

LEARNING TO REASON: DISTILLING HIERARCHY VIA SELF-SUPERVISION AND REINFORCEMENT LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

We present a hierarchical planning and control framework that enables an agent to perform various tasks and adapt to a new task flexibly. Rather than learning an individual policy for each particular task, the proposed framework, DISH, distills a hierarchical policy from a set of tasks by self-supervision and reinforcement learning. The framework is based on the idea of latent variable models that represent high-dimensional observations using low-dimensional latent variables. The resulting policy consists of two levels of hierarchy: (i) a planning module that *reasons* a sequence of latent intentions that would lead to optimistic future and (ii) a feedback control policy, shared across the tasks, that *executes* the inferred intention. Because the reasoning is performed in low-dimensional latent space, the learned policy can immediately be used to solve or adapt to new tasks without additional training. We demonstrate the proposed framework can learn compact representations (3-dimensional latent states for a 90-dimensional humanoid system) while solving a small number of imitation tasks, and the resulting policy is directly applicable to other types of tasks, i.e., navigation in cluttered environments.

1 INTRODUCTION

Reinforcement learning (RL) aims to compute the optimal control policy while an agent interacts with the environment. Recent advances in deep learning enable RL frameworks to utilize deep neural networks to efficiently represent and learn a policy having a flexible and expressive structure. As a result, we've been witnessing RL agents that already achieved or even exceeded human-level performances in particular tasks (Mnih et al., 2015; Silver et al., 2017). The core of intelligence, however, is not just to learn a policy for a particular problem instance, but to solve various multiple tasks or immediately adapt to a new task. Given that a huge computational burden makes it unrealistic to learn an individual policy for each task, an agent should be able to *reason* about its action. If predictions about consequences of actions are available, e.g., by using an internal model (Ha & Schmidhuber, 2018; Kaiser et al., 2019), an intelligent agent can plan a sequence of its actions. Involving planning procedures in a control policy could provide adaptiveness to an agent, but it is often not trivial to learn such a prediction & planning framework: First, it is difficult to obtain the exact internal dynamic model directly represented in high-dimensional state (observation) space. Model errors inevitably become larger in the high-dimensional space and are accumulated along the prediction (or planning) horizon. This prohibits planning methods from producing a valid prediction and so a sensible plan. Second, and perhaps more importantly, planning methods cannot help but relying on some dynamic programming or search procedures, which quickly become intractable for problems with high degrees of freedom (DOFs) because the size of search space grows exponentially with DOFs, i.e., the curse of dimensionality (LaValle, 2006).

Crucial evidence found in the cognitive science field is that there exists a certain type of hierarchical structure in the humans' motor control scheme addressing the aforementioned fundamental difficulty (Todorov & Ghahramani, 2003; Todorov, 2004). Such a hierarchical structure is known to utilize two levels of parallel control loops, operating in different time scales; in a coarser scale, the high-level loop generates task-relevant commands for the agent to perform a given task, and then in a finer time scale, the low-level loop maps those commands into control signals while actively reacting to disturbances that the high-level loop could not consider (e.g., the spinal cord) (Todorov & Ghahramani, 2003). Because the low-level loop does not passively generate control signals from high-level commands, the high-level loop is able to focus only on the task-relevant aspects of the envi-

ronment dynamics that can be represented in a low-dimensional form. Consequently, this hierarchical structure allows us for efficiently predicting and planning the future states to compute the commands.

Motivated by this evidence, we propose a framework, termed "DISH", that DISTills a Hierarchical structure for reasoning and control. As depicted in Fig. 1, the proposed framework has two levels of hierarchy. The high-level loop represents an agent's current state as a low-dimensional latent state and generates/reasons task-relevant high-level commands by predicting and planning the future in the latent space. The low-level loop receives the high-level commands as well as the current states and maps them into the high-dimensional control signal. Two different types of learning are required to build such a framework: (i) a low-dimensional latent representation for an internal model should be obtained from agent's own experiences via *self-supervised learning*; (ii) a control policy should be learned while interacting with the environment via *reinforcement learning*.

We combined these two learning problems by transforming a multitask RL problem into generative model learning using the control-inference duality (Levine, 2018; Todorov, 2008; Rawlik et al., 2012). In this perspective, an agent equipped with a low-level control policy is viewed as a generative model that outputs trajectories according to high-level commands. Reasoning the high-level commands is then considered as a posterior inference problem; we introduce a low-dimensional internal model to make this inference tractable. We demonstrate that the proposed framework can learn the compact representation (3-dimensional latent states for a humanoid robot having 90-dimensional states) and the control policy while solving a small number of imitation tasks, and the learned planning and control scheme is immediately applicable to new tasks, e.g., navigation through a cluttered environment.

2 RELATED WORK

Hierarchical RL: To apply task-specific policies learned from individual RL problems to various tasks, hierarchical structures are often considered where each learned policy serves as a low-level controller, i.e., as a "skill", and a high-level controller selects which skills to perform in the context the agent lies at (Peng et al., 2018; 2019; Merel et al., 2019a; Lee et al., 2019). Peng et al. (2018; 2019) trained robust control policies for imitating a broad range of example motion clips and integrated multiple skills into a composite policy capable of executing various tasks. Merel et al. (2019a) similarly trained many imitation policies and utilized them as individual skills that a high-level controller chooses based on the visual inputs. Lee et al. (2019) included transition policies which help the agent smoothly switch between the skills. Another line of approaches is using continuous-valued *latent* variables to represent skills (Co-Reyes et al., 2018; Gupta et al., 2018; Eysenbach et al., 2019; Florensa et al., 2017; Hausman et al., 2018). Co-Reyes et al. (2018) proposed an autoencoder-like framework where an encoder compresses trajectories into latent variables, a state decoder reconstructs trajectories, and a policy decoder provides a control policy to follow the reconstructed trajectory. Gupta et al. (2018); Eysenbach et al. (2019); Florensa et al. (2017) also introduced latent variables to efficiently represent various policies. Instead of using one *static* latent variable, Merel et al. (2019b) proposed a framework that encodes expert's demonstrations as latent *trajectories* and infers a latent trajectory from an unseen skill for one-shot imitation. Haarnoja et al. (2018a) proposed a hierarchical structure for RL problems where marginalization of low-level actions provides a new system for high-level action. In their framework, policies at all levels can be learned with different reward functions such that a high-level policy becomes easier to be optimized from the marginalization. Note that the above hierarchical RL approaches train the high-level policy by solving another RL problem; because the individual skill or the latent variables compress dynamics of the agent, variations of them provide efficient exploration for the high-level RL. Our framework also considers low-dimensional and continuous latent *trajectories* to represent various policies. Rather than learning a high-level policy, however, our framework learns an internal model with which the high-level module performs

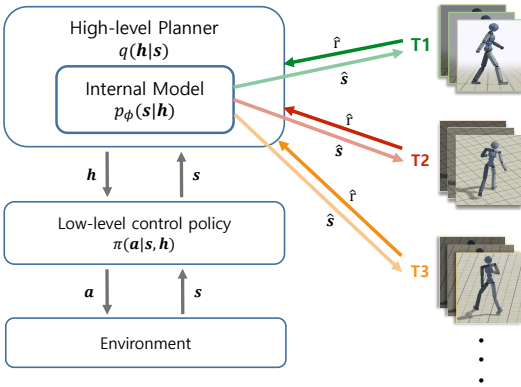


Figure 1: The DISH framework.

reasoning; the agent can efficiently reason its high-level commands by searching the low-dimensional latent space with the learned internal model. The learned planning/control structure is then directly applicable to new sets of tasks the agent hasn't met during training.

