://dx.doi.org/10.1214/193940307000000518>.

Bowman, S. R., Vilnis, L., Vinyals, O., Dai, A. M., Jozefowicz, R., and Bengio, S. Generating sentences from a continuous space. *CONLL*, 2016.

Buesing, L., Weber, T., Racaniere, S., Eslami, S., Rezende, D., Reichert, D. P., Viola, F., Besse, F., Gregor, K., Hassabis, D., et al. Learning and querying fast generative models for reinforcement learning. *arXiv preprint arXiv:1802.03006*, 2018.

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.

D. De Jong, E. The mnist sequence dataset. <https://edwin-de-jong.github.io/blog/mnist-sequence-data/>, 2016. Accessed: 2019-07-07.

Daum, F. and Huang, J. Curse of dimensionality and particle filters. In *2003 IEEE Aerospace Conference Proceedings (Cat. No. 03TH8652)*, volume 4, pp. 4_1979–4_1993. IEEE, 2003.

Dinh, L., Krueger, D., and Bengio, Y. Nice: Non-linear independent components estimation. In *ICLR*, 2015.

Dinh, L., Sohl-Dickstein, J., and Bengio, S. Density estimation using real nvp, 2017.

Doerr, A., Daniel, C., Schiegg, M., Duy, N.-T., Schaal, S., Toussaint, M., and Sebastian, T. Probabilistic recurrent state-space models. In *International Conference on Machine Learning*, pp. 1280–1289. PMLR, 2018.

- Gregor, K., Papamakarios, G., Besse, F., Buesing, L., and Weber, T. Temporal difference variational auto-encoder, 2019a.
- Gregor, K., Rezende, D. J., Besse, F., Wu, Y., Merzic, H., and Oord, A. v. d. Shaping belief states with generative environment models for rl. *arXiv preprint arXiv:1906.09237*, 2019b.
- Ha, D. and Schmidhuber, J. World models. *arXiv preprint arXiv:1803.10122*, 2018.
- Hafner, D., Lillicrap, T., Ba, J., and Norouzi, M. Dream to control: Learning behaviors by latent imagination. *arXiv preprint arXiv:1912.01603*, 2019a.
- Hafner, D., Lillicrap, T., Fischer, I., Villegas, R., Ha, D., Lee, H., and Davidson, J. Learning latent dynamics for planning from pixels, 2019b.
- Hafner, D., Lillicrap, T., Norouzi, M., and Ba, J. Mastering atari with discrete world models, 2021.
- Han, D., Doya, K., and Tani, J. Variational recurrent models for solving partially observable control tasks, 2019.
- Hausknecht, M. and Stone, P. Deep recurrent q-learning for partially observable mdps. *arXiv preprint arXiv:1507.06527*, 2015.
- Ho, J., Chen, X., Srinivas, A., Duan, Y., and Abbeel, P. Flow++: Improving flow-based generative models with variational dequantization and architecture design. In *International Conference on Machine Learning*, pp. 2722–2730. PMLR, 2019.
- Hollinger, G. A. and Sukhatme, G. S. Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9):1271–1287, 2014.
- Huang, C.-W., Krueger, D., Lacoste, A., and Courville, A. Neural autoregressive flows. In *ICML*, 2018.
- Igl, M., Zintgraf, L., Le, T. A., Wood, F., and Whiteson, S. Deep variational reinforcement learning for pomdps, 2018.
- Indelman, V., Carlone, L., and Dellaert, F. Planning in the continuous domain: A generalized belief space approach for autonomous navigation in unknown environments. *The International Journal of Robotics Research*, 34(7): 849–882, 2015.
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., and Saul, L. K. An introduction to variational methods for graphical models. *Machine learning*, 37(2):183–233, 1999.
- Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. Planning and acting in partially observable stochastic domains. *Artificial intelligence*, 1998.
- Karl, M., Soelch, M., Bayer, J., and Van der Smagt, P. Deep variational bayes filters: Unsupervised learning of state space models from raw data. *arXiv preprint arXiv:1605.06432*, 2016.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. and Dhariwal, P. Glow: Generative flow with invertible 1x1 convolutions. In *NeurIPS*, 2018.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improved variational inference with inverse autoregressive flow. In *NIPS*, 2016.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improving variational inference with inverse autoregressive flow, 2017.
- Kong, Z. and Chaudhuri, K. The expressive power of a class of normalizing flow models, 2020.
- Krishnan, R. G., Shalit, U., and Sontag, D. Deep kalman filters. *arXiv preprint arXiv:1511.05121*, 2015.
- Kumar, M., Babaeizadeh, M., Erhan, D., Finn, C., Levine, S., Dinh, L., and Kingma, D. Videoflow: A flow-based generative model for video. *arXiv preprint arXiv:1903.01434*, 2(5), 2019.
- Kurniawati, H. and Yadav, V. An online pomdp solver for uncertainty planning in dynamic environment. In *Robotics Research*, pp. 611–629. Springer, 2016.
- Kurniawati, H., Hsu, D., and Lee, W. S. Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces. In *Robotics: Science and systems*, volume 2008. Citeseer, 2008.
- Lee, A. X., Nagabandi, A., Abbeel, P., and Levine, S. Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model, 2020.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., and Wierstra, D. Continuous control with deep reinforcement learning. *arXiv preprint arXiv:1509.02971*, 2015.
- Ma, X., Chen, S., Hsu, D., and Lee, W. S. Contrastive variational reinforcement learning for complex observations, 2020a.
- Ma, X., Karkus, P., Hsu, D., and Lee, W. S. Particle filter recurrent neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 5101–5108, 2020b.

