
QueST: Self-Supervised Skill Abstractions for Learning Continuous Control

Atharva Mete¹, Haotian Xue¹, Albert Wilcox¹, Yongxin Chen^{1,2}, Animesh Garg^{1,2}

¹Georgia Institute of Technology, ²NVIDIA

Abstract

Generalization capabilities, or rather a lack thereof, is one of the most important unsolved problems in the field of robot learning, and while several large scale efforts have set out to tackle this problem, unsolved it remains. In this paper, we hypothesize that learning temporal action abstractions using latent variable models (LVMs), which learn to map data to a compressed latent space and back, is a promising direction towards low-level skills that can readily be used for new tasks. Although several works have attempted to show this, they have generally been limited by architectures that do not faithfully capture sharable representations. To address this we present Quantized Skill Transformer (QueST), which learns a larger and more flexible latent encoding that is more capable of modeling the breadth of low-level skills necessary for a variety of tasks. To make use of this extra flexibility, QueST imparts causal inductive bias from the action sequence data into the latent space, leading to more semantically useful and transferable representations. We compare to state-of-the-art imitation learning and LVM baselines and see that QueST’s architecture leads to strong performance on several multitask and few-shot learning benchmarks. Further results and videos are available at <https://quest-model.github.io>.

1 Introduction

One of the grand goals of robotic learning is a general-purpose model that can learn from complex multitask demonstration data and generalize to new tasks in a zero-shot or few-shot manner. While such general-purpose models have become ubiquitous in natural language (NLP) [66, 72, 52, 53, 63] and computer vision (CV) [77, 33, 7], they have eluded robotics researchers. Whereas CV and NLP can achieve positive transfer by scaling up models trained on internet scale datasets [33, 72, 56, 62, 39], even large scale robot data collection efforts [10, 11, 47, 20, 31] have been insufficient for this approach. To that end, we posit that in order to achieve positive transfer in robotics, it is important to design architectures that specifically lend themselves to efficient cross-task transfer.

There has recently been a surge of work towards the goal of learning generalist policies from large, diverse datasets. Several papers have used techniques such as action discretization [57, 16, 35, 10, 11, 47, 55] and implicit models [23, 15, 25] to model multimodal action distributions. In particular, the behavior transformer line of work shows that a carefully discretized action space combined with a GPT-style transformer leads to impressive capabilities modeling multimodal behavior distributions [57, 16, 35]. In another vein, several works have attempted to scale demonstration data and achieve positive transfer of low-level skills between high-level tasks [10, 11, 47, 46, 18, 12, 55, 3]. While these works have shown some transfer, for example applying policies for known tasks to unfamiliar objects, they have generally failed to achieve transfer of low-level skills to novel tasks [4]. We

¹Correspondence to: amete7@gatech.edu

hypothesize that in the relatively low-data regime of robot learning, it is promising to explicitly force the model to learn sharable representations. To that end we study latent variable models (LVMs), which learn to map data to a compressed latent space and back, introducing an information bottleneck which encourages the model to learn shared representations across the training data. Specifically we consider the application of LVMs to learn low-dimensional representations of action sequences. Such representations are termed temporal action abstractions or motion primitives – in this paper we refer to these abstractions as ‘skills’.

A wide body of work has considered the application of LVMs to robotics. One line of work learns temporal action abstractions (skills) in continuous latent space with a Gaussian prior [41, 50, 59]. While this line of work showed some initial promise of learning latent plans, it has failed to scale to difficult multitask settings due to the loose nature of the latent structure and posterior collapse issues that inhibit the learning of shared representations. On the other hand, recent work in CV and NLP has shown that vector-quantized discrete latent spaces are capable of learning semantically meaningful representations from data like phonetics in speech [9, 6] or melody in music [17, 2]. This insight, along with prior work showing that discretized action spaces can help to address the multimodality problem in when learning from large datasets [57, 16, 35, 14], motivates methods learning temporal action abstractions with discrete latent spaces. Several recent works have set out to do this [35, 74, 76, 30], showing some degree of positive transfer between tasks in multi-task and few-shot settings. However, they are generally limited by architectures that do not faithfully capture transferable representations [35, 74], or depend on state prediction and state-based objective functions which are impractical for many real robot tasks [76, 30].

In this paper, we present **Quantized Skill Transformer (QueST)**, a simple yet novel architecture for learning generalizable low-level skills within a discrete latent space. The key insight behind QueST is its ability to flexibly capture variable length motion primitives by representing them with a sequence of discrete codebook entries. We achieve this through a unique encoder-decoder architecture primarily designed to impart causal inductive bias in action sequence data into the latent space. Such formulation enables us to employ powerful sequence modeling approaches to plan and composably reason within the space of low-level skills. Through our experiments, we show that autoregressive modeling of these latent skills with a GPT-like transformer outperforms state-of-the-art baselines on challenging robotic manipulation benchmarks, where QueST shows an 8% improvement in multitask and 14% improvement in few-shot imitation learning over the next best baselines. We also conduct a detailed ablation and sensitivity study to validate our key architectural design decisions.

2 Related Works

The proposed framework in this paper introduces a methodology for self-supervised skill abstraction, followed by decision-making within this skill space. Several related works have explored similar sub-directions such as decision-making in the latent space and decision making with a transformer:

2.1 Learning from Offline Data

Behavior cloning (BC) [51] aims to learn a policy by directly mapping observations to actions, and is typically trained end-to-end using pre-collected pairs of observation and behavior data. While this on its surface is a simple supervised learning problem, there are several properties of robot demonstration data that should be considered when building BC systems. First, large BC datasets collected from a variety of human demonstrators tend to contain data sampled from multimodal distributions. To address this, some works opt to sample actions from Gaussian Mixture Models (GMM) [43], while others explore implicit models including those derived from energy-based models [23, 29] or diffusion models [15, 71, 14, 68, 28]. The Behavior Transformer (BeT) line of work [57, 16, 35] shows that transformer-based categorical policies in carefully discretized action spaces do a good job handling multimodal demonstrator distributions and QueST builds upon this by contributing a more capable discrete latent skill model.

Another key property of robot demonstration data is that sequential actions are often highly correlated with one another, and exploiting this can lead to stronger performance while ignoring it can lead to policies which are susceptible to temporally correlated confounders [61]. Recently several works have set out to handle this by predicting action chunks. For example, the Action Chunking Transformer (ACT) line of work [73, 24] shows that a transformer trained as a CVAE [60] to output chunks of

actions performs well for a wide variety of manipulation tasks, and diffusion policy [15] shows across the board improvements when predicting action chunks. As discussed in detail in Section 2.3, QueST builds on a long line of work which handles sequential correlations through temporally-extended action abstractions [41, 50, 59, 35, 74, 76, 30].

2.2 Multi-task and Few-shot Imitation Learning

In the past, robot learning researchers have approached multi-task decision making settings using a wide variety of methods such as supervised pre-training and fine-tuning [20, 42], meta-learning [22, 19] and action retrieval [48, 44]. There has recently been a large focus on multi-task language-conditioned imitation learning for robotics with several papers attempting to address the problem by training large models on large demonstration datasets [10, 11, 47, 46, 18, 12, 1]. While these papers achieve impressive multitask results, they mostly rely on sufficient data coverage and fail to generalize beyond their training distribution [4]. Thus, they lack abstractions that can readily be applied to learn new tasks, especially in a low-data regime. On the other hand LVMs like QueST are designed to learn sharable representations that can be applied to new tasks.