Model-based RL & Learning to Plan: Model-based RL algorithms attempt to learn the agent's dynamics and utilize the planning and control methods to perform tasks (Williams et al., 2017; Deisenroth et al., 2015; Chua et al., 2018). Williams et al. (2017); Chua et al. (2018) utilized deep neural networks to model the dynamics and adopted the model predictive control method on the learned dynamics; Deisenroth et al. (2015) used the Gaussian processes as system dynamics, which leads to the efficient and stable policy search. Though these methods have shown impressive results, they are not directly applicable to systems having high DOFs because high-dimensional modeling is hard to be exact and even advanced planning and control methods are not very scalable to such systems. One exceptional work was proposed by Ha & Schmidhuber (2018), where the variational autoencoder and the recurrent neural network are combined to model the dynamics of the observation. They showed that a simple linear policy w.r.t the low-dimensional latent state can control the low DOFs agent, but (i) high-DOFs systems require a more complicated policy structure to output high-dimensional actions and (ii) reasoning (or planning) by predicting the future is essential to solve a set of complex tasks. On the other hand, Ha et al. (2018a;b) trained the low-dimensional latent dynamics from expert's demonstrations and generated motion plans using the learned dynamics; the high-dimensional motion plans were able to be computed efficiently, but the control policy for executing those plans was not considered. Some recent works have attempted to build the policy network in such way that resembles the advanced planning and optimal control methods: Tamar et al. (2016) encoded the value iteration procedures into the network; Okada et al. (2017); Amos et al. (2018) wired the network so as to resemble the path-integral control and the iterative LQR methods, respectively. The whole policy networks are trained end-to-end and, interestingly, system dynamics and a cost function emerge during the learning procedure. However, these methods were basically designed just to mimic the expert's behaviors, i.e., addressing inverse RL problems, and also tried to find the control policy directly in the (possibly high-dimensional) state space.

3 DISH: DISTILLING HIERARCHY FOR PLANNING AND CONTROL

3.1 MULTITASK RL AS LATENT VARIABLE MODEL LEARNING

Suppose that a dynamical system with states $\mathbf{s} \in \mathcal{S}$ is controlled by actions $\mathbf{a} \in \mathcal{A}$, where the states evolve with the stochastic dynamics $p(\mathbf{s}_{k+1}|\mathbf{s}_k, \mathbf{a}_k)$ from the initial states $p(\mathbf{s}_1)$. Let $\tilde{r}_k(\mathbf{s}_k, \mathbf{a}_k)$ denote a reward function that the agent wants to maximize with the control policy $\pi_\theta(\mathbf{a}_k|\mathbf{s}_k)$. Reinforcement learning problems are then formulated as the following optimization problem:

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{q_\theta(\mathbf{s}_{1:K}, \mathbf{a}_{1:K})} \left[\sum_{k=1}^K \tilde{r}_k(\mathbf{s}_k, \mathbf{a}_k) \right], \quad (1)$$

where the *controlled* trajectory distribution q_θ is given by:

$$q_\theta(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}) \equiv p(\mathbf{s}_1) \prod_{k=1}^K p(\mathbf{s}_{k+1}|\mathbf{s}_k, \mathbf{a}_k) \pi_\theta(\mathbf{a}_k|\mathbf{s}_k). \quad (2)$$

By introducing an artificial binary random variable o_t , called the *optimality variable*, whose emission probability is given by exponential of a state-dependent reward, i.e. $p(O_k = 1|\mathbf{s}_k) = \exp(r_k(\mathbf{s}_k))$, and by defining an appropriate action prior $p(\mathbf{a})$ and corresponding the *uncontrolled* trajectory distribution, $p(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}) \equiv p(\mathbf{s}_1) \prod_{k=1}^K p(\mathbf{s}_{k+1}|\mathbf{s}_k, \mathbf{a}_k) p(\mathbf{a}_k)$, we can view the above RL problem as a probabilistic inference problem for a graphical model in Fig 2(a). The objective of such an inference problem is to find the optimal variational parameter, θ , such that the controlled trajectory distribution $q_\theta(\mathbf{s}_{1:K}, \mathbf{a}_{1:K})$ fits the posterior distribution $p(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}|O_{1:K} = 1)$ best. More detailed derivations of this duality can be found in Appendix A.2 or in the tutorial paper (Levine, 2018).

Rather than solving one particular task, i.e., one reward function, agents are often required to perform various tasks. Let \mathcal{T} be a set of tasks, and $\pi_{\theta_t^*}(\mathbf{a}_t|\mathbf{s}_t)$ be the optimal policy for t^{th} task, i.e.,

$$\theta_t^* = \arg \max_{\theta_t} \mathbb{E}_{q_{\theta_t}(\mathbf{s}_{1:K}, \mathbf{a}_{1:K})} \left[\sum_{k=1}^K \tilde{r}_k^{(t)}(\mathbf{s}_k, \mathbf{a}_k) \right], \quad \forall t \in \mathcal{T}. \quad (3)$$




Figure 2: (a) The conventional RL and (b) the proposed hierarchical RL framework. (c) The action-marginalized inference problem. (d) A low-dimensional LVM for the high-level reasoning.

For high DOF systems, where policies π_{θ_t} represent a mapping from high-dimensional state space to high-dimensional action space, it is computationally too expensive to optimize each policy individually. However, we can assume that tasks the agent needs to perform require similar solution properties and consequently the optimal policies have some sort of common structure. Then we can introduce a low-dimensional latent variable $\mathbf{h}^{(t)}$ to efficiently represent the policies as $\pi_{\theta}(\mathbf{a}_k|\mathbf{s}_k, \mathbf{h}^{(t)})$.

Such a hierarchical structure can be depicted as Fig. 2(b), where \mathbf{h} can be interpreted as high-level *commands*. We can then define the uncontrolled and the task t 's controlled trajectory distributions as

$$p(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}, \mathbf{h}_{1:K}) \equiv p(\mathbf{s}_1) \prod_{k=1}^K p(\mathbf{s}_{k+1}|\mathbf{s}_k, \mathbf{a}_k)p(\mathbf{a}_k)p(\mathbf{h}_k), \quad (4)$$

$$q_{\theta}^{(t)}(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}, \mathbf{h}_{1:K}) \equiv p(\mathbf{s}_1) \prod_{k=1}^K p(\mathbf{s}_{k+1}|\mathbf{s}_k, \mathbf{a}_k)\pi_{\theta}(\mathbf{a}_k|\mathbf{s}_k, \mathbf{h}_k)q^{(t)}(\mathbf{h}_k|\mathbf{s}_k), \quad (5)$$

receptively. In other words, the control policy π_{θ} is shared across all the tasks, actively mapping high-level commands \mathbf{h} , into actual actions, \mathbf{a} . Only high-level commands vary with the given task specifications. In the perspective of *control as inference*, a corresponding inference problem now has two parts: one for the policy parameter θ and another for the task-specific commands \mathbf{h} .