- Mazouze, B., Doan, T., Durand, A., Pineau, J., and Hjelm, R. D. Leveraging exploration in off-policy algorithms via normalizing flows. In *Conference on Robot Learning*, pp. 430–444. PMLR, 2020.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T. P., Harley, T., Silver, D., and Kavukcuoglu, K. Asynchronous methods for deep reinforcement learning, 2016.
- Mu, Y., Zhuang, Y., Wang, B., Zhu, G., Liu, W., Chen, J., Luo, P., Li, S. E., Zhang, C., and HAO, J. Model-based reinforcement learning via imagination with derived memory. In Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=jeATherHHGj>.
- Nishiyama, Y., Boularias, A., Gretton, A., and Fukumizu, K. Hilbert space embeddings of pomdps, 2012.
- Okada, M. and Taniguchi, T. Dreaming: Model-based reinforcement learning by latent imagination without reconstruction. *arXiv preprint arXiv:2007.14535*, 2020.
- Okada, M. and Taniguchi, T. Dreaming: Model-based reinforcement learning by latent imagination without reconstruction, 2021.
- Okada, M., Kosaka, N., and Taniguchi, T. Planet of the bayesians: Reconsidering and improving deep planning network by incorporating bayesian inference, 2020.
- Papamakarios, G., Pavlakou, T., and Murray, I. Masked autoregressive flow for density estimation, 2018.
- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows. In *ICML*, 2015.
- Salimans, T., Kingma, D., and Welling, M. Markov chain monte carlo and variational inference: Bridging the gap. In *International Conference on Machine Learning*, pp. 1218–1226. PMLR, 2015.
- Shani, G., Pineau, J., and Kaplow, R. A survey of point-based pomdp solvers. *Autonomous Agents and Multi-Agent Systems*, 27(1):1–51, 2013.
- Silver, D. and Veness, J. Monte-carlo planning in large pomdps. *Neural Information Processing Systems*, 2010.
- Smallwood, R. D. and Sondik, E. J. The optimal control of partially observable markov processes over a finite horizon. *Operations research*, 21(5):1071–1088, 1973.
- Sondik, E. J. *The optimal control of partially observable Markov processes*. Stanford University, 1971.
- Sutton, R. S. and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
- Tabak, E. G. and Turner, C. V. A family of nonparametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66(2):145–164, 2013.
- Tassa, Y., Doron, Y., Muldal, A., Erez, T., Li, Y., de Las Casas, D., Budden, D., Abdolmaleki, A., Merel, J., Lefrancq, A., Lillicrap, T., and Riedmiller, M. Deepmind control suite, 2018.
- Teshima, T., Ishikawa, I., Tojo, K., Oono, K., Ikeda, M., and Sugiyama, M. Coupling-based invertible neural networks are universal diffeomorphism approximators, 2020.
- Thrun, S. Monte carlo pomdps. In *NIPS*, volume 12, pp. 1064–1070, 1999.
- Tishby, N., Pereira, F. C., and Bialek, W. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.
- Touati, A., Satija, H., Romoff, J., Pineau, J., and Vincent, P. Randomized value functions via multiplicative normalizing flows. In *Uncertainty in Artificial Intelligence*, pp. 422–432. PMLR, 2020.
- Tschiatschek, S., Arulkumaran, K., Stühmer, J., and Hofmann, K. Variational inference for data-efficient model learning in pomdps. *arXiv preprint arXiv:1805.09281*, 2018.
- Van Den Berg, J., Patil, S., and Alterovitz, R. Motion planning under uncertainty using iterative local optimization in belief space. *The International Journal of Robotics Research*, 31(11):1263–1278, 2012.
- van den Berg, R., Hasenclever, L., Tomczak, J. M., and Welling, M. Sylvester normalizing flows for variational inference, 2019.
- Ward, P. N., Smofsky, A., and Bose, A. J. Improving exploration in soft-actor-critic with normalizing flows policies. *arXiv preprint arXiv:1906.02771*, 2019.
- Zhu, G., Zhang, M., Lee, H., and Zhang, C. Bridging imagination and reality for model-based deep reinforcement learning, 2020.
- Åström, K. Optimal control of markov processes with incomplete state information. *Journal of Mathematical Analysis and Applications*, 10(1):174–205, 1965. ISSN 0022-247X. doi: [https://doi.org/10.1016/0022-247X\(65\)90154-X](https://doi.org/10.1016/0022-247X(65)90154-X). URL <https://www.sciencedirect.com/science/article/pii/0022247X6590154X>.

A. Details of affine coupling layer for normalizing flow

In this section, we will introduce the details about the affine coupling layer (Dinh et al., 2017).

In the forward function, we split the input $\mathbf{x} \in \mathbb{R}^D$ into two parts according to the dimension: $\mathbf{x} = [\mathbf{x}_{1:k}, \mathbf{x}_{k+1:D}]$. Then, we let the first part $\mathbf{x}_{1:k}$ stay identical, so that the first k dimensions in the output $\mathbf{y} \in \mathbb{R}^D$ is $\mathbf{y}_{1:k} = \mathbf{x}_{1:k}$. After that, we use the identical part as the inputs to determine the transform parameters. In our case, we define two neural network $s, t : \mathbb{R}^k \rightarrow \mathbb{R}^{D-k}$, which stand for scale and translation functions. They receive $\mathbf{x}_{1:k}$ as inputs and output the affine parameters. As in (Dinh et al., 2017), the second part can be derived by:

$$\mathbf{y}_{k+1:D} = \mathbf{x}_{k+1:D} \odot \exp(s(\mathbf{x}_{1:k})) + t(\mathbf{x}_{1:k}) \quad (14)$$

Finally, the output, \mathbf{y} is the concatenation of the two parts: $\mathbf{y} = [\mathbf{y}_{1:k}, \mathbf{y}_{k+1:D}]$.