2.3 Decision Making in Learned Latent Spaces

LVMs, modeled by a paired encoder-decoder, have found extensive applications in computer vision [65, 32, 8] and generative models [56, 21, 13]. Recent studies also demonstrate the utility of latent space representations in robot decision-making, spanning offline RL [58, 50, 40], imitation learning [14, 67, 37, 35], and temporal action abstraction [57, 74, 75]. Most similar to our work are those that learn temporally abstracted discrete latent skill spaces. PRISE [74] learns single-step state-action abstractions within a discrete space and then does temporal abstraction by applying BPE tokenization. While this method shows promise in learning multi-task and few-shot policies, BPE is well suited for text and is known to suffer in domains with highly dynamic vocabularies, in robotics its equivalent to varying action distributions across unseen tasks. TAP [76] and H-GAP [30] utilize a self-supervised auto-encoder to learn skill codes, but their functionality relies on Model Predictive Control (MPC) using state prediction with ground-truth state-conditioned objective functions, which make it difficult to apply to real-world manipulation tasks. VQ-BeT [35] bears the strongest resemblance to QueST, also using a pre-trained discrete latent skill space to discretize the action space for a transformer-based policy prior. However, their quantization approach does not leverage the inherent structure in action sequences, limiting the representational capabilities of the latent space. Thus it’s shown to heavily rely on an continuous offset predictor for best performance. Unlike these works, QueST’s learned latent space is highly flexible yet structured and expressive, allowing it to effectively model action distribution across many distinct tasks in a meaningful shared representation.

3 Preliminary

3.1 Problem Setting

We consider a dataset $D = \{ \{ (O_1, a_1), \dots, (O_{T_i}, a_{T_i}) \}_{i=0}^{N_k}, L_k \}_{k=0}^M$ consisting of M robot interaction trajectories where a_t is a continuous-valued action and O_t is a tuple consisting of a high-dimensional sensory observation. The data is collected via either human teleoperation or scripted policies for M different task each with a label L_k . In our setting, O_t consist of RGB image observations from the front camera and gripper camera (if available) along with proprioceptive state of the agent. L_k is a natural language description of the k^{th} task but can also be a one-hot encoding.

3.2 Finite Scalar Quantization

We build on Finite Scalar Quantization (FSQ) [45] as a discrete bottleneck in our model. It’s a drop-in replacement for Vector Quantization (VQ) layers in VAEs with a simple scalar quantization scheme. Here the input representation e is projected to very few dimensions (typically 3 to 5) which are then bounded and rounded, creating an implicit codebook.

$$z = \text{round_ste}(f(e)), \text{ where } f \text{ is the bounding function} \tag{1}$$

Given a feature vector $e \in \mathbb{R}^d$ to quantize, instead of learning a parameterized codebook [65] and quantizing e by matching the nearest neighbor in the codebook, FSQ quantizes e by first bounding

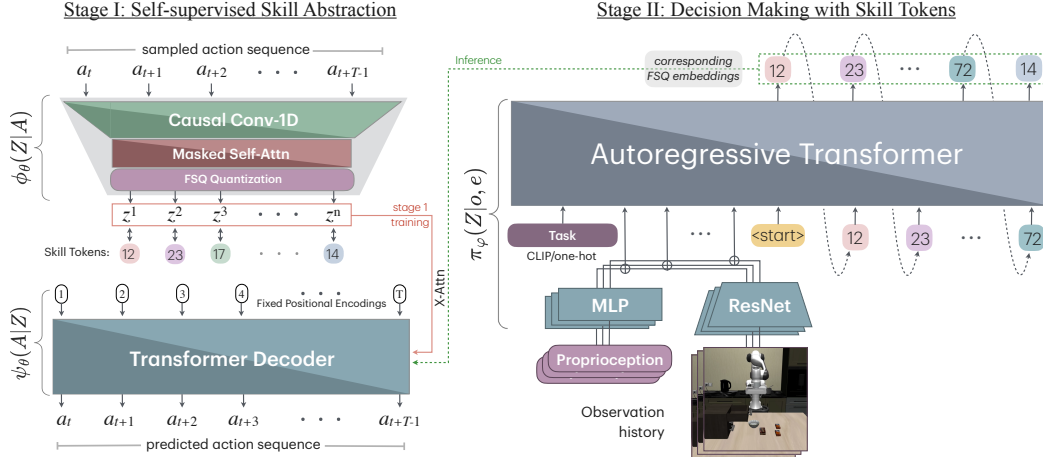


Figure 1: **Overview of Quantized Skill Transformer**: we factorize the policy that outputs action based on task descriptions e and observations o encoding into two parts: $\pi(A|o, e) = \psi_\theta(A|Z)\pi_\varphi(Z|o, e)$, where Z is a sequence of skill tokens for the action sequence A . In Stage I, we learn skill abstraction in a self-supervised way with a quantized autoencoder. In Stage II, we learn skill-based policy in the style of next-token prediction using a multi-modal transformer.

it into certain range with f (e.g. $f = \lfloor \alpha/2 \rfloor \odot \tanh(e)$) and then rounding each dimension into integer numbers directly with straight-through gradients (round_ste). $\alpha \in \mathbb{Z}^d$ defines the width of the codebook for each dimension (e.g. $\alpha = [8, 5, 5, 5], d = 4$). Finally, it is easy to see that the size of the quantization space is $\prod_{i=1}^d \alpha_i$. An MLP can be used to transform z into continuous space of required dimension further.

A common problem with vector-quantized codebooks (VQ) [65] is the under-utilization of the codebook. Recent works have attempted to address by heuristics like reinitializing the codebook, stochastic formulations, or some regularization [34, 69]. In contrast, FSQ achieves much better codebook utilization for large codebook sizes with much fewer parameters and simplified training without any auxiliary losses or aforementioned tricks. Due to its simplicity and proven benefits, we use FSQ in our main experiments, but since many prior works in this space use VQ we also perform an ablation with it (see section 5.6).

4 Method

In this section, we describe the key ideas behind Quantized Skill Transformer. In Section 4.1 we present our encoder-decoder architecture, which is designed to provide the flexibility to learn a wide range of skills with inductive biases to ensure that the learned skills are useful. In Section 4.2, we detail our skill prior, which we train to autoregressively predict codebook skills. Our full pipeline is shown in Figure 1.

4.1 Stage I: Learning the Skill Codebook

As a motivating example, consider the task of lifting a pot and placing it on a stove beside. This consist of primitives like reaching the pot, grasping it, lifting it to a certain height, reaching the stove and finally placing it on the stove. Each of these primitives are of variable lengths, and to properly model these skills it is important to learn a latent skill space with the flexibility to model all of them. At the same time, it is important that the learned skills are semantically meaningful so that they can be reused for new tasks, for example reusing the reaching skill for an object lifting task. In order to address these desiderata, we introduce the novel autoencoder architecture shown in Figure 1 consisting of an encoder ϕ_θ and decoder ψ_θ .

The input to the encoder ϕ_θ is an action sequence $a_{t:t+T-1}$ sampled from the dataset, which we pass through several 1D causal strided-convolution layers [64]. This step reduces the sequence length to achieve the desired temporal abstraction depending on the stride lengths and the number of layers.

We follow the convolutional layers with masked self-attention layers for sequence modeling. With a downsampling factor of F , the encoder outputs in total $n = T/F$ embeddings. The embeddings are then quantized using FSQ as per the equation 1 into n discrete latent codes $\{z^i\}$ termed as skill tokens:

$$(z^1, \dots, z^n) = \text{FSQ}(\phi_\theta(a_t, \dots, a_{t+T-1})). \quad (2)$$

Having an input sequence of actions mapped to multiple skill tokens gives this architecture more flexibility to model complex sequences of actions. At the same time, each component of the encoder is causal, meaning that an output representation at a position t cannot depend on input from any future timesteps. We found this inductive bias to encourage the model to learn semantically useful action representations by modeling the inherent causality in the action data. We validate this design choice in the ablations. (see section 5.6)

Typical autoencoder decoders are simply mirrored versions of the encoders, but this would prevent the decoder from attending to all quantized codes. This is important because individual codes do not represent anything meaningful but a sequence of codes represents a particular meaningful motion [45]. In order to maintain causality while attending to all codes, the decoder ψ_θ cross attends between fixed sinusoidal positional embedding inputs and the skill tokens, similarly with [73]. The architecture is a transformer decoder block consisting of alternate masked self-attention and cross-attention layers, after which the output embeddings are projected back to the original action dimension using an MLP layer. Thus, given a sequence Z of skill codes, ψ_θ reconstructs the original action

$$(\hat{a}_t, \dots, \hat{a}_{t+T-1}) = \psi_\theta(z^1, \dots, z^n) \quad (3)$$

As in [35], the autoencoder is trained by minimizing the ℓ_1 reconstruction loss:

$$\mathcal{L}_{\text{recon}}(\theta) = \|\psi_\theta(\text{FSQ}(\phi_\theta(a_{t:t+T-1}))) - a_{t:t+T-1}\|_1. \quad (4)$$

Unlike prior work which often conditions on the state as well as the actions [30, 5, 74, 37], we choose to learn state-independent abstractions that solely capture motion primitives irrespective of the current scene or task. Through our experiments we show that our model learns generalizable abstractions that are shared and can be transferred across tasks.