Similar to the latent model learning in Appendix A.3 and the control-inference duality in Appendix A.2, we can derive the following lower-bound of optimality likelihood $\mathcal{L}^{(t)}$ for a task t :

$$\begin{aligned} \log p_{\theta}(O_{1:K}^{(t)} = 1) &= \log \int p(O_{1:K}^{(t)} = 1|\mathbf{s}_{1:K})p(\tau) \frac{q_{\theta}^{(t)}(\tau)}{q_{\theta}^{(t)}(\tau)} d\tau \\ &\geq \mathbb{E}_{q_{\theta}^{(t)}(\tau)} \left[\sum_{k=1}^K r_k^{(t)}(\mathbf{s}_k) - \log \frac{\pi_{\theta}(\mathbf{a}_k|\mathbf{s}_k, \mathbf{h}_k)}{p(\mathbf{a}_k)} - \log \frac{q^{(t)}(\mathbf{h}_k|\mathbf{s}_k)}{p(\mathbf{h}_k)} \right] \equiv \mathcal{L}^{(t)}(\theta, q), \end{aligned} \quad (6)$$

with $\tau \equiv (\mathbf{s}_{1:K}, \mathbf{a}_{1:K}, \mathbf{h}_{1:K})$. This suggests a novel learning scheme of a hierarchical policy in (5). For a given task t , high-level commands \mathbf{h}_k are computed via variational inference. This inference procedure $q(\mathbf{h}|\mathbf{s})$ should take predictions about future rewards into account to generate \mathbf{h} , which can be interpreted as planning. To do so, we build an internal model via self-supervised learning and perform planning with the internal model. With the planning module equipped, a low-level policy $\pi_{\theta}(\mathbf{a}|\mathbf{s}, \mathbf{h})$ generates control actions \mathbf{a} as in RL problems, which can be trained using standard deep RL algorithms (Schulman et al., 2017; Haarnoja et al., 2018b).

3.2 SELF-SUPERVISED LEARNING OF INTERNAL MODEL

The role of $q(\mathbf{h}|\mathbf{s})$ is to compute the high-level commands that would lead to maximum accumulated rewards in the future; as shown in (6), this infers the commands that maximizes the likelihood of optimality variables when $O_{1:K} = 1$ were observed. Given that the ELBO gap is the KL-divergence between the posterior and variational distributions, it is obvious that more exact variational inference will make the lower bound tighter, thus directly leads to the agent's better performance as well as

the better policy learning. What would the exact posterior be like? Fig. 2(c) shows the graphical model of the inference problem that $q(\mathbf{h}|\mathbf{s})$ should address, which is obtained by marginalizing actions from Fig. 2(b); as also shown in (Haarnoja et al., 2018a), such a marginalization results in a new system with new control input \mathbf{h} , thus the inference problem in this level is again the RL/OC problem. To get the command at the moment, \mathbf{h}_1 , the inference procedure should compute the posterior command trajectories $\mathbf{h}_{1:T}$ by considering the dynamics and observations (the optimality variables), and marginalize the future commands $\mathbf{h}_{2:T}$ out. Though the dimension of \mathbf{h} is much smaller than that of \mathbf{a} , this inference problem is still not trivial to solve by two reasons: (i) The dynamics of states $p_\theta(s'|s, \mathbf{h}) = \int p(s'|s, \mathbf{a})\pi_\theta(\mathbf{a}|s, \mathbf{h})d\mathbf{a}$ contains the environment component of which information can be obtained only through expensive interactions with the environment. (ii) One might consider building a surrogate model $p_\phi(s'|s, \mathbf{h})$ via supervised learning with transition data obtained during low-level policy learning and utilizing the learned model for inference. However, learning high-dimensional transition model is hard to be accurate and the inference (planning) in high-dimensional space is intractable because of, e.g., the curse of dimensionality (Ha et al., 2018a).

However, we can reasonably assume that configurations that should be considered form some sort of low-dimensional manifold in the original space (Vernaza & Lee, 2012), and the closed-loop system with low-dimensional commands provides stochastic dynamics on that manifold. That is, we assume that a high-dimensional transition model in Fig. 2(c) can be represented as a latent variable model (LVM) in Fig. 2(d). Thus, our framework collects the trajectories from low-level policies and utilize them to learn a LVM for inference. Once this low-dimensional representation is obtained, any motion planning or inference algorithms can solve the variational inference problem very efficiently with the vastly restricted search space. The problem of learning LVMs can be formulated as a maximum likelihood estimation (MLE) problem. Suppose that we have collected a set of state trajectories and latent commands $\{\mathbf{s}_{1:K}^{(n)}, \mathbf{h}_{1:K}^{(n)}\}_{n=1, \dots, N}$. We then formulate the MLE problem as:

$$\phi^* = \arg \max_{\phi} \sum_n \log p_\phi(\mathbf{s}_{1:K}^{(n)}|\mathbf{h}_{1:K}^{(n)}). \quad (7)$$

Rather than using model in the state space directly, the states are assumed to be emerged from a latent dynamical system, where a latent state trajectory, $\mathbf{z}_{1:K}$, lies on a low-dimensional latent space \mathcal{Z} :

$$p_\phi(\mathbf{s}_{1:K}|\mathbf{h}_{1:K}) = \int p_\phi(\mathbf{s}_{1:K}|\mathbf{z}_{1:K})p_\phi(\mathbf{z}_{1:K}|\mathbf{h}_{1:K})d\mathbf{z}_{1:K}. \quad (8)$$

In particular, we consider the state space model where latent states follow stochastic transition dynamics with \mathbf{h} as inputs, i.e., the prior $p_\theta(\mathbf{z}_{1:K})$ is a probability measure of a following system:

$$\mathbf{z}_{k+1} = f_\phi(\mathbf{z}_k) + \sigma_\phi(\mathbf{z}_k)(\mathbf{h}_k + \mathbf{w}_k), \quad \mathbf{w}_k \sim \mathcal{N}(0, I) \quad (9)$$

and also a conditional likelihood of a state trajectory is assumed to be factorized along the time axis as: $\mathbf{s}_k \sim \mathcal{N}(\mu_\phi(\mathbf{z}_k), \Sigma_\phi(\mathbf{z}_k)) \forall k$. The resulting sequence modeling is self-supervised learning problem that has been extensively studied recently (Karl et al., 2017; Krishnan et al., 2017; Fraccaro et al., 2017; Ha et al., 2018b). In particular, we adopt the idea of Adaptive path-integral autoencoder in (Ha et al., 2018b), where the variational distribution is parameterized by the controls, \mathbf{u} , and an initial distribution, q_0 , i.e., the proposal $q_{\mathbf{u}}(\mathbf{z}_{[0,T]})$ is a probability measure of a following system:

$$\mathbf{z}_{k+1} = f_\phi(\mathbf{z}_k) + \sigma_\phi(\mathbf{z}_k)(\mathbf{h}_k + \mathbf{u}_k + \mathbf{w}_k), \quad \mathbf{w}_k \sim \mathcal{N}(0, I). \quad (10)$$

