Distilling a Hierarchical Policy for Planning and Control via Representation and Reinforcement Learning

Jung-Su Ha[*1,2], Young-Jin Park[*3], Hyeok-Joo Chae[4], Soon-Seo Park[4] and Han-Lim Choi[4]

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 representation 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 an optimistic future and (ii) a feedback control policy, shared across the tasks, that *executes* the inferred intention. Because the planning 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- and 1-dimensional latent states and commands for a humanoid with 197- and 36-dimensional state features and actions) 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.
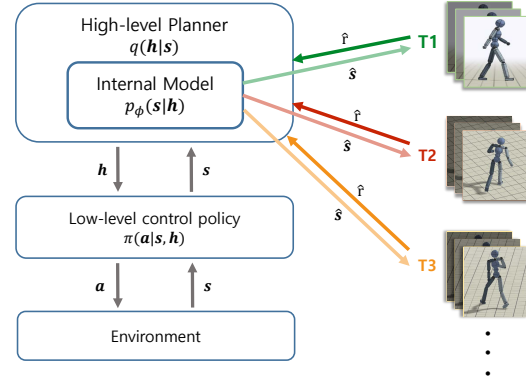
Fig. 1. The proposed DISH framework. The low-level control policy learned via RL maps high-level commands into actions, while the high-level planner reasons task-specific commands using the internal model learned via self-supervised learning.

I. 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, with which some of the deep RL agents have already achieved or even exceeded human-level performances in particular tasks [1], [2]. 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 of the RL algorithms makes it unrealistic to learn an individual policy for each task, an agent should be able to *reason* its action instead of *memorizing* the optimal behavior. This would be possible if predictions about consequences of actions are available, e.g., by using an internal model [3], [4]. Involving planning procedures in a control policy could provide adaptiveness to an agent, but learning such a prediction & planning framework is often not trivial: 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, which is accumulated along the prediction/planning horizon. This prohibits planning methods from producing a valid prediction and, as a result, 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 [5].

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 [6], [7]. 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, and then in a finer time scale, the (task-agnostic) 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) [6]. 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 environment 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 planning 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 plans/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

arXiv:2011.08345v2 [cs.LG] 6 Apr 2021

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 [8]–[10]. 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.

II. 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 controlller, i.e., as a "skill", and a high-level controller selects which skills to perform in the context the agent lies at [11]–[14]. [11], [12] 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. [13] similarly trained many imitation policies and utilized them as individual skills that a high-level controller chooses based on the visual inputs. [14] 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 [15]–[20]. [15] 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. [16]–[18] also introduced latent variables to efficiently represent various policies. Instead of using one *static* latent variable, [21] proposed a framework that encodes expert's demonstrations as latent *trajectories* and infers a latent trajectory from an unseen skill for one-shot imitation. [22] 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 planning; 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. Only a few recent works [23]–[26] incorporated reasoning processes into high-level modules, but most of those works did not exploit low-dimensional latent space for planning nor low-dimensional commands. Our ablation study in Section IV-A shows the effectiveness of utilizing both latent states and commands and, to our best knowledge, DISH is the first framework doing so.

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 [27]–[29]. [27], [29] utilized deep neural networks to model the dynamics and adopted the model predictive control method on the learned dynamics; [28] 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 [3], 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) planning (or reasoning) by predicting the future is essential to solve a set of complex tasks. On the other hand, [30], [31] 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 computations of planning and optimal control methods: [32] encoded the value iteration procedures into the network; [33], [34] 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.

III. DISH: DISTILLING HIERARCHY FOR PLANNING AND CONTROL

A. 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:

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

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

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, $p(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}) \equiv p(\mathbf{s}_1) \prod_{k=1}^{K} p(\mathbf{s}_{k+1}|\mathbf{s}_k, \mathbf{a}_k)p(\mathbf{a}_k)$, the above RL problem can be viewed 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 B or in the tutorial paper [8].

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 $\pi_{\theta_t^*}(\mathbf{a}_k|\mathbf{s}_k)$ be the optimal policy for t^{th} task, i.e.,

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

For high DOF systems, where policies π_{θ_t} represent a mapping from a high-dimensional state space to a high-dimensional action space, individually optimizing each policy is computationally too expensive. Instead of doing so, we can assume that tasks the agent needs to perform require similar solution properties, making the optimal policies possess common structures. We can then introduce a low-dimensional latent variable $\mathbf{h}^{(t)}$ that compresses a particular aspect of π_{θ_t} over all the policies and that each policy can be conditioned on as $\pi_\theta(\mathbf{a}_k|\mathbf{s}_k, \mathbf{h}^{(t)})$.

Fig. 2(b) depicts such a hierarchical structure, where $\mathbf{h}$ can be interpreted as a high-level *command*. Following the aforementioned duality, the uncontrolled and the task t's controlled trajectory distributions are defined as

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

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

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 the other for the task-specific commands $\mathbf{h}$. Note that, if high-level commands are computed via another explicit policy function $\bar{\pi}_\theta(\mathbf{h}|\mathbf{s})$ e.g. a neural network, the overall learning problem then becomes the standard Hierarchical RL (HRL). We instead introduce a planning module to generate high-level commands which infers the optimal $\mathbf{h}$ for a given task t by predicting futures. As often used in many HRL methods, the high-level module of the proposed framework operates in a coarser time scale than the low-level policy does.

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

$$\log p_\theta(O_{1:K}^{(t)} = 1) = \log \int p(O_{1:K}^{(t)} = 1|\mathbf{s}_{1:K})p(\tau)\frac{q_\theta^{(t)}(\tau)}{q_\theta^{(t)}(\tau)}d\tau$$

$$\geq \mathbb{E}_{q_\theta^{(t)}(\tau)} \left[\sum_{k=1}^{K} r_k^{(t)}(\mathbf{s}_k) - \log \frac{\pi_\theta(\mathbf{a}_k|\mathbf{s}_k, \mathbf{h}_k)}{p(\mathbf{a}_k)} \frac{q^{(t)}(\mathbf{h}_k|\mathbf{s}_k)}{p(\mathbf{h}_k)} \right]$$

$$\equiv \mathcal{L}^{(t)}(\theta, q), \quad (5)$$

where $\tau \equiv (\mathbf{s}_{1:K}, \mathbf{a}_{1:K}, \mathbf{h}_{1:K})$. Maximization of this lower bound suggests a novel learning scheme of the hierarchical policy in (4). **(i)** *Maximization w.r.t. q*: For a given task t and a fixed low-level policy π_θ, 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 representation learning with which planning is performed. **(ii)** *Maximization w.r.t. θ*: With the planning module equipped, the low-level policy $\pi_\theta(\mathbf{a}|\mathbf{s}, \mathbf{h})$ generates control actions $\mathbf{a}$ as in RL problems and is trained using standard deep RL algorithms [35], [36].