The affine coupling layer is an expressive transformation with easily-computed forward and reverse passes. The Jacobian of affine coupling layer is a triangular matrix, and its log determinant can also be efficiently computed.

B. Hyper Parameters and implementation details

Network Architecture We use the convolutional and deconvolutional networks that are similar to Dreamer (Hafner et al., 2019a), a GRU (Cho et al., 2014) with 200 units in the dynamics model, and implement all other functions as two fully connected layers of size 200 with ReLU activations. Base distributions in latent space are 30-dimensional diagonal Gaussians with predicted mean and standard deviation. As for the parameters network, we use a residual network composed of one fully connected layer, one residual block, and one fully connected layer. The residual network receives \mathbf{x}_a and c as input. The input is first concatenated with the context and passed into the network. The residual block passes the input through two fully connected layers and returns the sum of the input and the output. Finally the last layer outputs the parameters and we use 5 layers of affine coupling flows with a LU layer between them.

In our case, we use samples from the belief distribution as the inputs to the actor and value function as an approximation to the actor and value function with belief distribution as input. Calculating $V(b)$ needs to integrate through both the observation model and state transition model. Our approximation makes an assumption like in Qmdp, to avoid integrating through the observation model.

We use a GRU as the recurrent neural network to summary to temporal information. We assume an initial state s_0 to be a zero vector. After taking action a_t , we concatenate a_t with the previous state s_t and pass it through a small MLP to get $y_t = f(s_t, a_t)$, and use it as the input to the GRU: $h_{t+1}, z_{t+1} = GRU(h_t, y_t)$. We pass z_{t+1} through an MLP to get the base prior belief distribution p_0 (mean and variance) and then we sample from p_0 and pass it through a sequence of Normalizing Flow to get a sample from p_K . For the posterior distribution, we first use a CNN as encoder to encode the observation o_t into the feature x_t , and then concatenate z_{t+1} and x_t and pass them through an MLP to get the base posterior belief distribution q_0 and a sequence of Normalizing Flow. Similarly, we finally get a sample s_{t+1} from q_K .

Training Details We basically adopt the same data buffer updating strategy as in Dreamer (Hafner et al., 2019a). First, we use a small amount of S seed episodes ($S = 5$ in DMC experiments) with random actions to collect data. After that, we train the model for C update steps ($C = 100$ in DMC experiment) and conduct one additional episode to collect data with small Gaussian exploration noise added to the action. Algorithm 1 shows one update step in C update steps. After C update steps, we conduct one additional episode to collect data (this is not shown in Algorithm 1). When the agent interacts with the environment, we record the observations, actions, and rewards of the whole trajectory $((o_t, a_t, r_t)_{t=1}^T)$ and add it to data buffer \mathcal{B} .

Hyperparameters For DMControl tasks, we pre-process images by reducing the bit depth to 5 bits and draw batches of 50 sequences of length 50 to train the FORBES model, value model, and action model models using Adam (Kingma & Ba, 2014) with learning rates $\alpha_0 = 5 \times 10^{-4}$, $\alpha_1 = 8 \times 10^{-5}$, $\alpha_2 = 8 \times 10^{-5}$, respectively and scale down gradient norms that exceed 100. We clip the KL regularizers in \mathcal{J}_{Model} below 3.0 free nats as in Dreamer and PlaNet. The imagination horizon is $H = 15$ and the same trajectories are used to update both action and value models. We compute the TD- λ targets with $\gamma = 0.99$ and $\lambda = 0.95$. As for multiple imagined trajectories, we choose $N = 4$ across all environments.

For digit writing experiments in Section 4.1, we decrease the GRU hidden size to be 20, let the base distributions be a 2-dimensional diagonal Gaussian and only use 3 layers of affine coupling flows. For the image processing, we simply divide

the raw pixels by 255 and subtract 0.5 to make the inputs lie in $[-0.5, 0.5]$.

C. Extended information of Baselines

For model-free baselines, we compare with D4PG (Barth-Maron et al., 2018), a distributed extension of DDPG, and A3C (Mnih et al., 2016), the distributed actor-critic approach. D4PG is an improved variant of DDPG (Lillicrap et al., 2015) that uses distributed collection, distributional Q-learning, multi-step returns, and prioritized replay. We include the scores for D4PG with pixel inputs and A3C (Mnih et al., 2016) with vector-wise state inputs from DMCcontrol. For model-based baselines, we use PlaNet (Hafner et al., 2019b) and Dreamer (Hafner et al., 2019a), two state-of-the-art model-based RL. PlaNet (Hafner et al., 2019b) selects actions via online planning without an action model and drastically improves over D4PG and A3C in data efficiency. Dreamer (Hafner et al., 2019a) further improve the data efficiency by generating imaginary rollouts in the latent space.

D. Further Discussion on Related Works

This section further discusses the relationship between our work and some related works [Gregor et al. \(2019b\)](#); [Hafner et al. \(2019a; 2021\)](#).