Compared to the original formulation in (Ha et al., 2018b), the probability model here is conditioned on the commands, $\mathbf{h}_{1:K}$. Consequently, the problem has become conditional generative model learning (Sohn et al., 2015) and effectively the control input prior has changed from $\mathcal{N}(\mathbf{0}, I)$ to $\mathcal{N}(\mathbf{h}, I)$ as in (9) and (10) (Williams et al., 2017). Note that it is also possible to first obtain a low-dimensional representation considering each state (not trajectory) independently and then fit their dynamics using RNNs like World Model (Ha & Schmidhuber, 2018), or to stack two consecutive observations and learn the dynamical model considering the stacked data as one observation like E2C (Watter et al., 2015). However, Ha et al. (2018b) showed that the representations learned from the short horizon data easily fail to extract enough temporal information and a latent dynamical model suitable for planning can be well-obtained only when considering long trajectories.

3.3 PLANNING WITH LEARNED INTERNAL MODEL

Once the LVM is trained, we can efficiently explore the state space \mathcal{S} through the latent state \mathbf{z} and infer the latent commands $\mathbf{h}_{1:K}$ that are likely to result in high rewards; in particular, we adopt a

simple particle filter algorithm for inference, because it performs well for non-linear and non-Gaussian systems (Ha et al., 2018a; Piche et al., 2019). Particle filtering, which is also called the sequential Monte-Carlo, utilizes a set of samples and their weights to represent a posterior distribution. Starting from the initial state, it propagates a set of samples according to the dynamics (9) and updates the weights using the observation likelihood as $w' \propto w \times p(O_k = 1 | \mathbf{s}_k)$. It also resamples the low-weighted particles to maintain the effective sample size. In the perspective of this work, this procedure can be view as the agent simulating multiple future trajectories with the internal model, assigning each of them according to the reward, and reasoning the command that leads to the best-possible future. The detailed explanation is elaborated in Appendix A.4 and in Algorithm 2. Note that any other inference/planning algorithms compatible with Fig. 2(d) can be used instead.

3.4 MAIN DISH

In summary, the overall learning procedure is shown in Algorithm 1. The procedure consists of outer internal model learning loop and inner policy update loop. During the policy update stage, the algorithm collects trajectories by sampling a task and solve that using the hierarchical policy. At each timestep, high-level commands are computed by the planning module equipped with the internal latent model, and the low-level policy generates actions the agent takes. Utilizing the collected trajectories, the low-level policy is updated via reinforcement learning algorithm (e.g., policy gradient methods). After low-level policy update, DISH collects rollouts by random sampling a latent variable \mathbf{h} , and the latent model is learned via self-supervised learning of LVMs.

Algorithm 1 DIStilling Hierarchy for Planning and Control

```

1: Initialize policy  $\theta$  and latent model  $\phi$ 
2: for  $l = 1, \dots, L$  do
3:   for  $m = 1, \dots, M$  do
4:     for  $r = 1, \dots, N_{\text{rollout}}$  do
5:       Sample a task randomly  $t \in \mathcal{T}$ 
6:       Compute a high-level command  $\mathbf{h} \sim q_\phi(\mathbf{h} | \mathbf{s}; t)$ 
7:       Run policy  $\pi_\theta(\mathbf{a} | \mathbf{s}, \mathbf{h})$  and take action
8:       Store the trajectory  $\tau$  into experience buffer
9:     end for
10:    Optimize policy  $\pi_\theta$  using, e.g., PPO ▷ Eq. (6)
11:  end for
12:  Random sample  $\mathbf{h}$  and collect rollouts
13:  Learn latent model using, e.g., APIAE ▷ Eq. (7)
14: end for

```

4 EXPERIMENT


We evaluate the effectiveness of DISH framework on planning and control for the high dimensional system. The experiment has been conducted with humanoid model that has 34 degrees of freedom, 197 state features, and 36 action parameters. We built our low-level and high-level policy networks with shallow MLPs. Since humanoid model is highly unstable, low-level policy has been pre-trained by using motion imitation approach (Peng et al., 2018; 2019). Three motion capture clips, going straight ($t_i = 0$), turning right ($t_i = 1$), and left ($t_i = 2$) have been used for pre-training stage. During imitation learning, high-level controller is initialized by $q_\phi(\mathbf{h} | \mathbf{s}; t_i) = \mathbf{h}_{t_i}$ where $\mathbf{h}_0 = 0, \mathbf{h}_1 = +1, \mathbf{h}_2 = -1$. For training scheme, we adopted PPO and APIAE to train low-level and high-level modules, respectively.

4.1 LEARNING INTERNAL MODELS

By using rollouts from pre-trained low-level policy, we trained DISH internal model of 3-dimensional latent state and 1-dimensional latent control variable. Fig. 3(a) shows the learned latent space colored by the angular velocity of the motion. We can observe that the manifold forms a cylinder shape in 3-dimensional space, where the locomotion phase is encoded along the radius and the angular velocity is encoded as main axis of the cylinder. We compared our internal model with three base

Table 1: Comparison between different types of internal models.

	Reconstruction	$\ ref-tru\ $	$\ plan-ref\ $	$\ plan-tru\ $
DISH (ours, Fig. 1 & Fig.2(d))	0.3820	0.1638	0.1576	0.0930
sas' (w/o hierarchy & LVMs, Fig. 2(a))	0.1289	-	0.1771	-
shs' (w/o LVMs, Fig. 2(c))	0.1393	0.2226	0.1579	0.2231
zaz' (w/o hierarchy)	2.3351	-	0.2589	-

Figure 3: (a) Leanred latent model. (b), (c) Rollout trajectories of DISH, shs' . Color shows latent control value.

lines: (i) sas' that doesn't consider neither the hierarchical structure nor the LVMs, (ii) shs' that utilizes the hierarchy but doesn't learn the low-dimensional latent space, and (iii) zaz' that considers the latent space but without the hierarchical structure. For zaz' model, the variational autoencoder approach has been used to learn recognition and generative network between high-dimensional state and low-dimensional latent state. (Ha & Schmidhuber, 2018). Transition model of sas' , shs' , and zaz' has been trained by using rollout trajectories. The four models are compared on three imitation tasks, going straight, turning left and right. Table 1 reports the RMS errors for reconstruction and differences between reference, planned, and executed trajectories. Since training sas' is just a simple supervised learning problem, it had the lowest reconstruction error but the computed action from such the internal model couldn't make the humanoid walk. zaz' also failed to let the robot walk, because reasoning of the high-dimensional action can't be accurate enough. Only the internal models considering hierarchical structure could make the robot walk, and as shown in the table, our framework was able to generate the most executable and valuable plans. Fig.3(b) and 3(c) show the rollouts from each internal model, and we can observe our internal model can generate more meaningful trajectories. The process by which the internal model is trained can be seen in the supplementary video.¹

4.2 PLANNING AND CONTROL WITH LEARNED INTERNAL MODEL

To achieve better internal model and corresponding low-level policy, we performed one more training loop. Results shown in supplementary video illustrates DISH successfully solved long-horizon planning tasks with learned internal model.