4.2 Stage II: Learning the Skill Prior

After training the encoder ϕ_θ and decoder ψ_θ , we train a skill prior $\pi_\varphi(Z|e, o)$ to predict skills $Z = z^{1:n}$ corresponding to the demonstrator action distribution conditioned on a task embedding e and a length h sequence of image observations and proprioception inputs, $o = (i_{t-h}, p_{t-h}), \dots, (i_t, p_t)$. We encode image observations with a separate learned vision encoder for each camera view and encode proprioception using an MLP encoder, all of which are trained end-to-end with the rest of the skill prior. The observation token \mathcal{T}_t^o for a timestep t is obtained by concatenating outputs from all the aforementioned encoders. Task embeddings are designed specifically for each environment suite, as discussed in more detail in Section 5. See Appendix B for more details about the encoders.

Because skill tokens are highly dependent on one another according to the complex nonlinear representations learned by the autoencoder, it is important that the skill prior has the modeling capacity to reason about these dependencies. To achieve this, we employ a decoder-only transformer to model the distribution of skill tokens $\pi_\varphi(Z|\mathcal{T}_{t-h:t}^o, e)$ autoregressively as:

$$\pi_\varphi(Z|\mathcal{T}_{t-h:t}^o, e) = \prod_{i=1}^n \pi_\varphi(z^i | \langle \mathbf{s} \rangle, z^{1:i-1}, \mathcal{T}_{t-h:t}^o, e) \quad (5)$$

where $\langle \mathbf{s} \rangle$ is a learnable start token that marks the start of skill tokens. We add sinusoidal positional embeddings only to the skill tokens. To optimize the skill prior, we sample a sequence of demonstrator actions $a_{t:t+T-1}$ and use the trained encoder ϕ_θ to extract a latent skill vector $Z_t = z^{1:n}$ according to Equation 2. Then, we optimize π_φ using the following negative log-likelihood loss:

$$\mathcal{L}_{\text{task}}(\varphi) = -\log \pi_\varphi(Z_t|\mathcal{T}_{t-h:t}^o, e). \quad (6)$$

The full skill prior pipeline is shown in Figure 1.

Few-Shot Finetuning: For few-shot finetuning on new tasks, we use a model pre-trained on large set of tasks and finetune it on a small number of demonstrations (5 in our experiments) from the held-out task. Although finetuning only stage-2 is enough, we empirically found that finetuning the decoder on the predicted skill tokens gives a boost in the performance. Specifically, we finetune the decoder using following decoder loss:

$$\mathcal{L}_{\text{decoder}}(\theta) = \|\psi_{\theta}(\text{sg}(\hat{Z}_t)) - a_{t:t+T-1}\|_1 \quad (7)$$

where sg is the stop gradient operator. We present the results with and without decoder finetuning both. Additionally, we note that the encoder is still frozen in this setting.

4.3 Inference with Quantized Skill Transformer

At inference time, QueST uses the skill prior π_{φ} alongside the decoder ψ_{θ} to sample actions. Conditioned on the encoded observation history $\mathcal{T}_{t-h:t}^o$ and task embedding e , we use top- k sampling with a temperature of τ to autoregressively sample a skill vector $\hat{Z} \sim \pi_{\varphi}(\cdot | \mathcal{T}_{t-h:t}^o, e)$ from the skill prior. In practice, we find $k = 5$ and $\tau = 1$ to work well across all environments. Then, we use the decoder to map the skill vector back to the action space, producing a sequence of predicted actions $\hat{a}_{t:t+T-1} = \psi_{\theta}(\hat{Z})$. In a receding horizon fashion, we execute the first $T_a \leq T$ actions before replanning.

5 Experiments

We design the experiments to empirically evaluate the performance of Quantized Skill Transformer in three practical settings: (1) Multitask IL, (2) Few-shot transfer, and (3) Long-horizon IL. Lastly, we perform some ablations to empirically justify our model design choices.

5.1 Benchmarks and Baselines

We use the following benchmark suites to evaluate in the settings discussed above:

LIBERO [38] is a lifelong learning benchmark featuring several task suites consisting of a variety of language-labeled rigid- and articulated-body manipulation tasks. Specifically, we evaluate on the LIBERO-90 suite, which consists of 90 manipulation tasks, and the LIBERO-LONG suite, which consists of 10 long-horizon tasks composed of two tasks from the LIBERO-90 suite. As described in more detail below, we use the LIBERO benchmark to study the multitask IL, few-shot transfer and long-horizon IL settings. Because tasks from this benchmark are language-annotated, we use the output of a frozen CLIP [54] encoder for the task conditioning input e .

MetaWorld [70] features a wide range of manipulation tasks designed to test few-shot learning algorithms. We use the Meta-Learning 45 (ML45) suite which consists of 45 training tasks and 5 difficult held-out tasks which are structurally similar to the training tasks. We use this benchmark to test multi-task and few-shot learning. Because this benchmark does not include language labels, we use learned task embeddings e for task conditioning.

Baselines: We compare to the following baselines, which include similar discrete LVM pipelines as well as state of the art imitation learning algorithms:

1. The **ResNet-T** model from [38], which encodes observation and task instructions using ResNet-18 with FiLM [49], applies a transformer sequence model, and uses a GMM head to predict actions.
2. The UNet-based **diffusion policy** from [15], which uses a 1D convolutional UNet to map samples from a Gaussian prior to action samples from the demonstrator distribution according to a learned denoising process.
3. **ACT [73]**, which trains a transformer as a CVAE [60] to predict action chunks.
4. **VQ-BeT [35]**, which learns a discrete latent space using a VQ-VAE [65] and uses a transformer to predict discrete latent codes.
5. **PRISE [74]**, which first quantizes observation-action pairs and performs temporal abstraction using byte pair encoding (BPE) to learn a skill token vocabulary, which it uses as an action space for few-shot learning.

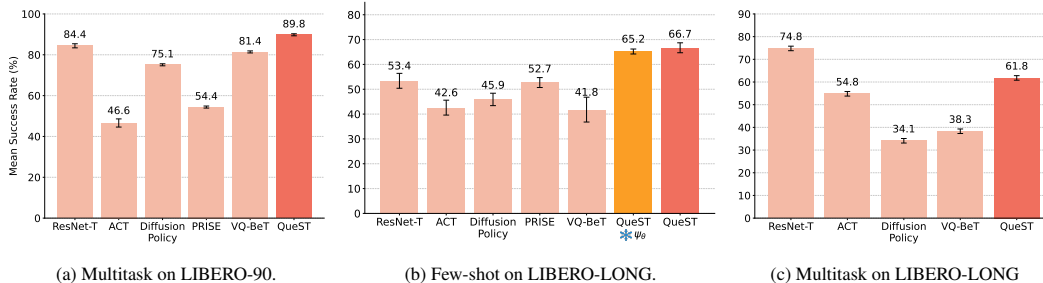


Figure 2: Multitask performance on LIBERO-90 (a) and LIBERO-LONG (c), and few shot performance on LIBERO-LONG (b). For (a) and (c) we train on the datasets described in Sections 5.2 and 5.4. For (b) we finetune the model from (a) on a condensed dataset as described in Section 5.3. Results show the mean and error bar represents standard error across four random seeds for multitask experiments and nine random seeds for fewshot experiments. Results for PRISE are taken from Zheng et al. [74], and the others we reimplemented and ran ourselves.