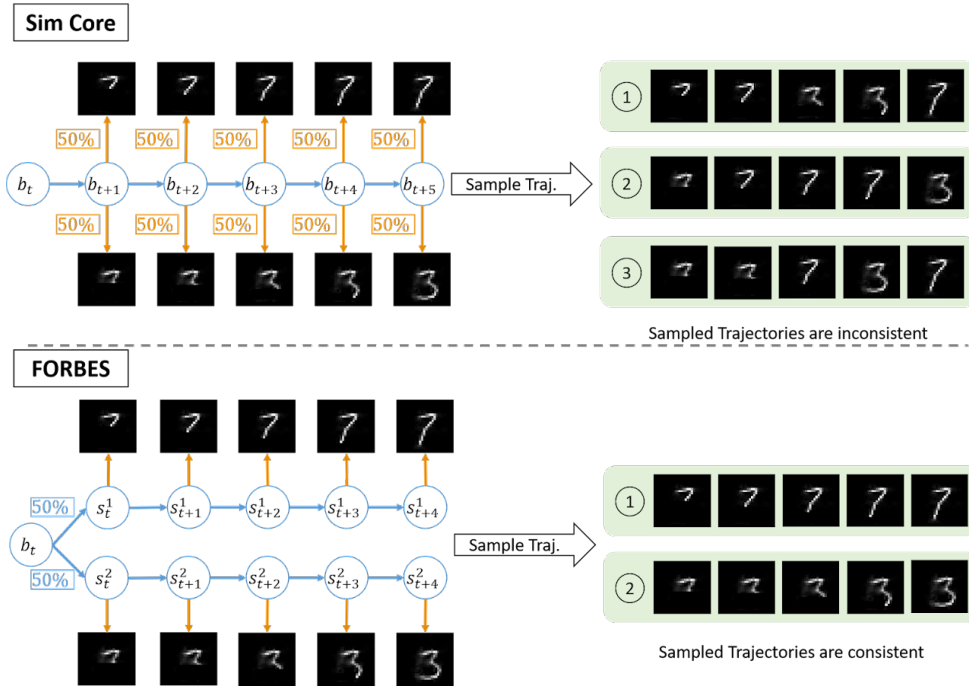


Figure 9: Comparison of FORBES and SimCore

First of all, [Gregor et al. \(2019b\)](#) proposes to use a flexible decoder to learn a compact vector representation of belief state and has promising results. Though it mentions using Normalizing Flow as a decoder, we believe it is orthogonal to our research. To summarize, we both aim to solve POMDP problems, but our method has little in common with ([Gregor et al., 2019b](#)), and our research directions are orthogonal: we use normalizing flows in different components for different purposes, which makes it unnecessary for citing and comparing with ([Gregor et al., 2019b](#)).

Specifically, our main contribution is to use Normalizing Flow to model accurate and flexible belief state distribution, and we prove its capability from theoretical and empirical perspectives. ([Gregor et al., 2019b](#)) did not model belief state distribution. Instead, they model the belief as a single state vector, and the belief transition is also deterministic. Their Normalizing Flow is only used on image reconstruction. They use expressive generative models, including normalizing flow, to reconstruct the images conditioned on the simulated future state. "Using a convolutional DRAW outperforms flows for learning a model of complex environment dynamics" is reported in the context of reconstruction.

What's more, directly modeling the observation model from the belief in the form of a deterministic vector may have the following deficiencies:

First of all, it seems that SimCore cannot make consistent trajectory predictions by directly modeling the belief observation model. As shown in Figure 9. For instance, in the sequential MNIST setting, suppose the SimCore([Gregor et al., 2019b](#)) is well trained, and we feed a beginning sequence which is the same as the last line in our Figure 4 into the SimCore. Assume it is equally possible to be '3' or '7', and the ConvDRAW will predict '3' and '7' each at 50% probability at every time step. However, we cannot consistently sample the same category ('3' or '7') in the same trajectory when we sample the future trajectory. However, since we explicitly model the state distribution for FORBES, we can first sample initial states and then rollout to sample multiple state trajectories, each covering a different category. This allows us to make diverse and consistent predictions.

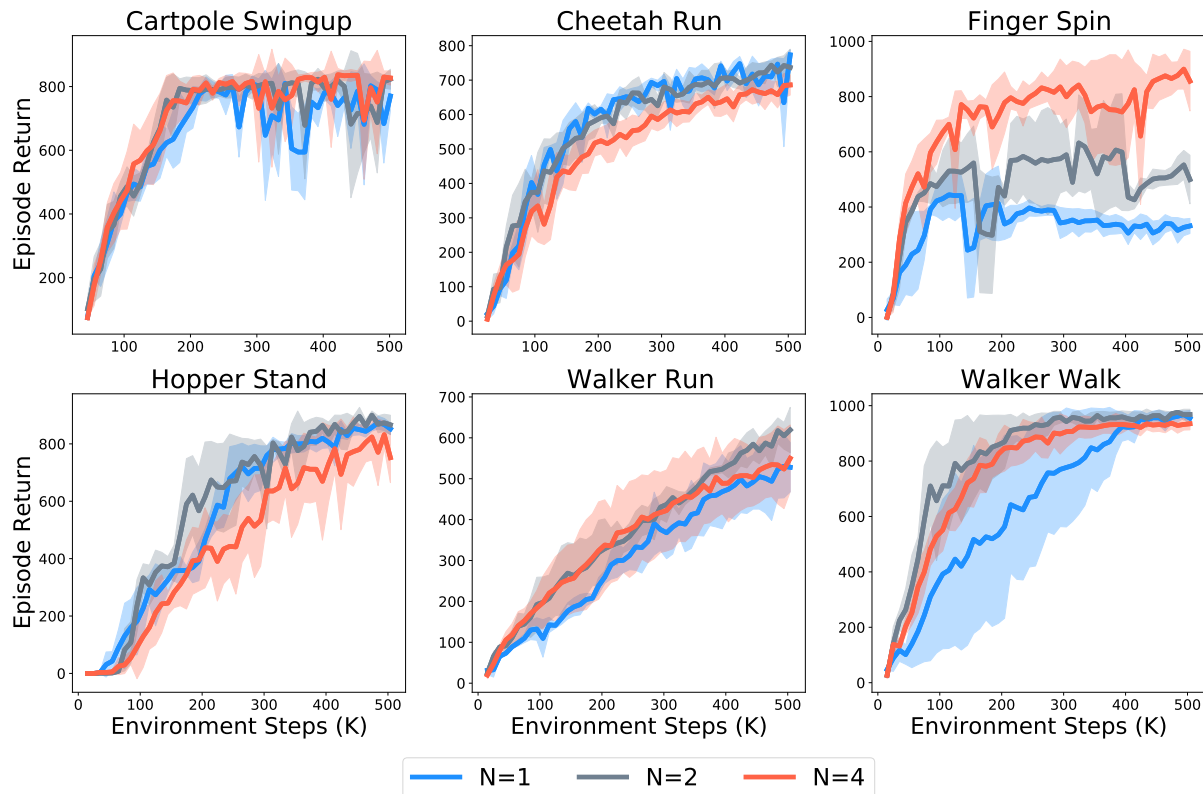
Secondly, accurately obtaining the belief state is the main challenge in solving the POMDP. Dreamer ([Hafner et al., 2021](#)) makes a strong isotropic Gaussian assumption to learn a continuous belief distribution, while Dreamer V2 ([Hafner et al.,](#)