5 CONCLUSION

We proposed a framework to learn a hierarchical policy for an RL agent. In the proposed policy, the high-level loop plans the agent's motion by predicting its low-dimensional "task-specific" futures and the low-level loop maps the high-level commands into actions while actively reacting to the environment using its own state feedback loop. Such the sophisticated separation was able to emerge

¹<https://drive.google.com/drive/folders/18AaPwxXZn3SqdHXEaoq5Gxz-ZObMXRW?usp=sharing>




Figure 4: Cluttered environments for navigation tasks.

because two loops operated in different scales; the high-level planning loop only focuses on task-specific low-dimensional aspects in a coarser timescale, which enables it to plan relatively long-term futures. In order to learn the internal model for planning, we took advantage of recent advances in self-supervised learning of sequential data, while the low-level control policy is learned using a deep RL algorithm. By alternately optimizing both the LVM and the policy, the proposed framework was able to construct a meaningful internal model as well as a versatile control policy.

As future works, it would be interesting to incorporate visual inputs into the high-level reasoning module as suggested in (Merel et al., 2019a). Though only continuous latent variables were considered in our framework, utilizing discrete variables such as "logic" or "modes" (Toussaint et al., 2018) also seems to be a promising direction. Also, besides imitation of experts, an agent should be able to learn from play as suggested in (Lynch et al., 2019).

REFERENCES

- Brandon Amos, Ivan Jimenez, Jacob Sacks, Byron Boots, and J Zico Kolter. Differentiable MPC for end-to-end planning and control. In *Advances in Neural Information Processing Systems*, pp. 8299–8310, 2018.
- Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. *Advances in Neural Information Processing Systems*, 2018.
- John D Co-Reyes, YuXuan Liu, Abhishek Gupta, Benjamin Eysenbach, Pieter Abbeel, and Sergey Levine. Self-consistent trajectory autoencoder: Hierarchical reinforcement learning with trajectory embeddings. In *International Conference on Machine Learning*, 2018.
- Marc Peter Deisenroth, Dieter Fox, and Carl Edward Rasmussen. Gaussian processes for data-efficient learning in robotics and control. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):408–423, 2015.
- Benjamin Eysenbach, Abhishek Gupta, Julian Ibarz, and Sergey Levine. Diversity is all you need: Learning skills without a reward function. In *International Conference on Learning Representations*, 2019.
- Carlos Florensa, Yan Duan, and Pieter Abbeel. Stochastic neural networks for hierarchical reinforcement learning. In *International Conference on Learning Representations*, 2017.
- Marco Fraccaro, Simon Kamronn, Ulrich Paquet, and Ole Winther. A disentangled recognition and nonlinear dynamics model for unsupervised learning. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 3604–3613, 2017.
- Abhishek Gupta, Russell Mendonca, YuXuan Liu, Pieter Abbeel, and Sergey Levine. Meta-reinforcement learning of structured exploration strategies. In *Advances in Neural Information Processing Systems*, 2018.
- David Ha and Jürgen Schmidhuber. Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems*, pp. 2455–2467, 2018.

- Jung-Su Ha, Hyeok-Joo Chae, and Han-Lim Choi. Approximate inference-based motion planning by learning and exploiting low-dimensional latent variable models. *IEEE Robotics and Automation Letters*, 3(4):3892–3899, 2018a.
- Jung-Su Ha, Young-Jin Park, Hyeok-Joo Chae, Soon-Seo Park, and Han-Lim Choi. Adaptive path-integral autoencoders: Representation learning and planning for dynamical systems. In *Advances in Neural Information Processing Systems*, pp. 8941–8952, 2018b.
- Tuomas Haarnoja, Kristian Hartikainen, Pieter Abbeel, and Sergey Levine. Latent space policies for hierarchical reinforcement learning. In *International Conference on Machine Learning*, 2018a.
- Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International Conference on Machine Learning*, 2018b.
- Karol Hausman, Jost Tobias Springenberg, Ziyu Wang, Nicolas Heess, and Martin Riedmiller. Learning an embedding space for transferable robot skills. In *International Conference on Learning Representations*, 2018.
- Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H Campbell, Konrad Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, et al. Model-based reinforcement learning for atari. *arXiv preprint arXiv:1903.00374*, 2019.
- Hilbert Johan Kappen and Hans Christian Ruiz. Adaptive importance sampling for control and inference. *Journal of Statistical Physics*, 162(5):1244–1266, 2016.
- Maximilian Karl, Maximilian Soelch, Justin Bayer, and Patrick van der Smagt. Deep variational bayes filters: Unsupervised learning of state space models from raw data. *International Conference on Learning Representations (ICLR)*, 2017.
- Rahul G Krishnan, Uri Shalit, and David Sontag. Structured inference networks for nonlinear state space models. In *AAAI Conference on Artificial Intelligence*, pp. 2101–2109, 2017.
- Steven M LaValle. *Planning algorithms*. Cambridge university press, 2006.
- Youngwoon Lee, Shao-Hua Sun, Sriram Somasundaram, Edward Hu, and Joseph J Lim. Composing complex skills by learning transition policies with proximity reward induction. In *International Conference on Learning Representations*, 2019.
- Sergey Levine. Reinforcement learning and control as probabilistic inference: Tutorial and review. *arXiv preprint arXiv:1805.00909*, 2018.
- Corey Lynch, Mohi Khansari, Ted Xiao, Vikash Kumar, Jonathan Tompson, Sergey Levine, and Pierre Sermanet. Learning latent plans from play. *arXiv preprint arXiv:1903.01973*, 2019.
- Josh Merel, Arun Ahuja, Vu Pham, Saran Tunyasuvunakool, Siqi Liu, Dhruva Tirumala, Nicolas Heess, and Greg Wayne. Hierarchical visuomotor control of humanoids. In *International Conference on Learning Representations*, 2019a.
- Josh Merel, Leonard Hasenclever, Alexandre Galashov, Arun Ahuja, Vu Pham, Greg Wayne, Yee Whye Teh, and Nicolas Heess. Neural probabilistic motor primitives for humanoid control. In *International Conference on Learning Representations*, 2019b.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidfjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529, 2015.
- Mustafa Mukadam, Jing Dong, Xinyan Yan, Frank Dellaert, and Byron Boots. Continuous-time gaussian process motion planning via probabilistic inference. *The International Journal of Robotics Research*, 37(11):1319–1340, 2018.
- Masashi Okada, Luca Rigazio, and Takenobu Aoshima. Path integral networks: End-to-end differentiable optimal control. *arXiv preprint arXiv:1706.09597*, 2017.

- Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. DeepMimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Trans. Graph.*, 37(4):143:1–143:14, July 2018. ISSN 0730-0301. doi: 10.1145/3197517.3201311. URL <http://doi.acm.org/10.1145/3197517.3201311>.
- Xue Bin Peng, Michael Chang, Grace Zhang, Pieter Abbeel, and Sergey Levine. Mpc: Learning composable hierarchical control with multiplicative compositional policies. In *Advances in Neural Information Processing Systems*, 2019.
- Alexandre Piche, Valentin Thomas, Cyril Ibrahim, Yoshua Bengio, and Chris Pal. Probabilistic planning with sequential monte carlo methods. In *International Conference on Learning Representations*, 2019.
- Konrad Rawlik, Marc Toussaint, and Sethu Vijayakumar. On stochastic optimal control and reinforcement learning by approximate inference. In *Robotics: Science and Systems*, 2012.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354, 2017.
- Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Advances in neural information processing systems*, pp. 3483–3491, 2015.
- Aviv Tamar, Yi Wu, Garrett Thomas, Sergey Levine, and Pieter Abbeel. Value iteration networks. In *Advances in Neural Information Processing Systems*, pp. 2154–2162, 2016.
- Emanuel Todorov. Optimality principles in sensorimotor control. *Nature neuroscience*, 7(9):907, 2004.
- Emanuel Todorov. General duality between optimal control and estimation. In *IEEE Conference on Decision and Control*, pp. 4286–4292. IEEE, 2008.
- Emanuel Todorov and Zoubin Ghahramani. Unsupervised learning of sensory-motor primitives. In *Engineering in Medicine and Biology Society, 2003. Proceedings of the 25th Annual International Conference of the IEEE*, volume 2, pp. 1750–1753. IEEE, 2003.
- Marc Toussaint. Robot trajectory optimization using approximate inference. In *Proceedings of the 26th annual international conference on machine learning*, pp. 1049–1056. ACM, 2009.
- Marc Toussaint, Kelsey Allen, Kevin Smith, and Joshua B Tenenbaum. Differentiable physics and stable modes for tool-use and manipulation planning. In *Robotics: Science and Systems*, 2018.
- Paul Vernaza and Daniel D Lee. Learning and exploiting low-dimensional structure for efficient holonomic motion planning in high-dimensional spaces. *The International Journal of Robotics Research*, 31(14):1739–1760, 2012.
- Manuel Watter, Jost Springenberg, Joschka Boedecker, and Martin Riedmiller. Embed to control: A locally linear latent dynamics model for control from raw images. In *Advances in Neural Information Processing Systems*, 2015.
- Grady Williams, Nolan Wagener, Brian Goldfain, Paul Drews, James M Rehg, Byron Boots, and Evangelos A Theodorou. Information theoretic MPC for model-based reinforcement learning. In *International Conference on Robotics and Automation (ICRA)*, 2017.
- Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, volume 8, pp. 1433–1438. Chicago, IL, USA, 2008.

A APPENDIX

A.1 CONTROL-INFERENCE DUALITY

One theoretical concept this work extensively takes advantage of is the duality between optimal control (OC) and probabilistic inference (Levine, 2018; Todorov, 2008; Rawlik et al., 2012). The idea is that, if we consider an artificial binary observation whose emission probability is given by the exponential of a negative cost, an OC problem can be reformulated as an equivalent inference problem. In this case, the objective is to find the trajectory or control policy that maximizes the likelihood of the observations along the trajectory. One advantage of this perspective is that in order to solve the OC or RL problems, we can adopt any powerful and flexible inference methods, e.g., the expectation propagation (Toussaint, 2009), the particle filtering (Ha et al., 2018a; Piche et al., 2019), or the inference for Gaussian processes (Mukadam et al., 2018). In addition to utilizing efficient inference methods, this work also enjoys the duality to transform a multi-task RL problem into a generative model learning problem, which enables an agent to distill a low-dimensional representation and a versatile control policy in a combined framework.

A.2 REINFORCEMENT LEARNING AS PROBABILISTIC INFERENCE

For easier reference, we restate the RL problem and the controlled trajectory distribution here:

$$\theta^* = \arg \max_{\theta} \mathbb{E}_{q_{\theta}(\mathbf{s}_{1:K}, \mathbf{a}_{1:K})} \left[\sum_{k=1}^K \tilde{r}_k(\mathbf{s}_k, \mathbf{a}_k) \right], \quad (11)$$

$$q_{\theta}(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}) \equiv p(\mathbf{s}_1) \prod_{k=1}^K p(\mathbf{s}_{k+1} | \mathbf{s}_k, \mathbf{a}_k) \pi_{\theta}(\mathbf{a}_k | \mathbf{s}_k), \quad (12)$$

respectively. It is well known in the literature that the above optimization (11) also can be viewed as a probabilistic inference problem for a certain type of graphical models (Levine, 2018; Todorov, 2008; Rawlik et al., 2012). Suppose we have an artificial binary random variable o_t , called the *optimality variable*, whose emission probability is given by exponential of a state-dependent reward, i.e.,

$$p(o_k = 1 | \mathbf{s}_k) = \exp(r_k(\mathbf{s}_k)), \quad (13)$$

and the action prior $p(\mathbf{a}_k)$ defines the *uncontrolled* trajectory distribution (see also Fig. 2(a)):

$$p(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}) \equiv p(\mathbf{s}_1) \prod_{k=1}^K p(\mathbf{s}_{k+1} | \mathbf{s}_k, \mathbf{a}_k) p(\mathbf{a}_k). \quad (14)$$

Then we can derive the evidence lower-bound (ELBO) for the variational inference:

$$\begin{aligned} \log p(O_{1:K}) &= \log \int p(O_{1:K} | \mathbf{s}_{1:K}) p(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}) d\mathbf{s}_{1:K} d\mathbf{a}_{1:K} \\ &= \log \int p(O_{1:K} | \mathbf{s}_{1:K}) p(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}) \frac{q_{\theta}(\mathbf{s}_{1:K}, \mathbf{a}_{1:K})}{q_{\theta}(\mathbf{s}_{1:K}, \mathbf{a}_{1:K})} d\mathbf{s}_{1:K} d\mathbf{a}_{1:K} \\ &\geq \mathbb{E}_{q_{\theta}(\mathbf{s}_{1:K}, \mathbf{a}_{1:K})} \left[\sum_{k=1}^K \left(\log p(O_k | \mathbf{s}_k) - \log \frac{\pi_{\theta}(\mathbf{a}_k | \mathbf{s}_k)}{p(\mathbf{a}_k)} \right) \right] \\ &= \mathbb{E}_{q_{\theta}(\mathbf{s}_{1:K}, \mathbf{a}_{1:K})} \left[\sum_{k=1}^K r_k(\mathbf{s}_k) - \log \frac{\pi_{\theta}(\mathbf{a}_k | \mathbf{s}_k)}{p(\mathbf{a}_k)} \right] \equiv \mathcal{L}(\theta). \end{aligned} \quad (15)$$

The ELBO maximization in (15) becomes equivalent to the reinforcement learning in (11) by choosing an action prior $p(\mathbf{a}_k)$ and parameterized policy family $\pi_{\theta}(\mathbf{a}_k | \mathbf{s}_k)$ to match $\tilde{r}_k = r_k - \log \frac{\pi_{\theta}}{p}$. Similar to (19), the above maximization means to find the control policy π_{θ} resulting in the variational distribution that best approximates the posterior trajectory distribution when all the optimality variables were observed $p(\mathbf{s}_{1:K}, \mathbf{a}_{1:K} | O_{1:K} = 1)$.