For a detailed discussion between QueST and the baseline methods, please refer to Appendix B.5.

5.2 Performance on Multitask BC

We evaluate the goal-conditioned multi-task imitation learning capabilities of QueST and the baselines using the LIBERO-90 and ML45 benchmark suites. For LIBERO-90, the learner receives 50 expert demonstrations per task from the author-provided dataset. For ML45, we use the scripted policies provided in the official Metaworld codebase to collect 100 demonstrations per task. We evaluate the model at the end of training and for each task run 40 evaluation rollouts (50 for MetaWorld) starting from the initial states selected sequentially from a predefined set. We report the aggregated results across 4 seeds (5 seeds for MetaWorld).

In Figure 2a, we present the average success rate across 90 tasks in LIBERO-90 against the aforementioned baselines. Quantized Skill Transformer achieves state-of-the-art results on LIBERO-90 benchmark, outperforming the baseline VQ-BeT and Diffusion Policy by a margin of 8 and 13% respectively. We attribute its performance to its learned latent space that enables effective knowledge sharing across tasks. While VQ-BeT also shows strong performance, we see that QueST’s architecture lends itself better to sharing representations across tasks. Our implementation of ResNet-T achieves significantly better performance than the reported number (16.8%) in [38] but is still lower than QueST. Figure 3a shows the average success rate across 45 tasks in ML45 benchmark. Being a simpler benchmark, all methods perform almost similar in Multitask-IL setting which is consistent with the trend observed in [74].

We attribute the reasonably good performance of the diffusion policy to its nature as a latent variable model, which employs a continuous latent variable with the same dimensionality as the actions. However, the consistent outperformance of QueST over the diffusion policy provides compelling evidence for the benefits of using a bottlenecked latent variable. This bottleneck encourages the model to learn shared representations, resulting in enhanced performance. While both VQ-BeT and PRISE employ a latent bottleneck, VQ-BeT’s architecture neglects the inherent inductive biases in the action data, which we believe results in a less well-structured latent space. PRISE incorporates this using a latent forward transition model, but is bottlenecked by the use of BPE which we posit is not suitable for such a dynamic latent space.

5.3 Few-shot Transfer to Unseen Tasks

In this setting we take the pretrained model from section 5.2 and test its 5-shot performance on unseen tasks from LIBERO-LONG and held-out set in ML45. We sample only five demonstrations for each task, generate the skill tokens using pretrained encoder and use them to finetune the skill prior and the decoder as described in Section 4.2. We also present the results without finetuning the decoder (frozen ψ_θ in figure 2b & 3b) to validate its generalization to unseen skill tokens sequences.

Figure 2b shows the average success rate for 5-shot IL across 8 unseen tasks in LIBERO-LONG. QueST achieves SOTA performance, surpassing all other baselines by an absolute margin of 14%.

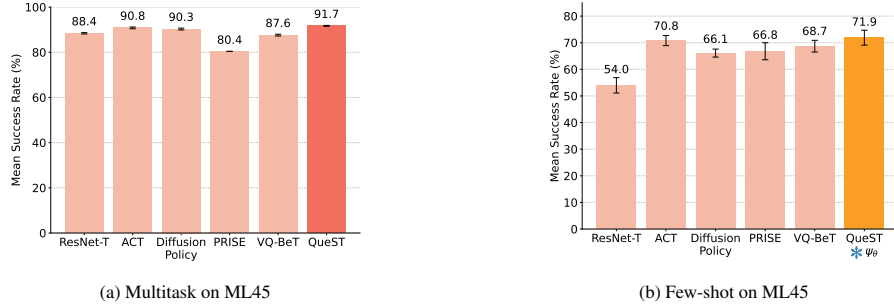


Figure 3: Multitask and few-shot success rate on the Metaworld ML45 task suite. In (a) we train on the dataset described in Section 5.2, and in (b) we finetune the model from (a) using 5 demonstrations each from a set of held out tasks. Results show the mean and error bar represents standard error across five random seeds. Results for PRISE are taken from Zheng et al. [74], and the others we reimplemented and ran ourselves.

	VQ	Obs. Cond.	Mirror Dec.	Ours
LIBERO-90	81.2 ± 0.6	81.9 ± 1.1	86.3 ± 0.9	88.6 ± 0.4
Few Shot	62.5 ± 2.0	61.3 ± 2.2	45.4 ± 2.0	68.8 ± 1.7

Table 1: Success rates after ablating design details of QueST. We present the mean across four random seeds and error tolerances show the standard error.

Though we see a marginal drop of 1.5% without decoder finetuning, it still outperforms all the baselines. These results highlight the superiority of QueST in learning transferable representations of action abstractions and effectively leveraging them for downstream decision making. For a fair comparison, we also tried fine-tuning the decoder of VQ-BeT but did not observe any gains from it. VQ-BeT struggles in this setting as it heavily relies on offset head to output continuous action corrections which requires more data-samples for sufficient coverage in the continuous action space. Figure 3b shows the average success rate for 5-shot IL across 5 unseen tasks in MetaWorld. Similar to multitask results, all methods perform comparably, with QueST showing a slight improvement over the others. QueST leverages its learned skill tokens to compositionally model their distribution for an unseen task in just 5 demonstration examples. For few-shot evaluation protocol please refer Appendix D.1.

5.4 Long-horizon BC

In this setting, we aim to purely study and compare the performance of our model on long-horizon tasks. We train the model (both stages) solely on LIBERO-LONG complete dataset (50 demonstrations per task) and evaluate with the same scheme as described earlier.

Figure 2c shows the average success rate across 10 LIBERO-LONG tasks. All LVMs perform significantly worse than the ResNet-T model. We attribute this to the relatively small size of LIBERO-LONG dataset which is just not enough to learn a good latent space in LVMs. QueST still outperforms all LVM baselines by a large margin demonstrating its long-horizon modeling capabilities.


Overall, we see that our model outperforms baselines like VQ-BeT in multitask settings, showing stronger modelling capacity. At the same time, it has the correct latent structure to outperform baselines like diffusion in few shot settings, especially even with frozen decoder, indicating strong generalization capabilities of learned skill-space.

5.5 Latency

Our pipeline runs at 33Hz, which is more than sufficient for vision-based real-robot control where most camera systems run at 30 fps. For comparison, our implementation of Resnet-T, VQ-BeT, ACT and Diffusion Policy run at 100Hz, 100Hz, 50Hz, and 12Hz respectively.

	Non Causal ϕ_θ	Non Causal ψ_θ	Fully Non Causal	Ours
LIBERO-90	82.0 ± 1.6	85.1 ± 1.8	78.5 ± 0.5	88.6 ± 0.4
Few Shot	58.8 ± 3.0	61.6 ± 2.5	56.1 ± 1.8	68.8 ± 1.7

Table 2: Success rates after ablating the causality in QueST. We present the mean across four random seeds and error tolerances show the standard error.



(a) Downsampling factor sensitivity. (b) Codebook size sensitivity.

Figure 4: We conduct a sensitivity experiment across downsampling factors (a) and codebook sizes (b) on the LIBERO benchmark. For (a) we fix a sequence length of $T = 32$. Overall, we see that the few-shot version is more sensitive to hyperparameters and that $F = 4$ with 1024 codebook vectors are good choices.

5.6 Ablations

We validate the proposed architecture by ablating some of its key design decisions. All the ablations are performed on LIBERO studying their effects on both multitask and fewshot IL settings.