2021) assumes discrete latent space. However, according to the theoretical analysis, our algorithm makes no assumption and can approach arbitrary continuous distribution. We believe that our methods can capture more subtle multimodal patterns without restricting the belief distribution to be discrete. This allows us to learn more general distribution (at least theoretically) and leaves great potential for future works.

To the best of our knowledge, we are the first to propose a normalizing flow based recurrent belief learning method to obtain the general continuous belief states in POMDP accurately. We provide theoretical analysis to illustrate that our algorithm has the potential of learning near perfectly accurate belief states. Through the sequential MNIST experiment, we empirically show the benefits of learning flexible belief distribution. Our method provides better reconstruction quality and can make multimodal future predictions. This flexible and accurate belief learning is essential for obtaining optimal solutions for POMDPs. As for the multiple imagined trajectories, we agree that the unimodal latent space leads to the lack of trajectory diversity, so that increasing the number of imagined trajectories will not effectively help. Our flexible belief distribution enables more accurate and multimodal future predictions by combining multiple imagined trajectories.

Therefore, we believe our proposed method is not merely a trivial combination of different components but a new framework for flexible and accurate belief distribution learning and POMDP RL with clear motivations and theoretical/empirical results.

E. An Ablation Study on the Number of Imagined Trajectories

Figure 10: An ablation study on the effect of different N on DMC environments.

To show the effect of N , we adjust the number of imagined trajectories on some DMC environments. We choose $N = 1, 2, 4$ and run 500K environment steps. We run $N = 1, 2$ with 3 different seeds, and $N = 4$ with 5 different seeds (we use the main DMC experiment results, where $N = 4$ here). The result shows that, in Finger Spin, the performance gain caused by multiple imagined trajectories is obvious. In finger spin, there are two objects and their interactions may result in complex locomotion patterns. When the environmental locomotion pattern itself is complex and flexible enough to incorporate diverse possibilities, then using FORBES allows the agent to make diverse predictions and using the multiple imagined trajectories technique will further exploit the advantages of FORBES. However, not all environments can show the advantages of multiple imaginations. In other environments, where there's only one agent and its behavior is relatively unimodal, a larger N does not effectively improve the performance, and different N results in similar performances.

F. Experiment Results on Other DMC Environments

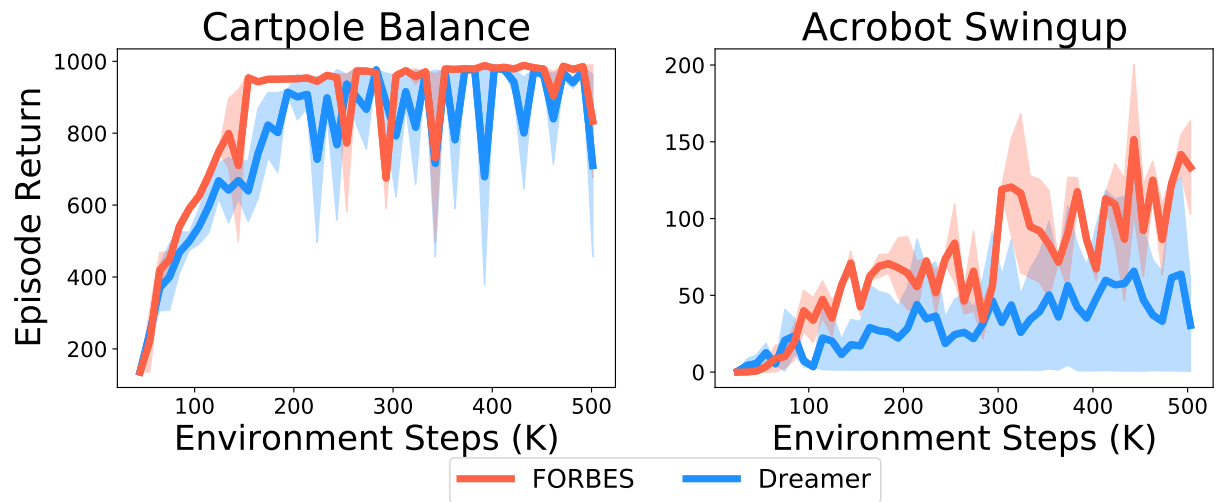


Figure 11: The training curve on Cartpole Balance and Acrobot Swingup.

We run our algorithm for on two more environments, namely Cartpole Balance and Acrobot Swingup and show the results in 11. FORBES achieves higher scores than Dreamer.