²For example, when $p(\mathbf{a}_k)$ and $\pi_{\theta}(\mathbf{a}_k | \mathbf{s}_k)$ are given as Gaussian with the same covariance, $\log \frac{\pi_{\theta}}{p}$ encodes quadratic penalty on the control effort; when $p(\mathbf{a})$ is given as an uninformative uniform distribution, $\log \frac{\pi_{\theta}}{p}$ becomes the entropy regularization term in the maximum entropy reinforcement learning (Ziebart et al., 2008; Haarnoja et al., 2018b).

A.3 SELF-SUPERVISED LEARNING OF LATENT DYNAMICAL MODELS

Self-supervised learning is an essential approach that allows an agent to learn underlying dynamics from sequential high-dimensional sensory inputs. The learned dynamical model can be utilized to predict and plan the future state of the agent. By assuming that observations were emerged from the low-dimensional latent states, the learning problems are formulated as latent model learning, which includes an intractable posterior inference of latent states for given input data (Karl et al., 2017; Krishnan et al., 2017; Fraccaro et al., 2017; Ha et al., 2018b).

Suppose that a set of observation sequences $\{\mathbf{s}_{1:K}^{(n)}\}_{n=1,\dots,N}$ is given, where $\mathbf{s}_{1:K}^{(n)} \equiv \{\mathbf{s}_k; \forall k = 1, \dots, K\}^{(n)}$ are i.i.d. sequences of observation that lie on (possibly high-dimensional) data space \mathcal{S} . The goal of the self-supervised learning problem of interest is to build a probabilistic model that well describes the given observations. The problem is formulated as a maximum likelihood estimation (MLE) problem by parameterizing a probabilistic model with ϕ :

$$\phi^* = \arg \max_{\phi} \sum_n \log p_{\phi}(\mathbf{s}_{1:K}^{(n)}). \quad (16)$$

For latent dynamic models, we assume that the observations are emerged from a latent dynamical system, where a latent state trajectory, $\mathbf{z}_{1:K} \equiv \{\mathbf{z}_k; \forall k \in 1, \dots, K\}$, lies on a (possibly low-dimensional) latent space \mathcal{Z} :

$$p_{\phi}(\mathbf{s}_{1:K}) = \int p_{\phi}(\mathbf{s}_{1:K}|\mathbf{z}_{1:K})p_{\phi}(\mathbf{z}_{1:K})d\mathbf{z}_{1:K}, \quad (17)$$

where $p_{\phi}(\mathbf{s}_{1:K}|\mathbf{z}_{1:K})$ and $p_{\phi}(\mathbf{z}_{1:K})$ are called a conditional likelihood and a prior distribution, respectively. Since the objective function (16) contains the intractable integration, it cannot be optimized directly. To circumvent the intractable inference, a variational distribution $q(\cdot)$ is introduced and then a surrogate loss function $\mathcal{L}(q, \phi; \mathbf{s}_{1:K})$, which is called the evidence lower bound (ELBO), can be considered alternatively:

$$\begin{aligned} \log p_{\phi}(\mathbf{s}_{1:K}) &= \log \int p_{\phi}(\mathbf{s}_{1:K}|\mathbf{z}_{1:K})p_{\phi}(\mathbf{z}_{1:K})d\mathbf{z}_{1:K} \\ &\geq \mathbb{E}_{q(\mathbf{z}_{1:K})} \left[\log \frac{p_{\phi}(\mathbf{s}_{1:K}|\mathbf{z}_{1:K})p_{\phi}(\mathbf{z}_{1:K})}{q(\mathbf{z}_{1:K})} \right] \\ &\equiv \mathcal{L}(q, \phi; \mathbf{s}_{1:K}), \end{aligned} \quad (18)$$

where $q(\cdot)$ can be any probabilistic distribution over \mathcal{Z} of which support includes that of $p_{\theta}(\cdot)$. Note that the gap between the log-likelihood and the ELBO is the Kullback-Leibler (KL) divergence between $q(\mathbf{z})$ and the posterior $p_{\theta}(\mathbf{z}_{1:K}|\mathbf{s}_{1:K})$:

$$\log p_{\phi}(\mathbf{s}_{1:K}) - \mathcal{L}(q, \phi; \mathbf{s}_{1:K}) = D_{KL}(q(\mathbf{z}_{1:K})||p_{\phi}(\mathbf{z}_{1:K}|\mathbf{s}_{1:K})). \quad (19)$$

One of the most general approaches is the expectation-maximization (EM) style optimization where, alternately, (i) E-step denotes an inference procedure where an optimal variational distribution q^* is computed for given ϕ and (ii) M-step maximizes the ELBO w.r.t. model parameter ϕ for given q^* .

Note that if we construct the whole inference and generative procedures as one computational graph, all the components can be learned by efficient end-to-end training (Karl et al., 2017; Krishnan et al., 2017; Fraccaro et al., 2017; Ha et al., 2018b). In particular, Ha et al. (2018b) proposed the adaptive path-integral autoencoder (APIAE), a framework that utilizes the optimal control method; this framework is suitable to this work because we want to perform the planning in the learned latent space. APIAE considers the state-space model in which the latent states are governed by a stochastic dynamical model, i.e., the prior $p_{\phi}(\mathbf{z}_{1:K})$ is a probability measure of a following system:

$$\mathbf{z}_{k+1} = f_{\phi}(\mathbf{z}_k) + \sigma_{\phi}(\mathbf{z}_k)\mathbf{w}_k, \quad \mathbf{z}_0 \sim p_0(\cdot), \quad \mathbf{w}_k \sim \mathcal{N}(0, I). \quad (20)$$

Additionally, a conditional likelihood of sequential observations is factorized along the time axis:

$$p_{\phi}(\mathbf{s}_{1:K}|\mathbf{z}_{1:K}) = \prod_{k=1}^K p_{\phi}(\mathbf{s}_k|\mathbf{z}_k). \quad (21)$$

If the variational distribution is parameterized by the control input $\mathbf{u}_{1:K-1}$ and the initial state distribution q_0 as:

$$\mathbf{z}_{k+1} = f_\phi(\mathbf{z}_k) + \sigma_\phi(\mathbf{z}_k) (\mathbf{u}_k + \mathbf{w}_k), \mathbf{z}_0 \sim q_0(\cdot), \mathbf{w}_k \sim \mathcal{N}(0, I), \quad (22)$$

the ELBO can be written in the following form:

$$\mathcal{L} = \mathbb{E}_{q_u} \left[\log p_\phi(\mathbf{s}_{1:K} | \mathbf{z}_{1:K}) + \log \frac{p_0(\mathbf{z}_0)}{q_0(\mathbf{z}_0)} - \sum_{k=1}^{K-1} \frac{1}{2} \|\mathbf{u}_k\|^2 + \mathbf{u}_k^\top \mathbf{w}_k \right]. \quad (23)$$