B. Self-supervised Learning of Internal Model

The role of $q(\mathbf{h}|\mathbf{s})$ is to compute the high-level commands that will lead to maximum accumulated rewards in the future; as shown in (5), this infers the commands that maximizes the likelihood of optimality variables when $O_{1:K} = 1$ were observed. Since the ELBO gap is the KL-divergence between the posterior and variational distributions, more exact variational inference will make the lower bound tighter, thereby directly leading 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); such 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 current step, $\mathbf{h}_1$, the inference procedure should compute the posterior command trajectories $q^*(\mathbf{h}_{1:K})$ by considering the dynamics

and observations (the optimality variables), and marginalize the future commands $\mathbf{h}_{2:K}$ out. Though the dimensionality of $\mathbf{h}$ is much lower than that of $\mathbf{a}$, this inference problem is still not trivial to solve by two reasons: (i) The dynamics of states $p_\theta(\mathbf{s}'|\mathbf{s}, \mathbf{h}) = \int p(\mathbf{s}'|\mathbf{s}, \mathbf{a})\pi_\theta(\mathbf{a}|\mathbf{s}, \mathbf{h})d\mathbf{a}$ 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 $p_\phi(\mathbf{s}'|\mathbf{s}, \mathbf{h})$ via supervised learning with transition data obtained during low-level policy learning. However, learning a 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 [30].

Based our prior work [31], we build a low-dimensional latent variable model and use it as an internal model for planning. Our framework collects the trajectories from low-level policies and utilize them to learn a LVM for inference, which is formulated as a maximum likelihood estimation (MLE) problem. Suppose that we have collected a set of state trajectories and latent commands $\{\mathbf{s}_{1:K}^{(n)}, \mathbf{h}_{1:K}^{(n)}\}_{n=1,...,N}$. We then formulate the MLE problem as:

$$\phi^* = \arg\max_\phi \sum_n \log p_\phi(\mathbf{s}_{1:K}^{(n)}|\mathbf{h}_{1:K}^{(n)}). \tag{6}$$

As in Fig. 2(d), the states are assumed to emerge from a latent dynamical system, where a latent state trajectory, $\mathbf{z}_{1:K}$, lies on a low-dimensional latent space $\mathcal{Z}$:

$$p_\phi(\mathbf{s}_{1:K}|\mathbf{h}_{1:K}) = \int p_\phi(\mathbf{s}_{1:K}|\mathbf{z}_{1:K})p_\phi(\mathbf{z}_{1:K}|\mathbf{h}_{1:K})d\mathbf{z}_{1:K}. \tag{7}$$

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

$$\mathbf{z}_{k+1} = f_\phi(\mathbf{z}_k) + \sigma_\phi(\mathbf{z}_k)(\mathbf{h}_k + \mathbf{w}_k), \ \mathbf{w}_k \sim \mathcal{N}(0, I) \tag{8}$$

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 has a form of self-supervised learning problems that have been extensively studied recently [31], [37]–[39]. In particular, we adopt the idea of Adaptive path-integral autoencoder in [31], 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:

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

Compared to the original formulation in [31], the probability model here is conditioned on the commands, $\mathbf{h}_{1:K}$, making the learning problem conditional generative model learning [40].[1]

C. Planning with Learned Internal Model

Once the LVM is trained, a planning module can efficiently explore the state space $\mathcal{S}$ through the latent state $\mathbf{z}$ and

[1]Effectively it only shifts the control input prior from $\mathcal{N}(\mathbf{0}, I)$ to $\mathcal{N}(\mathbf{h}, I)$ as written in (8) and (9) [27].

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 which is known to perform well with non-linear and non-Gaussian systems [30], [41]. The detailed

---

Algorithm 1 PF for Planning with Internal Model

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

---

procedure is shown in Algorithm 1. At each time step, the high-level planner takes the current state as an argument and outputs the commands by predicting the future trajectory and corresponding reward $r_k(\cdot)$. 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 4–5), and the corresponding configurations are generated through the learned model (line 6). The weights of all particles are then updated based on the reward of the generated configurations (line 7 and 9) such that the particles that induce higher reward values get higher weights. If only a few samples have weights effectively, the algorithm resamples the particles from the current approximate posterior distribution to maintain the effective sample size (line 10). After the forward recursion over the planning horizon, the optimal commands are computed as a linear combination of the initial disturbances; that is, it is given by the expected disturbance under the posterior transition dynamics [42].

In the perspective of this work, this procedure can be viewed as the agent simulating multiple future state trajectories with the internal model, assigning each of them according to the reward, and planning/reasoning the command that leads to the best-possible future. Note that if we iterate the whole procedure multiple times to improve the commands, the algorithm becomes equivalent to the adaptive path integral method [27], [31], [42]. If the resampling procedure is eliminated, this method reduces to the widely-used cross entropy method [24].

D. Learning Algorithm

The overall learning procedure is summarized in Algorithm 2. It consists of an outer internal model learning loop and an inner policy update loop. During the policy update stage (inner loop), the algorithm samples a task,

Algorithm 2 DIStilling Hierarchy for Planning and Control

1: Initialize policy θ and latent model ϕ
2: **for** $l = 1, ..., L$ **do**
3:     **while** not converged **do**
4:         Sample a task $t \in \mathcal{T}$
5:         Run the policy $\mathbf{a} \sim \pi_\theta(\mathbf{a}|\mathbf{s}, \mathbf{h})$, $\mathbf{h} \sim q_\phi(\mathbf{h}|\mathbf{s})$
6:         Store trajectories τ into the experience buffer
7:         Train the policy π_θ using e.g. PPO $\triangleright$ Eq. (5)
8:     **end while**
9:     Random sample $\mathbf{h}$ and collect rollouts
10:     Train the internal model using e.g. APIAE $\triangleright$ Eq. (6)
11: **end for**
---

executes the action using the hierarchical policy, and collects trajectories into the experience buffer. At each time step, the low-level policy decides actions the agent takes under the high-level commands determined by the planning module. Using transition data in the buffer, the low-level policy is updated via a deep RL algorithm (e.g., policy gradient methods). After the low-level policy update, DISH collects another set of rollouts by random sampling a latent variable $\mathbf{h}$, and the internal model is learned via self-supervised representation learning. These two learning procedures are then iterated for L times.