G. Extended Results on DMC

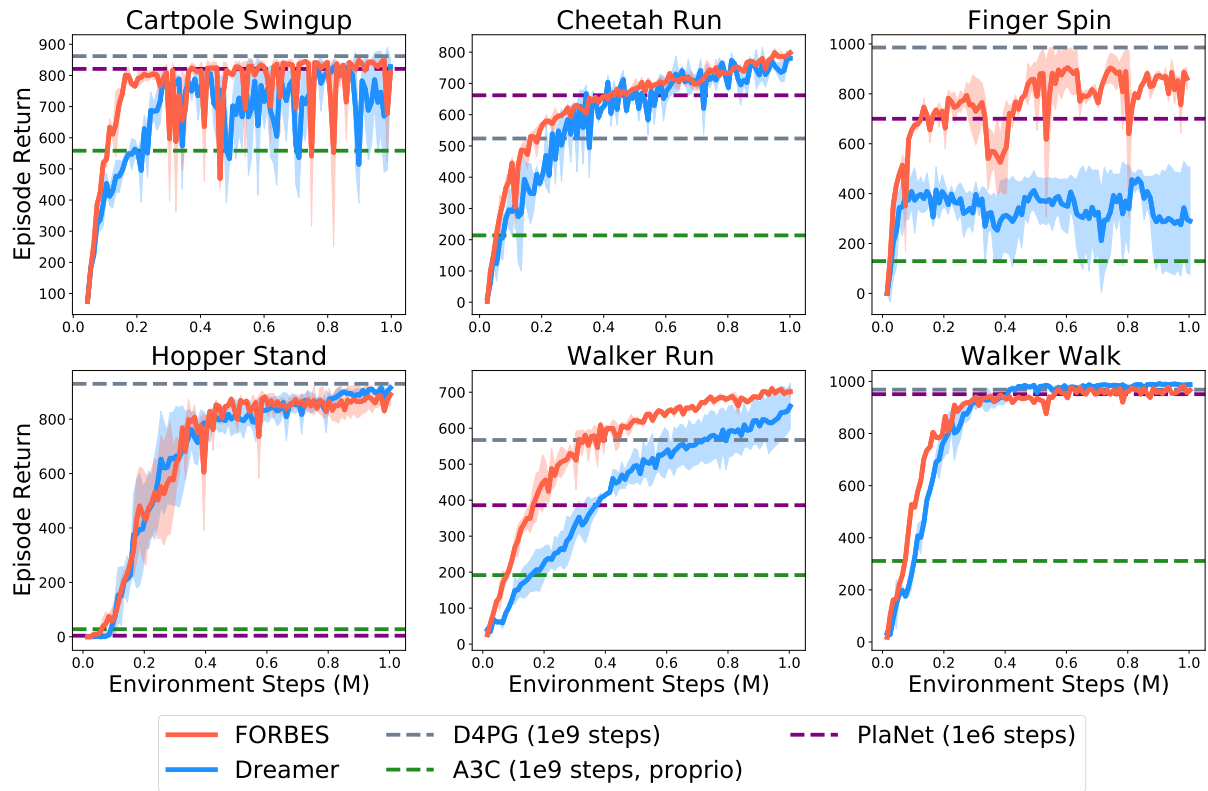


Figure 12: The training curve on DMC environment for 1M environment steps.

We run our algorithm for 1M environment steps and show the curve in Figure 12. We choose 1M environment steps because most of the curves have converged in most of the environments. FORBES achieves higher scores than Dreamer in most of the tasks.

H. An Ablation Study on the Model Parameters

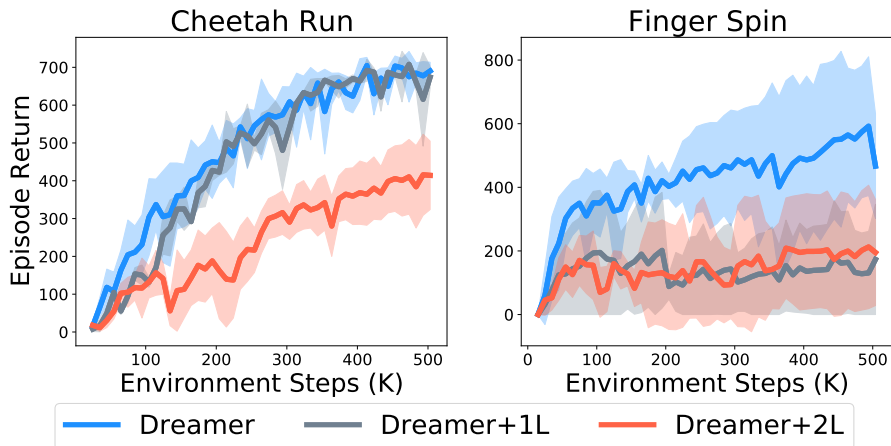


Figure 13: An ablation study on the effect of adding parameters to Dreamer on two DMC environments.

In this section, we show that having a flexible belief state distribution is the key to improving performance, rather than introducing more parameters. Having more parameters do not necessarily mean better performance. Increasing parameters may also make it difficult to converge and negatively affect the sample efficiency.

We add an ablation study that adds more parameters to Dreamer to test the effectiveness of having more parameters. We add 1, 2 hidden layer(s) to all the MLP in RSSM, and the result is shown in Figure 13. The results show that simply adding parameters cannot improve the performance.

I. Comparison of ELBO on FORBES and RSSM on DMC




Figure 14: The ELBO of FORBES and RSSM.

We provide the ELBO in DMC environments and FORBES in Figure 14, and FORBES has higher ELBO.