- Vector Quantization:** We replace the FSQ layer with a Vector Quantization layer of nearly the same codebook size, shown in the VQ column of Table 1. We see that FSQ’s superior codebook utilization leads to an improvement in performance.
- Observation Conditioned Decoder:** Many prior condition the action decoder with current observation [76, 30]. We experiment with this by appending observation tokens to the skill tokens and allowing the transformer decoder to jointly cross-attend to both, shown in Table 1. We see that conditioning on observations leads to a deterioration in performance.
- Mirrored Decoder:** Following a typical autoencoder design, we use a decoder that mirrors the encoder, using transposed convolutions instead of strided convolutions, and with the strides in reverse order as in the encoder. This decoder directly takes skill-token embeddings as input and outputs the continuous actions, and results are shown in Table 1. We see that this method performs worse, suggesting attending to all quantized codes in z , as our decoder does, is important for faithfully predicting actions.
- Causality:** We ablate the use of causal layers in various parts of our network in Table 2. We see that removing causality from any part of our architecture leads to worse performance, suggesting that causal masking imparts the model with a helpful inductive bias for modeling robot action data.

We also perform a sensitivity experiment over several hyperparameters including downsampling factor and codebook size in Figure 4. Across the board we see that the hyperparameters are more important in the difficult few-shot learning setting. In Figure 4a we see that both algorithms have the best performance with a modest downsampling factor of $F = 4$, and in Figure 4b we see that QueST does well with a 1024 codebook vectors. For more discussion on ablations please refer Appendix C.

5.7 Latent Skill-Space Analysis

We present a t-SNE visualization (Figure 5) illustrating the learned skill-space across multiple set of similar tasks. We consider four different combinations of similar tasks to effectively examine the z -embeddings corresponding to their trajectories. Each data point in the plot represents a vector of n z -embeddings at a specific timestep throughout the entire episode, with decreasing transparency indicating temporal progression. We show that the QueST encoder learns a semantically meaningful skill-space that encodes shared representations of similar motion primitives across different tasks. This analysis includes the first 11 tasks from LIBERO-90. For better comprehension, we encourage readers to review the corresponding rollouts on the website. Notably, the skill-space learning happens in the first stage training which does not make use of any task labels.

6 Conclusion

We present Quantized Skill Transformer, a novel LVM architecture for learning sharable skills in a discrete latent space. The key idea behind QueST is to represent action sequences as a series

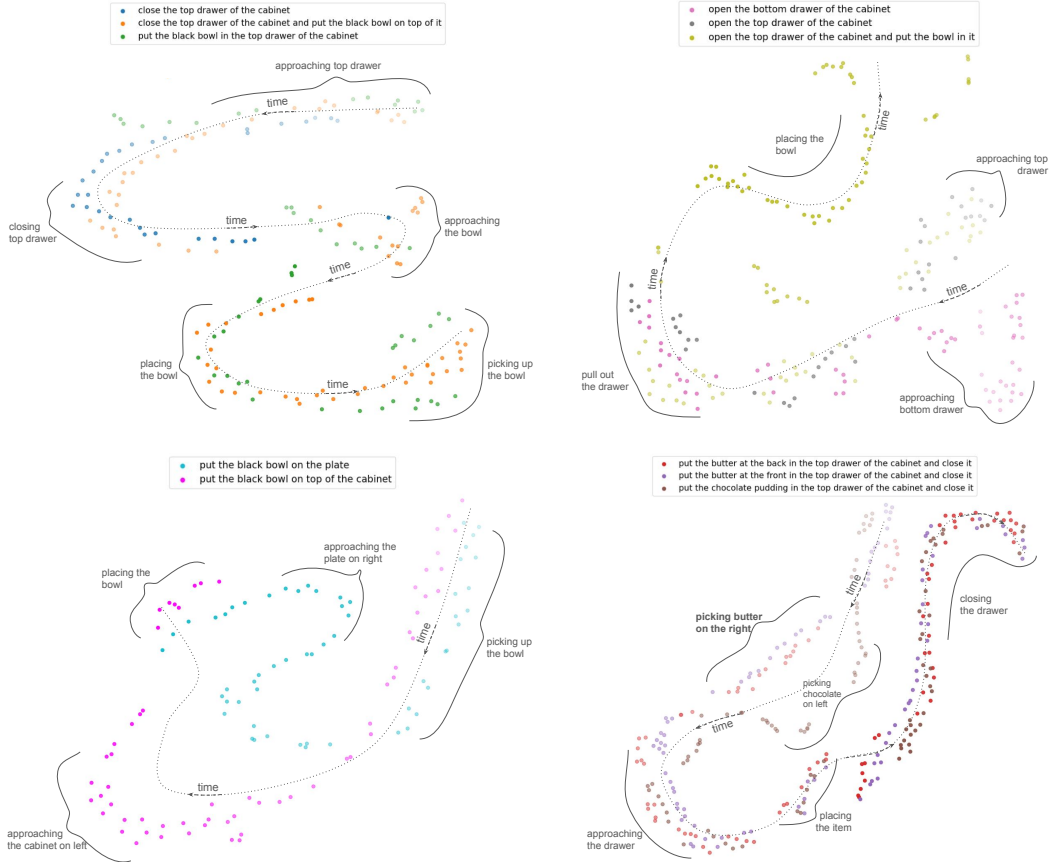


Figure 5: t-SNE visualization of skill-token embeddings. Here, the transparency decreases as the episode progresses. The overall patterns clearly shows how similar motion primitives like approaching, picking and placing from different tasks are aligned with one another.

of codebook vectors, and we demonstrate that using causal convolutions and masked transformers provides an inductive bias that encourages the model to learn useful shared representations. We evaluate QueST across 145 robot manipulation tasks, and show that it outperforms several state-of-the-art baselines in multitask and few-shot learning settings. Our results highlight the usefulness of QueST’s encoder (decoder) as semantically-sound, task-agnostic tokenizer (detokenizer) for continuous actions, and its potential to leverage Large Multi-modal Language Models in stage-2.

Limitations: While the benchmarks we consider encompass a wide variety of tasks, the held-out tasks are still structurally similar to the pretraining set, which makes few-shot transfer feasible. In scenarios with a more diverse task, current model may struggle to solve new tasks solely within the learned skill space. A promising direction is to train stage-1 on larger datasets, such as Open X-Embodiment [47], with an expanded codebook that could capture more diverse motion primitives. Additionally, our current architecture only accounts for causality. Future work should explore other inductive biases, like geometric invariance and dynamic consistency, to enhance abstraction learning.

A statement on societal impact: This paper works towards the broader goal of automating a wide range of manipulation tasks. While this can have positive impacts, such as helping people with mobility impairments or performing menial tasks humans would rather not do, it can also have negative impacts such as automating peoples’ jobs away and further concentrating wealth in the hands of a handful of companies. It is important that we in the machine learning community advocate for equitable use of the technology we develop.