Note that, for complex systems, tasks can be selected more carefully (at line 4) for a more stable learning landscape; for example, in earlier phases where the agent couldn't yet learn a valid policy and/or an internal model, the agent can first learn them through imiation learning of expert's demonstrations [11] or play data [43], or through intrinsic motivations to acquire useful skills [23]. As more challenging tasks are gradually provided to the agent, the internal model will be learned to cover wider ranges of state space for those tasks and the low-level policy will be trained such that it can execute more complicated high-level commands.

IV. EXPERIMENT

In this section, we demonstrate the effectiveness of the proposed framework on performing planning and control for the high dimensional humanoid robot [11] which has 197 state features and 36 action parameters, simulated by 1.2kHz Bullet physics engine [44]. The low-level control policy and the internal latent model are trained through the imitation learning, where three locomotion data from the Carnegie Mellon University motion capture (CMU mocap) database are used as target motions of imitation. The control policy is trained with the DeepMimic imitation reward [11] by using proximal policy optimization (PPO) [35], while the internal model is learned to maximize the likelihood of experience data (i.e. (6)) by using the APIAE approach [31]. The internal model of DISH is constructed to have a 3-dimensional latent state and a 1-dimensional latent command for all experiments. The low-level policy and the internal model are operated in different time scales, 30Hz and 1Hz, respectively. The learned hierarchical model is then evaluated on trajectory following and navigation tasks in Section IV-A and IV-B, respectively. For planning and execution, a 5-second trajectory is planned

(a) DISH   (b) zaz' (VAE)
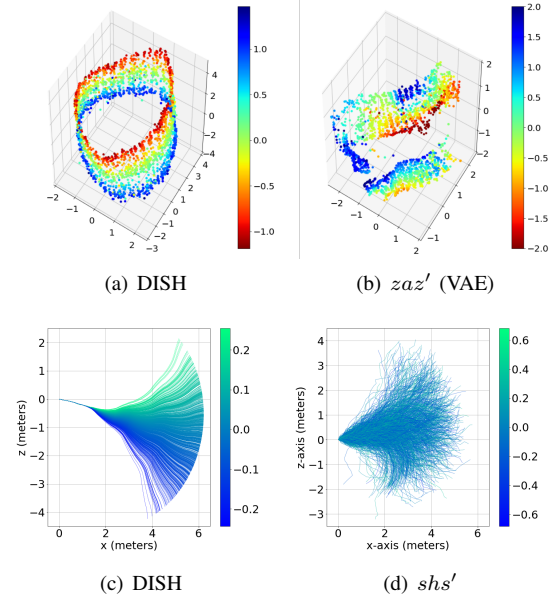
(c) DISH   (d) shs'

Fig. 3. (a), (b) Learned latent models colored by angular velocity. (c), (d) Rollout samples in horizontal (x-z plane) colored by latent command value.

and only the first high-level command is applied to the policy at 1Hz and 4Hz for each task.

We refer to the appendix for the reward functions, hyper-parameters, and network architectures (Appendix D and E), task configurations (Appendix F), and more experimental results (Appendix H). Our TensorFlow [45] implementation are available online[2]. The supplementary video also summarizes the training procedure and visualizes the resulting policy[3].

A. Ablation Study: Learning Hierarchical Structure

In the first experiment, we examine how effectively the proposed framework learns and exploits the internal model. To investigate the effectiveness of each component introduced, we conduct ablation studies by considering three baselines: (i) sas' that does not have neither the hierarchical structure nor LVMs (Fig. 2(a)), (ii) shs' that utilizes the hierarchical policy but doesn't learn the low-dimensional latent dynamics (Fig. 2(c)), and (iii) zaz' that considers the latent dynamics but without the hierarchical structure (no latent commands, a LVM version of Fig. 2(a); Fig.2(d) depicts a LVM version of Fig. 2(c).). Given the rollouts $\{\tau^{(i)}\} = \{\mathbf{s}_{1:K}^{(i)}, \mathbf{a}_{1:K}^{(i)}, \mathbf{h}_{1:K}^{(i)}\}$, sas' and shs' are simply supervised learning problems. For the zaz' model, the variational autoencoder (VAE) approach [46] is taken to train mappings between the observation and the latent space, and then the latent dynamics is trained via supervised learning, following the idea of [3]. All models including DISH are trained using the same rollouts for the fair comparison. Note that most HRL frameworks can be categorized as either zaz' e.g., [3], [24] or shs' e.g., [23]. The similar network structures are used for the baselines (See Appendix E). The first three columns of Table I summarizes the different features of the models with the related works.

Qualitatively, Figs. 3(a) and 3(b) show the learned latent space colored by the moving-averaged angular velocity of

TABLE I

COMPARISON BETWEEN DIFFERENT TYPES OF INTERNAL MODELS AND QUANTITATIVE COMPARISON FOR TRAJECTORY FOLLOWING TASKS. 'F'
DENOTES THAT IT WAS NOT ABLE TO RECORD THE TRUE TRAJECTORY SINCE THE AGENT KEPT FALLING.

	command, $\mathbf{h}$	LVM, $\mathbf{z}$	Reconstruction	‖ref-true‖	‖plan-ref‖	‖plan-true‖
DISH (ours, L=1, Fig. 1 & Fig.2(d))	✓	✓	0.3820	0.1638	0.1576	**0.0930**
DISH+ (ours, L=2)			0.8414	**0.1452**	**0.1509**	0.1105
sas' (w/o command & LVM, Fig. 2(a))	✗	✗	**0.1289**	F	0.1771	F
shs' (w/o LVM, Fig. 2(c))	✓	✗	0.1393	0.2226	0.1579	0.2231
zaz' (w/o command)	✗	✓	2.3351	F	0.2589	F

the ground truth motion. In the case of DISH, the latent state forms a manifold of a cylindrical shape in 3-dimensional space where the locomotion phase and the angular velocity are well encoded along the manifold. In contrast, the latent state structure of the zaz' model does not capture the phase information and failed to construct a periodic manifold, which prevents a valid latent dynamics from being learned. Figs. 3(c) and 3(d) show the rollout trajectories from each internal model colored by the values of high-level commands, $\mathbf{h}$. The high-level commands of DISH are learned to control the heading direction of the humanoid so that the agent can make the structural exploration in the configuration space. The shs' model, on the other hand, fails to learn a valid controlled dynamics (since its space is too large) and consequently just generates noisy trajectories.