J. Evidence Lower Bound Derivations

The variational bound for latent dynamics models $p(o_{1:T}, s_{1:T} | a_{1:T}) = \prod_t p(s_t | s_{t-1}, a_{t-1})p(o_t | s_t)$ and a variational posterior $q(s_{1:T} | o_{1:T}, a_{1:T}) = \prod_t q(s_t | o_{\leq t}, a_{< t})$ follows from importance weighting and Jensen's inequality as shown,

$$\begin{aligned}
 \log p(o_{1:T} | a_{1:T}) &= \log \mathbb{E}_{p(s_{1:T}|a_{1:T})} \left[\prod_{t=1}^T p(o_t | s_t) \right] \\
 &= \log \mathbb{E}_{q(s_{1:T}|o_{1:T}, a_{1:T})} \left[\prod_{t=1}^T p(o_t | s_t) p(s_t | s_{t-1}, a_{t-1}) / q(s_t | o_{\leq t}, a_{< t}) \right] \\
 &\geq \mathbb{E}_{q(s_{1:T}|o_{1:T}, a_{1:T})} \left[\sum_{t=1}^T \log p(o_t | s_t) + \log p(s_t | s_{t-1}, a_{t-1}) - \log q(s_t | o_{\leq t}, a_{< t}) \right].
 \end{aligned} \tag{15}$$

We use the same factorization of $q(s_{1:T} | \tau_t, o_t)$ in ELBO derivations and algorithm design as in (Hafner et al., 2019b;a).

K. Proofs of Theorem

Theorem 1: The approximation error of the lower bound is

$$\log p(o_{1:T}, r_{1:T} | a_{1:T}) - \mathcal{J}_{\text{Model}} = \mathbb{E}_{q_K(s_{1:T} | \tau_T, o_T)} \left[\sum_{t=1}^T D_{\text{KL}}(q(s_t | \tau_t, o_t) \| p(s_t | \tau_t, o_t)) \right]$$

where $p(s_t | \tau_t, o_t)$ is the true posterior.

Proof:

$$\begin{aligned} & D_{\text{KL}}(q(s_t | \tau_t, o_t) \| p(s_t | s_{t-1}, a_{t-1}, o_t)) | a_{1:T} \\ &= \int q(s_t | \tau_t, o_t) \log \frac{q(s_t | \tau_t, o_t)}{p(s_t | s_{t-1}, a_{t-1}, o_t)} ds_t \\ &= \int q(s_t | \tau_t, o_t) \log \frac{q(s_t | \tau_t, o_t)}{\frac{p(s_t | s_{t-1}, a_{t-1})p(o_t | s_t)}{p(o_t | a_{1:T})}} ds_t \\ &= \int q(s_t | \tau_t, o_t) \log q(s_t | \tau_t, o_t) ds_t + \int q(s_t | \tau_t, o_t) \log p(o_t | a_{1:T}) ds_t \\ &\quad - \int q(s_t | \tau_t, o_t) \log [p(s_t | s_{t-1}, a_{t-1})p(o_t | s_t)] ds_t \\ &= \log p(o_t | a_{1:T}) + \int q(s_t | \tau_t, o_t) \log q(s_t | \tau_t, o_t) ds_t - \int q(s_t | \tau_t, o_t) \log [p(s_t | s_{t-1}, a_{t-1})p(o_t | s_t)] ds_t \\ &= \log p(o_t | a_{1:T}) + \int q(s_t | \tau_t, o_t) \log q(s_t | \tau_t, o_t) ds_t - \int q(s_t | \tau_t, o_t) \log p(s_t | s_{t-1}, a_{t-1}) ds_t \\ &\quad - \int q(s_t | \tau_t, o_t) \log p(o_t | s_t) ds_t \\ &= \log p(o_t | a_{1:T}) + \int q(s_t | \tau_t, o_t) \log \frac{q(s_t | \tau_t, o_t)}{p(s_t | s_{t-1}, a_{t-1})} ds_t - \int q(s_t | \tau_t, o_t) \log p(o_t | s_t) ds_t \\ &= \log p(o_t | a_{1:T}) + D_{\text{KL}}(q(s_t | \tau_t, o_t) \| p(s_t | s_{t-1}, a_{t-1}, o_t)) - \mathbb{E}_{q(s_{1:t} | \tau_t, o_t)}[\log p(o_t | s_t)] \end{aligned} \tag{16}$$

For a sequence from time 1 to T, we have

$$\begin{aligned} & \sum_t D_{\text{KL}}(q(s_t | \tau_t, o_t) \| p(s_t | s_{t-1}, a_{t-1}, o_t)) \\ &= \log p(o_{1:T} | a_{1:T}) - \mathbb{E}_{q(s_{1:t} | \tau_t, o_t)} \left[\sum_{t=1}^T (\log p(o_t | s_t) - D_{\text{KL}}(q(s_t | \tau_t, o_t) \| p(s_t | s_{t-1}, a_{t-1}, o_t))) \right] \end{aligned} \tag{17}$$

Then we can derive the Theorem 1 with (17):

$$\begin{aligned} & \log p(o_{1:T}, r_{1:T} | a_{1:T}) \\ &= \mathbb{E}_{q_K(s_{1:T} | \tau_T, o_T)} \left[\sum_t D_{\text{KL}}(q(s_t | \tau_t, o_t) \| p(s_t | s_{t-1}, a_{t-1}, o_t)) \right] \\ &\quad + \mathbb{E}_{q(s_{1:T} | \tau_T, o_T)} \left[\sum_{t=1}^T (\log p(o_t | s_t) + \log p(r_t | s_t) - D_{\text{KL}}(q(s_t | \tau_t, o_t) \| p(s_t | s_{t-1}, a_{t-1}, o_t))) \right] \\ &= \mathbb{E}_{q_K(s_{1:T} | \tau_T, o_T)} \left[\sum_t D_{\text{KL}}(q(s_t | \tau_t, o_t) \| p(s_t | s_{t-1}, a_{t-1}, o_t)) \right] + \mathcal{J}_{\text{Model}} \\ &= \mathbb{E}_{q_K(s_{1:T} | \tau_T, o_T)} \left[\sum_t D_{\text{KL}}(q(s_t | \tau_t, o_t) \| p(s_t | \tau_t, o_t)) \right] + \mathcal{J}_{\text{Model}} \end{aligned} \tag{18}$$

where $p(s_t | s_{t-1}, a_{t-1}, o_t) = p(s_t | \tau_t, o_t)$ given the sampled s_{t-1} from $q(s_{1:t} | \tau_t, o_t)$.