References

- [1] Q-transformer: Scalable offline reinforcement learning via autoregressive q-functions. In *7th Annual Conference on Robot Learning*, 2023.
- [2] A. Agostinelli, T. I. Denk, Z. Borsos, J. Engel, M. Verzetti, A. Caillon, Q. Huang, A. Jansen, A. Roberts, M. Tagliasacchi, M. Sharifi, N. Zeghidour, and C. Frank. Musiclm: Generating music from text, 2023.
- [3] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, A. Herzog, D. Ho, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, E. Jang, R. J. Ruano, K. Jeffrey, S. Jesmonth, N. J. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, K.-H. Lee, S. Levine, Y. Lu, L. Luu, C. Parada, P. Pastor, J. Quiambao, K. Rao, J. Rettinghouse, D. Reyes, P. Sermanet, N. Sievers, C. Tan, A. Toshev, V. Vanhoucke, F. Xia, T. Xiao, P. Xu, S. Xu, M. Yan, and A. Zeng. Do as i can, not as i say: Grounding language in robotic affordances, 2022.
- [4] M. Ahn, D. Dwibedi, C. Finn, M. G. Arenas, K. Gopalakrishnan, K. Hausman, B. Ichter, A. Irpan, N. Joshi, R. Julian, S. Kirmani, I. Leal, E. Lee, S. Levine, Y. Lu, I. Leal, S. Maddineni, K. Rao, D. Sadigh, P. Sanketi, P. Sermanet, Q. Vuong, S. Welker, F. Xia, T. Xiao, P. Xu, S. Xu, and Z. Xu. Autort: Embodied foundation models for large scale orchestration of robotic agents, 2024.
- [5] A. Ajay, A. Kumar, P. Agrawal, S. Levine, and O. Nachum. Opal: Offline primitive discovery for accelerating offline reinforcement learning, 2021.
- [6] A. Baevski, H. Zhou, A. Mohamed, and M. Auli. wav2vec 2.0: A framework for self-supervised learning of speech representations. *CoRR*, abs/2006.11477, 2020. URL <https://arxiv.org/abs/2006.11477>.
- [7] Y. Bai, X. Geng, K. Mangalam, A. Bar, A. Yuille, T. Darrell, J. Malik, and A. A. Efros. Sequential modeling enables scalable learning for large vision models. *arXiv preprint arXiv:2312.00785*, 2023.
- [8] H. Bao, L. Dong, S. Piao, and F. Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.
- [9] Z. Borsos, R. Marinier, D. Vincent, E. Kharitonov, O. Pietquin, M. Sharifi, D. Roblek, O. Teboul, D. Grangier, M. Tagliasacchi, and N. Zeghidour. Audioldm: a language modeling approach to audio generation, 2023.
- [10] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, et al. Rt-1: Robotics transformer for real-world control at scale. *arXiv preprint arXiv:2212.06817*, 2022.
- [11] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, K. Choromanski, T. Ding, D. Driess, A. Dubey, C. Finn, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023.
- [12] A. Buckler, L. Figueredo, S. Haddadin, A. Kapoor, S. Ma, S. Vemprala, and R. Bonatti. Latte: Language trajectory transformer. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7287–7294. IEEE, 2023.
- [13] H. Chang, H. Zhang, L. Jiang, C. Liu, and W. T. Freeman. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11315–11325, 2022.
- [14] L. Chen, S. Bahl, and D. Pathak. Playfusion: Skill acquisition via diffusion from language-annotated play. In *Conference on Robot Learning*, pages 2012–2029. PMLR, 2023.
- [15] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song. Diffusion policy: Visuomotor policy learning via action diffusion. *arXiv preprint arXiv:2303.04137*, 2023.
- [16] Z. J. Cui, Y. Wang, N. M. M. Shafiullah, and L. Pinto. From play to policy: Conditional behavior generation from uncurated robot data. *arXiv preprint arXiv:2210.10047*, 2022.
- [17] P. Dhariwal, H. Jun, C. Payne, J. W. Kim, A. Radford, and I. Sutskever. Jukebox: A generative model for music, 2020.
- [18] D. Driess, F. Xia, M. S. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. Vuong, T. Yu, et al. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023.
- [19] Y. Duan, M. Andrychowicz, B. C. Stadie, J. Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba. One-shot imitation learning, 2017.

- [20] F. Ebert, Y. Yang, K. Schmeckpeper, B. Bucher, G. Georgakis, K. Daniilidis, C. Finn, and S. Levine. Bridge data: Boosting generalization of robotic skills with cross-domain datasets, 2021.
- [21] P. Esser, R. Rombach, and B. Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021.
- [22] C. Finn, T. Yu, T. Zhang, P. Abbeel, and S. Levine. One-shot visual imitation learning via meta-learning, 2017.
- [23] P. Florence, C. Lynch, A. Zeng, O. A. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson. Implicit behavioral cloning. In *Conference on Robot Learning*, pages 158–168. PMLR, 2022.
- [24] Z. Fu, T. Z. Zhao, and C. Finn. Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. *arXiv preprint arXiv:2401.02117*, 2024.
- [25] H. Ha, P. Florence, and S. Song. Scaling up and distilling down: Language-guided robot skill acquisition. In *Proceedings of the 2023 Conference on Robot Learning*, 2023.
- [26] S. Haldar, Z. Peng, and L. Pinto. Baku: An efficient transformer for multi-task policy learning, 2024. URL <https://arxiv.org/abs/2406.07539>.
- [27] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [28] M. Janner, Y. Du, J. B. Tenenbaum, and S. Levine. Planning with diffusion for flexible behavior synthesis. *arXiv preprint arXiv:2205.09991*, 2022.
- [29] D. Jarrett, I. Bica, and M. van der Schaar. Strictly batch imitation learning by energy-based distribution matching. *Advances in Neural Information Processing Systems*, 33:7354–7365, 2020.
- [30] Z. Jiang, Y. Xu, N. Wagener, Y. Luo, M. Janner, E. Grefenstette, T. Rocktäschel, and Y. Tian. H-gap: Humanoid control with a generalist planner. *arXiv preprint arXiv:2312.02682*, 2023.
- [31] A. Khazatsky, K. Pertsch, S. Nair, A. Balakrishna, S. Dasari, S. Karamcheti, S. Nasiriany, M. K. Srirama, L. Y. Chen, K. Ellis, P. D. Fagan, J. Hejna, M. Itkina, M. Lepert, Y. J. Ma, P. T. Miller, J. Wu, S. Belkhale, S. Dass, H. Ha, A. Jain, A. Lee, Y. Lee, M. Memmel, S. Park, I. Radosavovic, K. Wang, A. Zhan, K. Black, C. Chi, K. B. Hatch, S. Lin, J. Lu, J. Mercat, A. Rehman, P. R. Sanketi, A. Sharma, C. Simpson, Q. Vuong, H. R. Walke, B. Wulfe, T. Xiao, J. H. Yang, A. Yavary, T. Z. Zhao, C. Agia, R. Bajjal, M. G. Castro, D. Chen, Q. Chen, T. Chung, J. Drake, E. P. Foster, J. Gao, D. A. Herrera, M. Heo, K. Hsu, J. Hu, D. Jackson, C. Le, Y. Li, K. Lin, R. Lin, Z. Ma, A. Maddukuri, S. Mirchandani, D. Morton, T. Nguyen, A. O’Neill, R. Scalise, D. Seale, V. Son, S. Tian, E. Tran, A. E. Wang, Y. Wu, A. Xie, J. Yang, P. Yin, Y. Zhang, O. Bastani, G. Berseth, J. Bohg, K. Goldberg, A. Gupta, A. Gupta, D. Jayaraman, J. J. Lim, J. Malik, R. Martín-Martín, S. Ramamoorthy, D. Sadigh, S. Song, J. Wu, M. C. Yip, Y. Zhu, T. Kollar, S. Levine, and C. Finn. Droid: A large-scale in-the-wild robot manipulation dataset. 2024.
- [32] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [33] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023.
- [34] D. Lee, C. Kim, S. Kim, M. Cho, and W.-S. Han. Autoregressive image generation using residual quantization, 2022.
- [35] S. Lee, Y. Wang, H. Etukuru, H. J. Kim, N. M. M. Shafiqullah, and L. Pinto. Behavior generation with latent actions. *arXiv preprint arXiv:2403.03181*, 2024.
- [36] S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end training of deep visuomotor policies, 2016.
- [37] Z. Liang, Y. Mu, H. Ma, M. Tomizuka, M. Ding, and P. Luo. Skilldiffuser: Interpretable hierarchical planning via skill abstractions in diffusion-based task execution. *arXiv preprint arXiv:2312.11598*, 2023.
- [38] B. Liu, Y. Zhu, C. Gao, Y. Feng, Q. Liu, Y. Zhu, and P. Stone. Libero: Benchmarking knowledge transfer for lifelong robot learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- [39] H. Liu, C. Li, Q. Wu, and Y. J. Lee. Visual instruction tuning, 2023.
- [40] J. Luo, P. Dong, J. Wu, A. Kumar, X. Geng, and S. Levine. Action-quantized offline reinforcement learning for robotic skill learning. In *Conference on Robot Learning*, pages 1348–1361. PMLR, 2023.