To quantitatively evaluate the planning performance of DISH and its ability to flexibly perform different tasks, we compare DISH to the baseline models on three trajectory following tasks: going straight, turning left and right. Table I reports the RMS errors for reconstruction and differences between the reference, planned, and executed trajectories. There are three things we can observe from the table: (i) Although sas' has the lowest reconstruction error, the computed action from its internal model even cannot make the humanoid walk. This is because the humanoid has a highly unstable dynamics and reasoning of the high-dimensional action is not accurate enough to stabilize the humanoid dynamics, i.e., searching over the 36-dimensional action space with the limited number of particles (1024 in this case) is not feasible. For the same reason, zaz' also fails to let the humanoid walk. (ii) Only the models considering the hierarchical policy structure can make the humanoid walk, and the DISH framework generates the most executable and valuable plans; the humanoid with the shs' model walks just in random directions rather than following a planned trajectory (see Fig. 3(d)), which implies that the high-level command $\mathbf{h}$ does not provide any useful information regarding the task. (iii) By iterating the low-level policy and the internal model learning further, DISH+ becomes able to reason better plans as well as execute them better.

B. Planning and Control with Learned Hierarchy

In the second experiment, we further demonstrate the capability of DISH framework to perform navigation tasks in cluttered environments (shown in Fig. 4). Since the humanoid with the baseline models either kept falling or failed to walk in a desired direction, we omit the comparisons in this task. The navigation reward is designed as a sum of
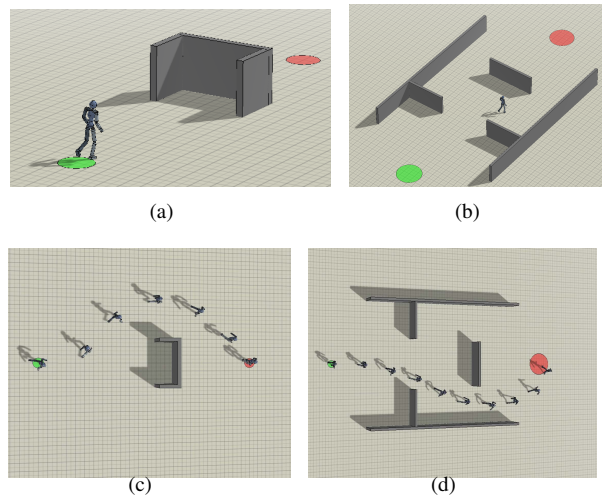

Fig. 4. Cluttered environments for navigation tasks.

two components: penalty for distance from the goal and penalty for collision with obstacles. As shown in Figs. 4(c) and 4(d) as well as in the supplementary video, the humanoid equipped with the DISH policy is able to not only escape a bug trap which cannot be overcome with greedy algorithms (i.e. without planning), but also navigate through obstacle regions successfully. Unlike the HRL algorithms, the proposed hierarchical policy trained from the imitation tasks can be directly applied to the navigation tasks. It shows the generalization power of planning process; utilizing the internal model and the command-conditioned policy enables the agent to directly adapt to changing tasks and environments. On the other hand, to address more challenging problems which require more complex motions, the *curriculum* would be designed more carefully, as discussed in Section III-D.

V. CONCLUSION

We proposed a framework to learn a hierarchical policy for an RL agent, where 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. 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 trained using a deep RL algorithm. By alternately optimizing both the LVM and the policy, the framework was able to construct a meaningful internal model as well as a versatile control policy.

REFERENCES

[1] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, p. 529, 2015.

[2] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, *et al.*, "Mastering the game of go without human knowledge," *Nature*, vol. 550, no. 7676, p. 354, 2017.

[3] D. Ha and J. Schmidhuber, "Recurrent world models facilitate policy evolution," in *Advances in Neural Information Processing Systems*, 2018, pp. 2455–2467.

[4] L. Kaiser, M. Babaeizadeh, P. Milos, B. Osinski, R. H. Campbell, K. Czechowski, D. Erhan, C. Finn, P. Kozakowski, S. Levine, *et al.*, "Model-based reinforcement learning for atari," *arXiv preprint arXiv:1903.00374*, 2019.

[5] S. M. LaValle, *Planning algorithms*. Cambridge university press, 2006.

[6] E. Todorov and Z. 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*, vol. 2. IEEE, 2003, pp. 1750–1753.

[7] E. Todorov, "Optimality principles in sensorimotor control," *Nature neuroscience*, vol. 7, no. 9, p. 907, 2004.

[8] S. Levine, "Reinforcement learning and control as probabilistic inference: Tutorial and review," *arXiv preprint arXiv:1805.00909*, 2018.

[9] E. Todorov, "General duality between optimal control and estimation," in *IEEE Conference on Decision and Control*. IEEE, 2008, pp. 4286–4292.

[10] K. Rawlik, M. Toussaint, and S. Vijayakumar, "On stochastic optimal control and reinforcement learning by approximate inference," in *Robotics: Science and Systems*, 2012.

[11] X. B. Peng, P. Abbeel, S. Levine, and M. van de Panne, "DeepMimic: Example-guided deep reinforcement learning of physics-based character skills," *ACM Trans. Graph.*, vol. 37, no. 4, pp. 143:1–143:14, July 2018. [Online]. Available: http://doi.acm.org/10.1145/3197517.3201311

[12] X. B. Peng, M. Chang, G. Zhang, P. Abbeel, and S. Levine, "Mcp: Learning composable hierarchical control with multiplicative compositional policies," in *Advances in Neural Information Processing Systems*, 2019.

[13] J. Merel, A. Ahuja, V. Pham, S. Tunyasuvunakool, S. Liu, D. Tirumala, N. Heess, and G. Wayne, "Hierarchical visuomotor control of humanoids," in *International Conference on Learning Representations*, 2019.

[14] Y. Lee, S.-H. Sun, S. Somasundaram, E. Hu, and J. J. Lim, "Composing complex skills by learning transition policies with proximity reward induction," in *International Conference on Learning Representations*, 2019.

[15] J. D. Co-Reyes, Y. Liu, A. Gupta, B. Eysenbach, P. Abbeel, and S. Levine, "Self-consistent trajectory autoencoder: Hierarchical reinforcement learning with trajectory embeddings," in *International Conference on Machine Learning*, 2018.

[16] A. Gupta, R. Mendonca, Y. Liu, P. Abbeel, and S. Levine, "Meta-reinforcement learning of structured exploration strategies," in *Advances in Neural Information Processing Systems*, 2018.

[17] B. Eysenbach, A. Gupta, J. Ibarz, and S. Levine, "Diversity is all you need: Learning skills without a reward function," in *International Conference on Learning Representations*, 2019.

[18] C. Florensa, Y. Duan, and P. Abbeel, "Stochastic neural networks for hierarchical reinforcement learning," in *International Conference on Learning Representations*, 2017.