L. More Results on Digit Writing Experiments

In this section, we show more results of the predictions on the digit writing experiment in Figure 15.

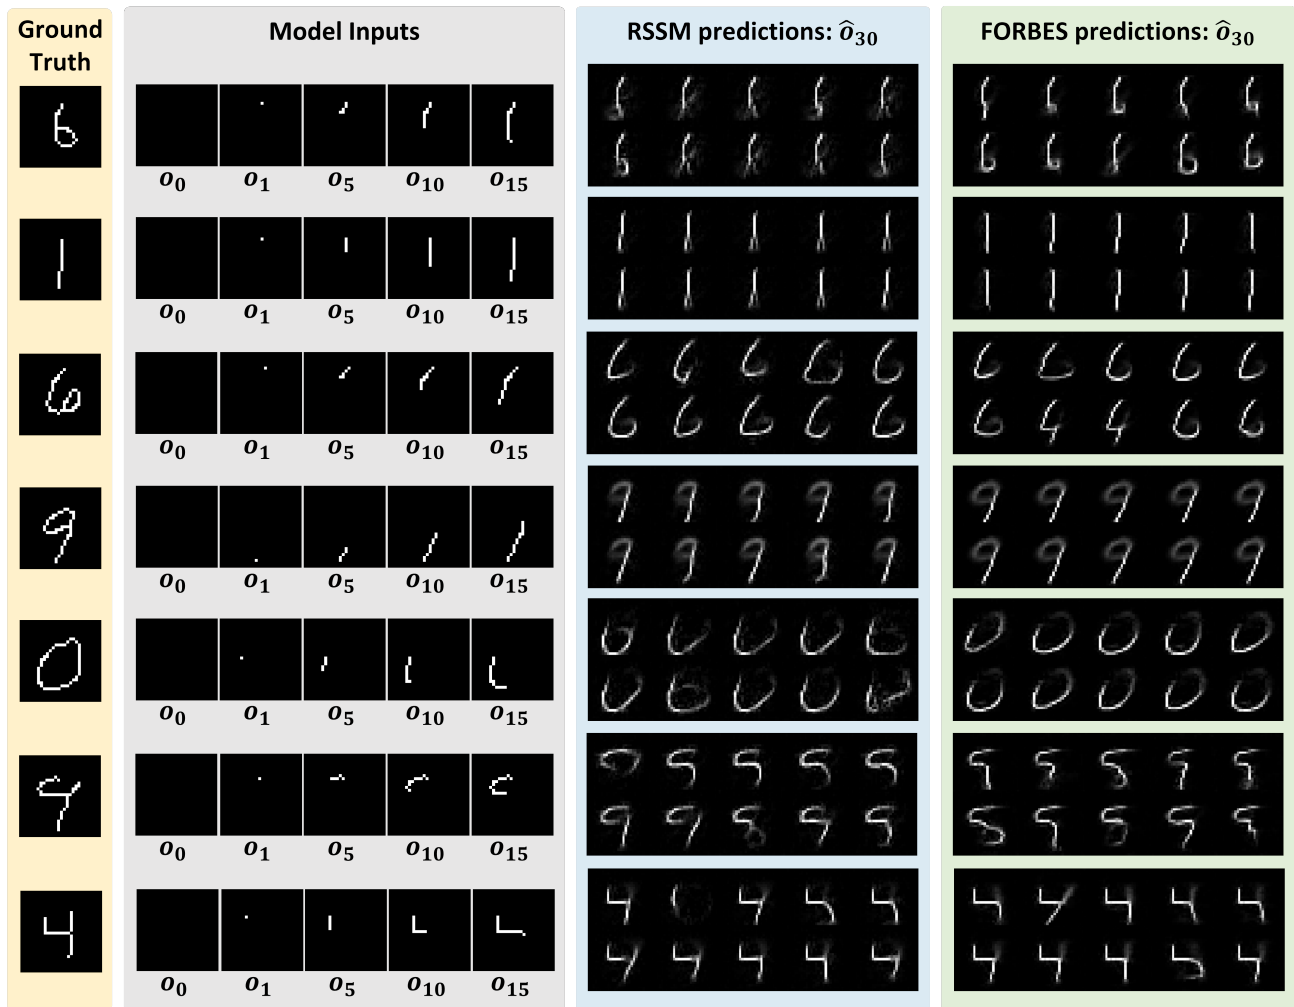



Figure 15: Additional prediction results on sequential MNIST of two models.

M. Reconstructions of the visual control tasks

In this section, we show the reconstructions of the visual control tasks during the evaluating phase.

For each environment, we use 10 frames. The left one is the original picture for each frame, and the right one is the reconstruction picture. The following results in Figure 16 show that FORBES can make high-quality reconstructions. The corresponding videos can be found in the supplementary material.



(a) Cartpole Swing Up



(b) Cheetah Run




(c) Finger Spin



(d) Hopper Stand



(e) Walker Run



(f) Walker Walk

Figure 16: The reconstruction results on of FORBES six environments from DeepMind Control Suite(Tassa et al., 2018).