- [41] C. Lynch, M. Khansari, T. Xiao, V. Kumar, J. Tompson, S. Levine, and P. Sermanet. Learning latent plans from play. *Conference on Robot Learning (CoRL)*, 2019. URL <https://arxiv.org/abs/1903.01973>.
- [42] Z. Mandi, F. Liu, K. Lee, and P. Abbeel. Towards more generalizable one-shot visual imitation learning. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 2434–2444. IEEE, 2022.
- [43] A. Mandlekar, D. Xu, J. Wong, S. Nasiriany, C. Wang, R. Kulkarni, L. Fei-Fei, S. Savarese, Y. Zhu, and R. Martín-Martín. What matters in learning from offline human demonstrations for robot manipulation. *arXiv preprint arXiv:2108.03298*, 2021.
- [44] E. Mansimov and K. Cho. Simple nearest neighbor policy method for continuous control tasks, 2018. URL <https://openreview.net/forum?id=ByL48G-AW>.
- [45] F. Mentzer, D. Minnen, E. Agustsson, and M. Tschannen. Finite scalar quantization: Vq-vae made simple. *arXiv preprint arXiv:2309.15505*, 2023.
- [46] Octo Model Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, C. Xu, J. Luo, T. Kreiman, Y. Tan, P. Sanketi, Q. Vuong, T. Xiao, D. Sadigh, C. Finn, and S. Levine. Octo: An open-source generalist robot policy. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, 2024.
- [47] A. Padalkar, A. Pooley, A. Jain, A. Bewley, A. Herzog, A. Irpan, A. Khazatsky, A. Rai, A. Singh, A. Brohan, et al. Open x-embodiment: Robotic learning datasets and rt-x models. *arXiv preprint arXiv:2310.08864*, 2023.
- [48] N. D. Palo and E. Johns. Dinobot: Robot manipulation via retrieval and alignment with vision foundation models. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [49] E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.
- [50] K. Pertsch, Y. Lee, and J. Lim. Accelerating reinforcement learning with learned skill priors. In *Conference on robot learning*, pages 188–204. PMLR, 2021.
- [51] D. A. Pomerleau. Alvin: An autonomous land vehicle in a neural network. *Advances in neural information processing systems*, 1, 1988.
- [52] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- [53] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, and I. Sutskever. Language models are unsupervised multitask learners. 2019.
- [54] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [55] S. Reed, K. Zolna, E. Parisotto, S. G. Colmenarejo, A. Novikov, G. Barth-Maron, M. Gimenez, Y. Sulsky, J. Kay, J. T. Springenberg, T. Eccles, J. Bruce, A. Razavi, A. Edwards, N. Heess, Y. Chen, R. Hadsell, O. Vinyals, M. Bordbar, and N. de Freitas. A generalist agent, 2022.
- [56] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- [57] N. M. Shafiullah, Z. Cui, A. A. Altanzaya, and L. Pinto. Behavior transformers: Cloning k modes with one stone. *Advances in neural information processing systems*, 35:22955–22968, 2022.
- [58] L. X. Shi, J. J. Lim, and Y. Lee. Skill-based model-based reinforcement learning. *arXiv preprint arXiv:2207.07560*, 2022.
- [59] A. Singh, H. Liu, G. Zhou, A. Yu, N. Rhinehart, and S. Levine. Parrot: Data-driven behavioral priors for reinforcement learning. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=Ysuv-W0FeKR>.
- [60] K. Sohn, H. Lee, and X. Yan. Learning structured output representation using deep conditional generative models. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. URL https://proceedings.neurips.cc/paper_files/paper/2015/file/8d55a249e6baa5c06772297520da2051-Paper.pdf.

- [61] G. Swamy, S. Choudhury, J. A. Bagnell, and Z. S. Wu. Causal imitation learning under temporally correlated noise, 2022.
- [62] G. Team. Gemini: A family of highly capable multimodal models, 2024.
- [63] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, D. Bikel, L. Blecher, C. C. Ferrer, M. Chen, G. Cucurull, D. Esiobu, J. Fernandes, J. Fu, W. Fu, B. Fuller, C. Gao, V. Goswami, N. Goyal, A. Hartshorn, S. Hosseini, R. Hou, H. Inan, M. Kardas, V. Kerkez, M. Khabsa, I. Kloumann, A. Korenev, P. S. Koura, M.-A. Lachaux, T. Lavril, J. Lee, D. Liskovich, Y. Lu, Y. Mao, X. Martinet, T. Mihaylov, P. Mishra, I. Molybog, Y. Nie, A. Poulton, J. Reizenstein, R. Rungta, K. Saladi, A. Schelten, R. Silva, E. M. Smith, R. Subramanian, X. E. Tan, B. Tang, R. Taylor, A. Williams, J. X. Kuan, P. Xu, Z. Yan, I. Zarov, Y. Zhang, A. Fan, M. Kambadur, S. Narang, A. Rodriguez, R. Stojnic, S. Edunov, and T. Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023.
- [64] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu. Wavenet: A generative model for raw audio. In *Arxiv*, 2016. URL <https://arxiv.org/abs/1609.03499>.
- [65] A. Van Den Oord, O. Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- [66] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [67] W. Wan, Y. Zhu, R. Shah, and Y. Zhu. Lotus: Continual imitation learning for robot manipulation through unsupervised skill discovery. *arXiv preprint arXiv:2311.02058*, 2023.
- [68] Z. Xian, N. Gkanatsios, T. Gervet, T.-W. Ke, and K. Fragkiadaki. Chaineddiffuser: Unifying trajectory diffusion and keypose prediction for robotic manipulation. In *7th Annual Conference on Robot Learning*, 2023.
- [69] J. Yu, X. Li, J. Y. Koh, H. Zhang, R. Pang, J. Qin, A. Ku, Y. Xu, J. Baldridge, and Y. Wu. Vector-quantized image modeling with improved VQGAN. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=pfNyExj7z2>.
- [70] T. Yu, D. Quillen, Z. He, R. Julian, K. Hausman, C. Finn, and S. Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Conference on robot learning*, pages 1094–1100. PMLR, 2020.
- [71] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu. 3d diffusion policy. *arXiv preprint arXiv:2403.03954*, 2024.
- [72] S. Zhang, L. Dong, X. Li, S. Zhang, X. Sun, S. Wang, J. Li, R. Hu, T. Zhang, F. Wu, et al. Instruction tuning for large language models: A survey. *arXiv preprint arXiv:2308.10792*, 2023.
- [73] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn. Learning fine-grained bimanual manipulation with low-cost hardware, 2023.
- [74] R. Zheng, C.-A. Cheng, H. Daumé III, F. Huang, and A. Kolobov. Prize: Learning temporal action abstractions as a sequence compression problem. *arXiv preprint arXiv:2402.10450*, 2024.
- [75] R. Zheng, X. Wang, Y. Sun, S. Ma, J. Zhao, H. Xu, H. Daumé III, and F. Huang. Taco: Temporal latent action-driven contrastive loss for visual reinforcement learning. *Advances in Neural Information Processing Systems*, 36, 2024.
- [76] zhengyao jiang, T. Zhang, M. Janner, Y. Li, T. Rocktäschel, E. Grefenstette, and Y. Tian. Efficient planning in a compact latent action space. In *3rd Offline RL Workshop: Offline RL as a "Launchpad"*, 2022. URL <https://openreview.net/forum?id=pVBETTS2av>.
- [77] X. Zou, J. Yang, H. Zhang, F. Li, L. Li, J. Wang, L. Wang, J. Gao, and Y. J. Lee. Segment everything everywhere all at once. *Advances in Neural Information Processing Systems*, 36, 2024.