[19] K. Hausman, J. T. Springenberg, Z. Wang, N. Heess, and M. Riedmiller, "Learning an embedding space for transferable robot skills," in *International Conference on Learning Representations*, 2018.

[20] O. Nachum, S. S. Gu, H. Lee, and S. Levine, "Data-efficient hierarchical reinforcement learning," in *Advances in Neural Information Processing Systems*, 2018, pp. 3303–3313.

[21] J. Merel, L. Hasenclever, A. Galashov, A. Ahuja, V. Pham, G. Wayne, Y. W. Teh, and N. Heess, "Neural probabilistic motor primitives for humanoid control," in *International Conference on Learning Representations*, 2019.

[22] T. Haarnoja, K. Hartikainen, P. Abbeel, and S. Levine, "Latent space policies for hierarchical reinforcement learning," in *International Conference on Machine Learning*, 2018.

[23] A. Sharma, S. Gu, S. Levine, V. Kumar, and K. Hausman, "Dynamics-aware unsupervised discovery of skills," *arXiv preprint arXiv:1907.01657*, 2019.

[24] D. Hafner, T. P. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson, "Learning latent dynamics for planning from pixels," in *ICML*, 2019, pp. 2555–2565. [Online]. Available: http://proceedings.mlr.press/v97/hafner19a.html

[25] S. Nasiriany, V. Pong, S. Lin, and S. Levine, "Planning with goal-conditioned policies," 2019.

[26] K. Liu, T. Kurutach, C. Tung, P. Abbeel, and A. Tamar, "Hallucinative topological memory for zero-shot visual planning," in *International Conference on Machine Learning*. PMLR, 2020, pp. 6259–6270.

[27] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, "Information theoretic MPC for model-based reinforcement learning," in *International Conference on Robotics and Automation (ICRA)*, 2017.

[28] M. P. Deisenroth, D. Fox, and C. E. Rasmussen, "Gaussian processes for data-efficient learning in robotics and control," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 37, no. 2, pp. 408–423, 2015.

[29] K. Chua, R. Calandra, R. McAllister, and S. Levine, "Deep reinforcement learning in a handful of trials using probabilistic dynamics models," *Advances in Neural Information Processing Systems*, 2018.

[30] J.-S. Ha, H.-J. Chae, and H.-L. Choi, "Approximate inference-based motion planning by learning and exploiting low-dimensional latent variable models," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3892–3899, 2018.

[31] J.-S. Ha, Y.-J. Park, H.-J. Chae, S.-S. Park, and H.-L. Choi, "Adaptive path-integral autoencoders: Representation learning and planning for dynamical systems," in *Advances in Neural Information Processing Systems*, 2018, pp. 8941–8952.

[32] A. Tamar, Y. Wu, G. Thomas, S. Levine, and P. Abbeel, "Value iteration networks," in *Advances in Neural Information Processing Systems*, 2016, pp. 2154–2162.

[33] M. Okada, L. Rigazio, and T. Aoshima, "Path integral networks: End-to-end differentiable optimal control," *arXiv preprint arXiv:1706.09597*, 2017.

[34] B. Amos, I. Jimenez, J. Sacks, B. Boots, and J. Z. Kolter, "Differentiable MPC for end-to-end planning and control," in *Advances in Neural Information Processing Systems*, 2018, pp. 8299–8310.

[35] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.

[36] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *International Conference on Machine Learning*, 2018.

[37] M. Karl, M. Soelch, J. Bayer, and P. van der Smagt, "Deep variational bayes filters: Unsupervised learning of state space models from raw data," *International Conference on Learning Representations (ICLR)*, 2017.

[38] R. G. Krishnan, U. Shalit, and D. Sontag, "Structured inference networks for nonlinear state space models." in *AAAI Conference on Artificial Intelligence*, 2017, pp. 2101–2109.

[39] M. Fraccaro, S. Kamronn, U. Paquet, and O. Winther, "A disentangled recognition and nonlinear dynamics model for unsupervised learning," in *Advances in Neural Information Processing Systems (NIPS)*, 2017, pp. 3604–3613.

[40] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," in *Advances in neural information processing systems*, 2015, pp. 3483–3491.

[41] A. Piche, V. Thomas, C. Ibrahim, Y. Bengio, and C. Pal, "Probabilistic planning with sequential monte carlo methods," in *International Conference on Learning Representations*, 2019.

[42] H. J. Kappen and H. C. Ruiz, "Adaptive importance sampling for control and inference," *Journal of Statistical Physics*, vol. 162, no. 5, pp. 1244–1266, 2016.

[43] C. Lynch, M. Khansari, T. Xiao, V. Kumar, J. Tompson, S. Levine, and P. Sermanet, "Learning latent plans from play," *arXiv preprint arXiv:1903.01973*, 2019.

[44] E. Coumans *et al.*, "Bullet physics library," *Open source: bulletphysics. org*, vol. 15, no. 49, p. 5, 2013.

[45] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals,

P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," 2015, software available from tensorflow.org. [Online]. Available: https://www.tensorflow.org/

[46] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," *arXiv preprint arXiv:1312.6114*, 2013.

[47] M. Toussaint, "Robot trajectory optimization using approximate inference," in *Proceedings of the 26th annual international conference on machine learning*. ACM, 2009, pp. 1049–1056.

[48] M. Mukadam, J. Dong, X. Yan, F. Dellaert, and B. Boots, "Continuous-time gaussian process motion planning via probabilistic inference," *The International Journal of Robotics Research*, vol. 37, no. 11, pp. 1319–1340, 2018.

[49] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning." in *AAAI*, vol. 8. Chicago, IL, USA, 2008, pp. 1433–1438.

A. Control-Inference Duality

One theoretical concept this work extensively takes advantage of is the duality between optimal control (OC) and probabilistic inference [8]–[10]. 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 [47], the particle filtering [30], [41], or the inference for Gaussian processes [48]. 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.