Then, the problem of finding the optimal variational parameters \mathbf{u}^* and q_0^* (or equivalently, the best approximate posterior) can be formulated as a stochastic optimal control (SOC) problem:

$$\mathbf{u}^*, q_0^* = \arg \min_{\mathbf{u}, q_0} \mathbb{E}_{q_u(\mathbf{z}_{1:K})} \left[V(\mathbf{z}_{1:K}) + \sum_{k=1}^{K-1} \frac{1}{2} \|\mathbf{u}_k\|^2 + \mathbf{u}_k^\top \mathbf{w}_k \right], \quad (24)$$

where $V(\mathbf{z}_{1:K}) \equiv -\log \frac{p_0(\mathbf{z}(0))}{q_0(\mathbf{z}(0))} - \sum_{k=1}^K \log p_\phi(\mathbf{s}_k | \mathbf{z}_k)$ serves as a state cost of the SOC problem. Ha et al. (2018b) constructed the differentiable computational graph that resembles the path-integral control procedure to solve the above SOC problem, and trained the whole architecture including the latent dynamics, $p_0(\mathbf{z})$, $f_\phi(\mathbf{z})$ and $\sigma_\phi(\mathbf{z})$, and the generative network, $p_\phi(\mathbf{s} | \mathbf{z})$ through the end-to-end training.

A.4 PLANNING BY PARTICLE FILTERING

Algorithm 2 Particle Filtering with Internal Model for Planning

- 1: Initialize $\forall i \in \{1, \dots, N_{\text{particle}}\} : \mathbf{z}_1^{(i)} \sim q_\phi(\cdot | \mathbf{s}_{\text{cur}})$ and $w_1^{(i)} = 1/N_{\text{particle}}$
 - 2: **for** $k = 2, \dots, K_{\text{plan}}$ **do**
 - 3: **for** $i = 1, \dots, N_{\text{particle}}$ **do**
 - 4: $\mathbf{z}_k^{(i)} = f_\phi(\mathbf{z}_{k-1}^{(i)}) + \sigma_\phi(\mathbf{z}_{k-1}^{(i)}) (\mathbf{h}_{k-1}^{(i)} + \mathbf{w}_{k-1}^{(i)})$, $\mathbf{w}_{k-1}^{(i)} \sim \mathcal{N}(0, I)$
 - 5: $\mathbf{s}_k^{(i)} \sim \mathcal{N}(\mu_\phi(\mathbf{z}_k^{(i)}), \Sigma_\phi(\mathbf{z}_k^{(i)}))$.
 - 6: $w_k^{(i)} = w_{k-1}^{(i)} \exp(r_k(\mathbf{s}_k^{(i)}))$
 - 7: **end for**
 - 8: $w_k^{(i)} = w_k^{(i)} / \sum_j w_k^{(j)}$, $\forall i \in \{1, \dots, N_{\text{particle}}\}$
 - 9: Resample $\{\mathbf{z}_{1:k}^{(i)}, \mathbf{w}_{1:k}^{(i)}\}$ if $(\sum_i (w_k^{(i)})^2)^{-1} < N_{\text{particle}}/3$
 - 10: **end for**
 - 11: **return** $\mathbf{h}_1^* = \sum_i w_{K_{\text{plan}}}^{(i)} \mathbf{w}_1^{(i)}$
-

At each time step, the high-level planner takes the current state as an argument and required to output the commands by predicting the future trajectory and corresponding reward. We adopted the particle filter algorithm to perform such the reasoning and the pseudo code is shown in Algorithm 2. The particle filter algorithm attempts to represent the posterior distribution using a set of samples. The algorithm first samples N_{particle} initial latent states using the inference network (which is a part of the learned internal model) and assigns the same weights for them. During the forward recursion, the particles are propagated using the latent dynamics of the internal model (line 4), and the corresponding configurations are generated through the learned model (line 5). The weights of all particles are then updated based on the reward of the generated configurations (line 6 and 8); i.e., the particles that induce higher reward values will get higher weights. If only a few samples have weights effectively, i.e., if the weights collapse, the algorithm resamples the particles from the current approximate posterior distribution to maintain the effective sample size (line 9). After the forward recursion over the planning horizon, the optimal commands are computed as a linear combination of the initial disturbances; i.e., it is given by the expected disturbance under the posterior transition dynamics (Kappen & Ruiz, 2016).

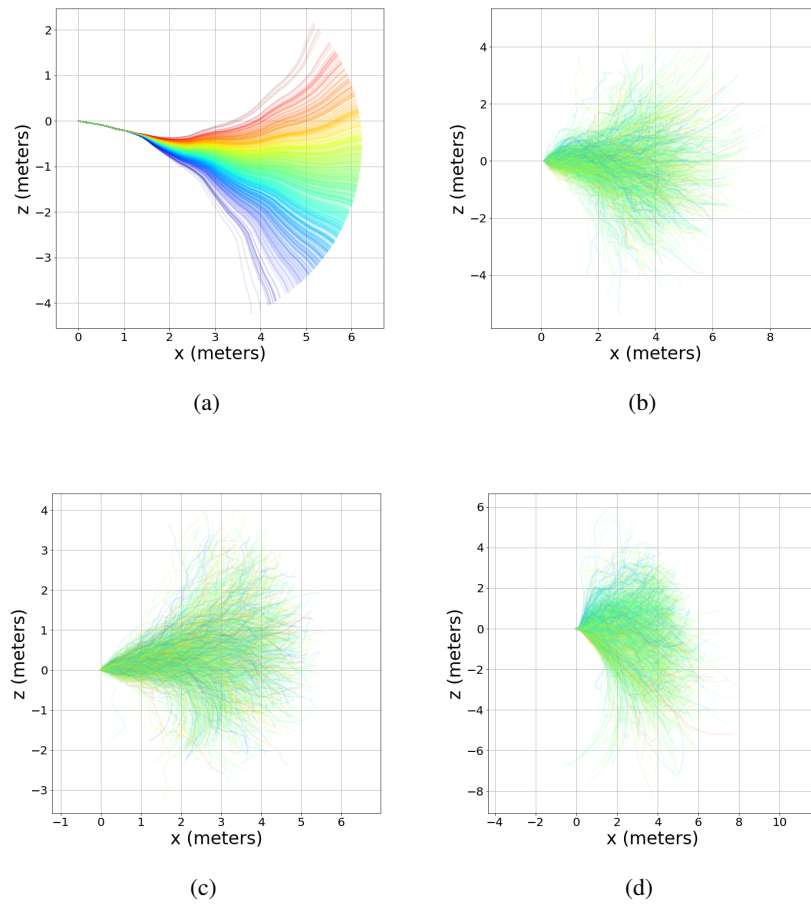


Figure 5: Rollouts from the different internal models. (a) DISH (ours) (b) sas' (c) shs' (d) zaz'