B. Reinforcement Learning as Probabilistic Inference

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

$$\theta^* = \operatorname*{argmax}_{\theta} \mathbb{E}_{q_\theta(\mathbf{s}_{1:K}, \mathbf{a}_{1:K})} \left[\sum_{k=1}^{K} \tilde{r}_k(\mathbf{s}_k, \mathbf{a}_k) \right], \tag{10}$$

$$q_\theta(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}) \equiv p(\mathbf{s}_1) \prod_{k=1}^{K} p(\mathbf{s}_{k+1}|\mathbf{s}_k, \mathbf{a}_k) \pi_\theta(\mathbf{a}_k|\mathbf{s}_k), \tag{11}$$

respectively. It is well known in the literature that the above optimization ((10)) also can be viewed as a probabilistic inference problem for a certain type of graphical models [8]–[10]. 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.,

$$p(o_k = 1|\mathbf{s}_k) = \exp\left(r_k(\mathbf{s}_k)\right), \tag{12}$$

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

$$p(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}) \equiv p(\mathbf{s}_1) \prod_{k=1}^{K} p(\mathbf{s}_{k+1}|\mathbf{s}_k, \mathbf{a}_k) p(\mathbf{a}_k). \tag{13}$$

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

$$\begin{aligned}
\log p(O_{1:K}) &= \log \int p(O_{1:K}|\mathbf{s}_{1:K}) p(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}) d\mathbf{s}_{1:K} d\mathbf{a}_{1:K} \\
&= \log \int p(O_{1:K}|\mathbf{s}_{1:K}) p(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}) \frac{q_\theta(\mathbf{s}_{1:K}, \mathbf{a}_{1:K})}{q_\theta(\mathbf{s}_{1:K}, \mathbf{a}_{1:K})} d\mathbf{s}_{1:K} d\mathbf{a}_{1:K} \\
&\geq \mathbb{E}_{q_\theta(\mathbf{s}_{1:K}, \mathbf{a}_{1:K})} \left[\sum_{K=1}^{K} \left(\log p(O_k|\mathbf{s}_k) - \log \frac{\pi_\theta(\mathbf{a}_k|\mathbf{s}_k)}{p(\mathbf{a}_k)} \right) \right] \\
&= \mathbb{E}_{q_\theta(\mathbf{s}_{1:K}, \mathbf{a}_{1:K})} \left[\sum_{k=1}^{K} r_k(\mathbf{s}_k) - \log \frac{\pi_\theta(\mathbf{a}_k|\mathbf{s}_k)}{p(\mathbf{a}_k)} \right] \\
&\equiv \mathcal{L}(\theta). \tag{14}
\end{aligned}$$

The ELBO maximization in (14) becomes equivalent to the reinforcement learning in (10) by choosing an action prior $p(\mathbf{a}_k)$ and parameterized policy family $\pi_\theta(\mathbf{a}_k|\mathbf{s}_k)$ to match $\tilde{r}_k = r_k - \log \frac{\pi_\theta}{p}$[4]. Similar to (18), 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 $p(\mathbf{s}_{1:K}, \mathbf{a}_{1:K}|O_{1:K} = 1)$.

---

[4]For 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 [36], [49].

C. Self-supervised Learning of Latent Dynamical Models

Self-supervised representation learning is an essential approach that allows an agent to learn underlying dynamics only 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 [31], [37]–[39].

Suppose that a set of observation sequences $\{\mathbf{s}_{1:K}^{(n)}\}_{n=1,...,N}$ is given, where $\mathbf{s}_{1:K}^{(n)} \equiv \{\mathbf{s}_k; \forall k = 1, ..., K\}^{(n)}$ 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:

$$\phi^* = \underset{\phi}{\operatorname{argmax}} \sum_n \log p_\phi(\mathbf{s}_{1:K}^{(n)}). \tag{15}$$

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

$$p_\phi(\mathbf{s}_{1:K}) = \int p_\phi(\mathbf{s}_{1:K}|\mathbf{z}_{1:K})p_\phi(\mathbf{z}_{1:K})d\mathbf{z}_{1:K}, \tag{16}$$

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 ((15)) 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:

$$\log p_\phi(\mathbf{s}_{1:K}) = \log \int p_\phi(\mathbf{s}_{1:K}|\mathbf{z}_{1:K})p_\phi(\mathbf{z}_{1:K})d\mathbf{z}_{1:K}$$
$$\geq \mathbb{E}_{q(\mathbf{z}_{1:K})}\left[\log \frac{p_\phi(\mathbf{s}_{1:K}|\mathbf{z}_{1:K})p_\phi(\mathbf{z}_{1:K})}{q(\mathbf{z}_{1:K})}\right]$$
$$\equiv \mathcal{L}(q, \phi; \mathbf{s}_{1:K}), \tag{17}$$

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 $p_\theta(\mathbf{z}_{1:K}|\mathbf{s}_{1:K})$:

$$\log p_\phi(\mathbf{s}_{1:K}) - \mathcal{L}(q, \phi; \mathbf{s}_{1:K}) = D_{KL}(q(\mathbf{z}_{1:K})||p_\phi(\mathbf{z}_{1:K}|\mathbf{s}_{1:K})). \tag{18}$$

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 [31], [37]–[39]. In p articular, [31] 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:

$$\mathbf{z}_{k+1} = f_\phi(\mathbf{z}_k) + \sigma_\phi(\mathbf{z}_k)\mathbf{w}_k, \ \mathbf{z}_0 \sim p_0(\cdot), \ \mathbf{w}_k \sim \mathcal{N}(0, I). \tag{19}$$

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

$$p_\phi(\mathbf{s}_{1:K}|\mathbf{z}_{1:K}) = \prod_{k=1}^{K} p_\phi(\mathbf{s}_k|\mathbf{z}_k). \tag{20}$$

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

$$\mathbf{z}_{k+1} = f_\phi(\mathbf{z}_k) + \sigma_\phi(\mathbf{z}_k)(\mathbf{u}_k + \mathbf{w}_k), \ \mathbf{z}_0 \sim q_0(\cdot), \ \mathbf{w}_k \sim \mathcal{N}(0, I), \tag{21}$$

the ELBO can be written in the following form:

$$\mathcal{L} = \mathbb{E}_{q_\mathbf{u}}\left[\log p_\phi(\mathbf{s}_{1:K}|\mathbf{z}_{1:K}) + \log \frac{p_0(\mathbf{z}_0)}{q_0(\mathbf{z}_0)} - \sum_{k=1}^{K-1} \frac{1}{2}\|\mathbf{u}_k\|^2 + \mathbf{u}_k^\top \mathbf{w}_k\right]. \tag{22}$$

(a) low-level policy network   (b) transition network   (c) generative network
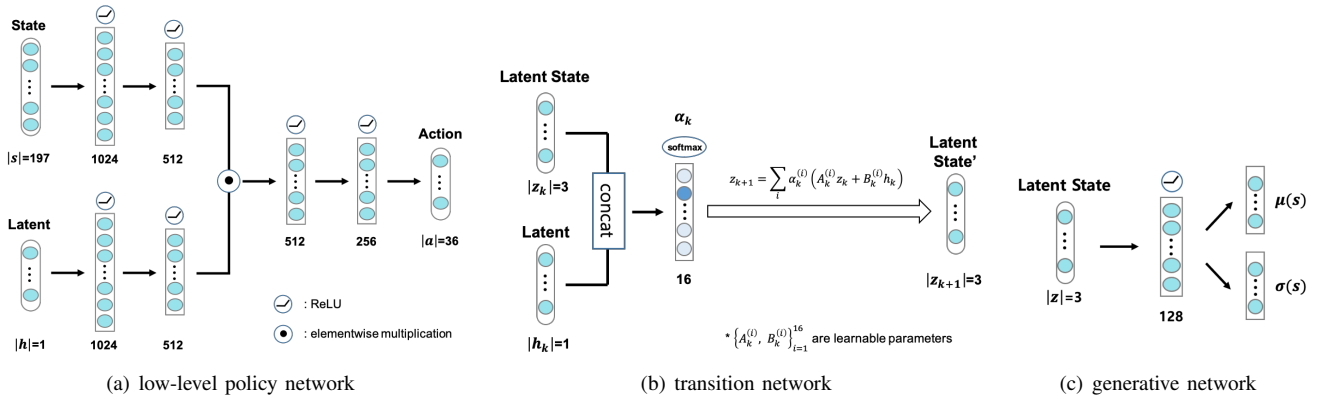
Fig. 5.   The network architectures of policy and internal model.

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:

$$\mathbf{u}^*, q_0^* = \underset{\mathbf{u}, q_0}{\operatorname{argmin}} \, \mathbb{E}_{q_{\mathbf{u}}(\mathbf{z}_{1:K})} \left[V(\mathbf{z}_{1:K}) + \sum_{k=1}^{K-1} \frac{1}{2} \|\mathbf{u}_k\|^2 + \mathbf{u}_k^\top \mathbf{w}_k \right], \tag{23}$$

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. [31] 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.

D. Training Low-level Policy

For the training algorithm for low-level policy network (π_θ), we extend motion imitation approach [11] to multi-task scheme; we construct value networks parameterized by neural network with size $[197, 1024, 512, 1]$ for each task (three in our experiments), and the low-level policy network (actor network) taking a state feature $\mathbf{s}$ and a latent variable $\mathbf{h}$ as inputs to determine an action $\mathbf{a}$ as illustrated in Fig. 5(a). The imitation reward is given as following:

$$r_t = 0.3 r_t^{\text{root}} + 0.2 r_t^{\text{pose}} + 0.15 r_t^{\text{vel}} + 0.15 r_t^{\text{ee}} + 0.2 r_t^{\text{com}},$$

$$r_t^{\text{root}} = \exp\left(-0.5 \|\hat{\mathbf{p}}_t^r - \mathbf{p}_t^r\|_2^2 - 0.5 \|\hat{\dot{\mathbf{p}}}_t^r - \dot{\mathbf{p}}_t^r\|_2^2 - 0.5 \|\hat{\mathbf{q}}_t^r - \mathbf{q}_t^r\|_2^2 - 0.05 \|\hat{\dot{\mathbf{q}}}_t^r - \dot{\mathbf{q}}_t^r\|_2^2 \right),$$

$$r_t^{\text{pose}} = \exp\left(-2 \sum \|\hat{\mathbf{q}}_t^j - \mathbf{q}_t^j\|_2^2 \right),$$

$$r_t^{\text{vel}} = \exp\left(-0.1 \sum \|\hat{\dot{\mathbf{q}}}_t^j - \dot{\mathbf{q}}_t^j\|_2^2 \right),$$

$$r_t^{\text{ee}} = \exp\left(-40 \sum \|\hat{\mathbf{p}}_t^e - \mathbf{p}_t^e\|_2^2 \right),$$

$$r_t^{\text{com}} = \exp\left(-\|\hat{\dot{\mathbf{p}}}_t^c - \dot{\mathbf{p}}_t^c\|_2^2 \right) \tag{24}$$

where $\mathbf{q}_t$ and $\mathbf{p}_t$ represent angle and global position while $\hat{}$ represent those of the reference.[5] As reference motion data, three motion capture clips, turning left ($\mathbf{t} = [1, 0, 0]$), going straight ($\mathbf{t} = [0, 1, 0]$), turning right ($\mathbf{t} = [0, 0, 1]$) from the Carnegie Mellon University motion capture (CMU mocap) database are utilized. Following the reference, PPO with same hyperparameters is used for RL algorithm.

Since the internal model does not exist at the first iteration ($l = 1$), the high-level planner is initialized by $q_\phi(\mathbf{h}|\mathbf{s};t) = \mathbf{w}^T \mathbf{t}$ where $\mathbf{w} = [-1, 0, 1]$. After the first iteration, high-level planner computes a command $\mathbf{h}_t$ that makes the model to best follow the horizontal position of the reference motion for 5 seconds ($\delta t = 0.1s$ and $K_{plan} = 50$). The corresponding reward function is as following:

$$r_k = -\|\hat{\mathbf{p}}_{h,k}^r - \mathbf{p}_{h,k}^r\|_2^2 \tag{25}$$

where $\mathbf{p}_{h,k}$ is the horizontal components of position vector at time step k.

---

[5]Each superscript denotes as following: r: root (pelvis), j: local joints, e: end-effectors (hands and feet), c: center-of-mass.

TABLE II

INPUT VARIABLES OF BASELINE TRANSITION MODELS

input variables	DISH	sas'	shs'	zaz'
$\mathbf{x}_k$	$\mathbf{z}_k$	$\mathbf{s}_k$	$\mathbf{s}_k$	$\mathbf{z}_k$
$\mathbf{y}_k$	$\mathbf{h}_k$	$\mathbf{a}_k$	$\mathbf{h}_k$	$\mathbf{a}_k$

TABLE III

COMPARISON OF RMS ERRORS BETWEEN REFERENCE, PLANNED, AND EXECUTED TRAJECTORIES FOR DIFFERENT TYPES OF INTERNAL MODELS.

	task1 (turn left)			task2 (go straight)			task3 (turn right)		
	‖ref-true‖	‖plan-ref‖	‖plan-true‖	‖ref-true‖	‖plan-ref‖	‖plan-true‖	‖ref-true‖	‖plan-ref‖	‖plan-true‖
DISH	**0.1290**	**0.1364**	**0.0743**	0.2073	0.1705	**0.1073**	**0.1550**	**0.1661**	**0.0974**
DISH+	0.1466	0.1704	0.1223	**0.1177**	0.1075	0.1177	0.1711	0.1747	0.0988
sas'	F	0.2474	F	F	0.0660	F	F	0.2178	F
shs'	0.2525	0.2036	0.3167	0.1385	**0.0561**	0.1280	0.2767	0.2140	0.2247
zaz'	F	0.2731	F	F	0.1994	F	F	0.3044	F

E. Training Internal Models

Internal models of DISH is trained to maximize the ELBO in (22) by using the APIAE approach [31] with hyperparameters as following: 3 adaptations ($R = 4$), 10 time steps ($K = 10$), 32 samples ($L = 32$), and time interval of 0.1s ($\delta t = 0.1$). The network architectures of transition network and inference network are shown in Fig 5(b) - 5(c).

For the baselines, the transition functions, $f_\phi(\mathbf{x}_{k+1}|\mathbf{x}_k, \mathbf{y}_k)$, were parameterized by neural networks having same architectures as DISH except for the input variables as shown in Table II. The loss function for baseline is as following:

$$\mathcal{L}_{sas'}(\phi) = \sum_{k=1}^{K} ||\mathbf{s}_k - \tilde{\mathbf{s}}_k||_2^2, \; \tilde{\mathbf{s}}_k = f_\phi^{sas'}(\tilde{\mathbf{s}}_{k-1}, \mathbf{a}_{k-1}),$$

$$\mathcal{L}_{shs'}(\phi) = \sum_{k=1}^{K} ||\mathbf{s}_k - \tilde{\mathbf{s}}_k||_2^2, \; \tilde{\mathbf{s}}_k = f_\phi^{shs'}(\tilde{\mathbf{s}}_{k-1}, \mathbf{h}_{k-1}),$$

$$\mathcal{L}_{zaz'}(\phi) = \sum_{k=1}^{K} ||\mathbf{s}_k - g(\tilde{\mathbf{z}}_k)||_2^2, \; \tilde{\mathbf{z}}_k = f_\phi^{zaz'}(\tilde{\mathbf{z}}_{k-1}, \mathbf{a}_{k-1}), \quad (26)$$

where $\tilde{\mathbf{s}}_1 = \mathbf{s}_1$, $\tilde{\mathbf{z}}_1$ is latent state for $\mathbf{s}_1$ inferred by VAE, and $g(\cdot)$ is learned generative network of VAE.

F. Task Configurations

Trajectory Following Tasks: Planning reward r_t penalizes the distance between the horizontal position of the root of humanoid character $\mathbf{p}_k^r$ and the that of reference trajectory $\bar{\mathbf{p}}_k$:

$$r_k = -||\bar{\mathbf{p}}_k - \mathbf{p}_{h,k}^r||_2^2. \quad (27)$$

Navigation Tasks: Planning cost r_t penalizes the distance between the horizontal position of the root of humanoid character $\mathbf{p}_k^r$ and the that of the goal $\mathbf{p}_{\text{goal}}$ and the collision with obstacles, while giving a reward on arrival:

$$r_k = -||\mathbf{p}_{\text{goal}} - \mathbf{p}_{h,k}^r||_2^2 - 10^5 \times (\text{IS\_CRASHED}) + 10^4 \times (\text{IS\_REACHED}). \quad (28)$$

G. High-level Planning

Empirically found that too short planning horizon (i.e., the greedy algorithm) prevents the agent to perform a task that requires far-sighted reasoning, such as a bug trap example. Meanwhile, planning with a too-long trajectory significantly reduces the sample efficiency, increases a computational burden, and finally degrades the performance. Therefore, the planning horizon and frequency were experimentally determined according to the characteristics of each task; a 5-second trajectory is planned at 1Hz and 4Hz for trajectory following task and navigation task, respectively.

H. Additional Results

Table III reports the RMS between reference, planned, and executed trajectories for each tasks. As illustrated in the table, DISHs showed the best performance. Although shs' sometimes showed smaller errors for the difference between the planed and reference trajectories, the errors between the reference and executed trajectory of DISHs are always smallest. This demonstrates that DISHs best learn the internal dynamics of the humanoid, making the most valid predictions for future
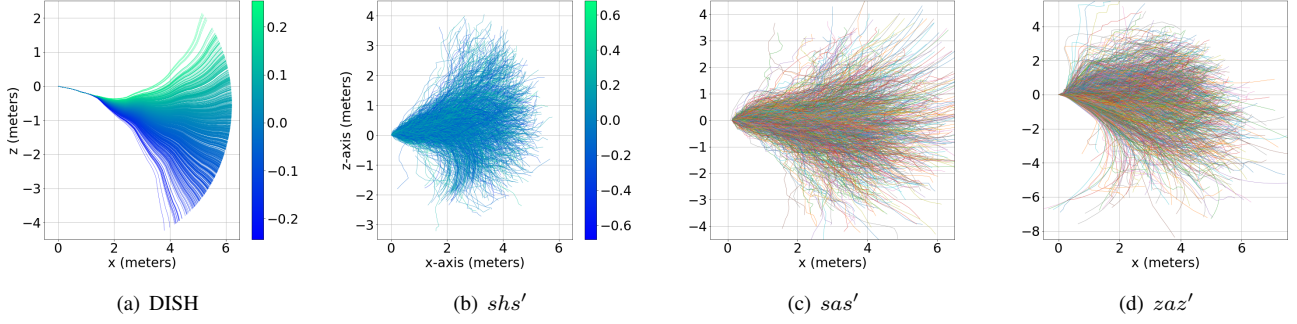
Fig. 6. Rollout samples from different types of internal models. (a) and (b) is colored by latent control value.

motion. Comparing DISH ($L = 1$) and DISH+ ($L = 2$), we can observe that DISH outperforms in the turning tasks while showing the worse performance in going straight. This is because the high-level planner is initialized to output only one of $\{-1, 0, 1\}$ (as shown in Appendix D), so the corresponding low-level policy of DISH is trained only around $\mathbf{h} \in \{-1, 0, 1\}$ rather than along the continuous $\mathbf{h}$ values. As a result, the DISH agent is only good at radical turns (not smooth turns), making it difficult to stabilize its heading direction properly. The ability to turn smoothly is obtained in the next iteration where the proper planning module is equipped, thus, although it lost some ability to turn fast, the DISH+ agent achieves the better ability to walk straight and the increased average performance (see Table I).

Fig. 6 shows rollout samples by varying the control values. Primarily, we fixed the control values h during each rollout to see the effect of latent control on the agents' trajectory. As compared in (a) and (b), the latent control in the DISH agent is observed to manipulates the angular velocity of the agent while we could not find any clear role of the latent control in the $sh's$ model; despite the existence of the hierarchical structure in the policies $sh's$ failed to learn the desirable internal model. Overall, except for DISHs, the generated trajectories are very noisy, which indicates that the baseline internal models are not suitable for planning the future movements of the humanoid.