Appendix

A Website

For further results and videos please see our website. <https://quest-model.github.io>

B Experiment Details

B.1 Hyperparameters:

We present hyperparameters in the following tables:

Table 3: Stage 1 Parameters

Parameter	Value
encoder dim	256
decoder dim	256
sequence length (T)	16/32
encoder heads	4
encoder layers	2
decoder heads	4
decoder layers	4
attention dropout	0.1
fsq level	[8, 5, 5, 5]
conv layers	3
downsampling factor	2/4

Table 4: Stage 2 Parameters

Parameter	Value
vocab size	1000
block size (n)	8
number of layers	6
number of heads	6
embedding dimension	384
attention dropout	0.1
beam size	5
temperature	1.0
decoder loss scale	0/10
execution horizon (T_a)	8
observation history	1

B.2 Architecture Implementation:

For vision encoder we used a shallow Convolutional Neural Network (CNN), consisting of the first four layers of ResNet18 [27] followed by a spatial softmax [36]. In encoder, we use causal convolution layers from [30]. For transformer blocks, we used the transformers library from hugging face <https://huggingface.co/docs/transformers/> with appropriate masking for ensuring causality.

B.3 Baseline Implementation:

To ensure fair comparison of different model architectures, we use same input modalities and same observation & task encoders for all baselines, for both LIBERO and Metaworld experiments. VQ-BeT needs a goal image, we instead give it task embedding as goal. Same as QueST, we concatenate observation embeddings for all modalities at any timestep and project them to respective model’s hidden dimension.

Depending on the dataset, we also tune some key hyperparameters for the baselines and present the results for best performing ones.

1. **ResNet-T:** Transformer trunk’s hidden dimension and number of layers determines the model capacity. Original implementation [38] uses the hidden dimension of 64 with 4 layers. We observed improved performance for the hidden dimension of 256 with 6 layers and hence report all results for that. As per original implementation we use an observation history of 10 timesteps.
2. **Diffusion Policy:** The model capacity is determined by hidden dimension of U-Net layers. Most widely used implementations use [256, 512, 1024], we ablate a larger model with [256, 256, 512, 512, 1024] but did not observe any performance gains. We also ablate prediction (T) and execution horizon (T_a) with 16, 32 and 8, 16 respectively and observed best performance for $T = 32, T_a = 16$ on LIBERO and $T = 16, T_a = 8$ for MetaWorld. As per original paper ablations [15] an observation history of 1 was used.

3. **VQ-BeT:** Since LIBERO and MetaWorld are larger datasets as compared to the benchmarks in original VQ-BeT paper, we ablate some parameters to increase the model capacity. Specifically, the stage 1 encoder by default is a single MLP layer of dimension 128. We ablate this with 2, 4 layers and with 256, 512 dimensions but observed worse reconstruction loss with increase in capacity. We use residual-VQ configuration of $32/2 \approx 1024$ sized codebook which is close to the codebook size of 1000 for QueST. We use an observation window size of 10 and ablate the action window size (T) with 1, 5, 32. On LIBERO, the performance was lowest for $T = 1$, and highest for $T = 5$. VQ-BeT maps the whole input sequence to just one embedding leading to extreme compression for larger sequence length and thus performs worse with $T = 32$.

B.4 Compute:

The models are implemented in PyTorch. For all our experiments we use a server consisting of 8 Nvidia RTX 1080Ti 10GB memory each. And all our models easily fit on one GPU for training.

B.5 Discussion on Baselines

Many recent works use discrete latent variable models as a mechanism to learn shared abstractions over continuous low-level skills. QueST, VQ-BeT and PRISE all do this and perform two- staged learning. However, in this work we propose several important architecture choices leading to QueST’s strong performance. Specifically, QueST encodes actions to a sequence of n encodings (skill tokens) using a novel autoencoder that captures temporal correlation within an action sequence with causal convolution and masked self-attention layers.

Like QueST, VQ-BeT performs temporal abstraction by encoding a sequence of actions into one state-independent latent vector. It concatenates input actions and encodes the sequence to just one single latent encoding using an MLP which does not explicitly model any temporal correlation. A sampled action sequence might contain multiple motion primitives of variable length and start point, and capturing them with a single encoding is limiting as it restricts abstraction at different levels of granularity. This is validated by the poorer (-14%) performance of VQ-BeT with a larger chunk size of 32. QueST flexibly captures this variability within n encodings using its encoder, specifically designed to model temporal correlation within input actions. With an effective codebook size of C , QueST’s effective latent space is C^n while that of VQ-BeT is C . Its relatively small latent space limits its expressive capacity, and thus limits its ability to learn sharable representations between tasks, as evidenced by its worse few-shot performance and reliance on continuous offset prediction (whereas QueST can achieve high success rate using only the output from its action decoder).

PRISE performs temporal abstraction by learning discrete codes for state-action pairs and using BPE to group common token sequences into higher-level skills. However, BPE is known to suffer with evolving language leading to a suboptimal character-level tokenization, and it might struggle to effectively encode new action sub-sequences that the model has not seen before (eg. stitching sequences in between two tasks in LIBERO-LONG). This is primarily why decoder finetuning is necessary in PRISE. On the contrary, QueST is more end-to-end as it lets the encoder handle temporal abstraction in the encoding phase itself and gracefully encodes new action-sequences as combinations of its learned latent codes. Our few-shot results without decoder-finetuning shows that such a unified approach learns more generalized abstractions than the baselines and can more effectively represent new action sequences in unseen tasks.

Despite tuning the hyperparameters, the multitask performance of ACT is very low (54%) on LIBERO. Moreover, as per BAKU [26], ACT’s performance drastically reduces when number of demos are reduced to 35-per-task in MetaWorld. This suggests ACT is very data-hungry and hence struggles in low-data regime. Another major issue with ACT is tuning the action chunk size to which the method is very sensitive (-14% in metaworld fewshot with chunk size of 32). While we report QueST results for chunk size 32, we did observe very little variation ($<1\%$ in LIBERO) with size 16,48 and 64, indicating robustness to this hyperparameter.

C Discussion on Ablations

For aiding this discussion we present the ablation results again in table 5 and table 6 below.

	VQ	Obs. Cond.	Mirror Dec.	Ours
LIBERO-90	81.2 ± 0.6	81.9 ± 1.1	86.3 ± 0.9	88.6 ± 0.4
Few Shot	62.5 ± 2.0	61.3 ± 2.2	45.4 ± 2.0	68.8 ± 1.7

Table 5: Success rates after ablating design details of QueST.

- Replacing FSQ with VQ still outperforms VQ-BeT in few-shot setting suggesting that QueST’s superior performance is not only due to a better quantization scheme but also due to its architecture that flexibly maps an input sequence to multiple embeddings and allows for efficient transfer.
- It’s tempting to ground the mapping between z-tokens and actions with observation tokens with an intuition that z-tokens will define a coarse set of actions and observation tokens will aid finer action decoding. But we observe worse performance with this. We hypothesize that the reconstruction objective forces encoder and decoder for most optimal quantization at the bottleneck layer but with extra observation information the decoder might focus more on observation tokens in turn hurting the quantization. This observation goes hand-in-hand with a closely related prior work SPiRL[58] that tried same ablation and found that state conditioned decoder hurts downstream RL.
- We observe a poorer performance in both multitask and few-shot settings with a conventional stage 1 autoencoder. This validates the QueST’s cross-attention architecture that allows for attending to all z-tokens and maintaining causality at the same time.

	Non Causal ϕ_θ	Non Causal ψ_θ	Fully Non Causal	Ours
LIBERO-90	82.0 ± 1.6	85.1 ± 1.8	78.5 ± 0.5	88.6 ± 0.4
Few Shot	58.8 ± 3.0	61.6 ± 2.5	56.1 ± 1.8	68.8 ± 1.7

Table 6: Success rates after ablating the causality in QueST.

- We observe that a fully-causal stage-1 is most optimal and a non-causal decoder does not hurt as much as a non-causal encoder does. This can be explained with a simplistic setting where the input to stage-1 are 2D trajectories of a point agent. Consider an anti-clockwise circular trajectory and an S-shaped one where the first half of the later overlaps with the first half (semi-circle) of the former. When both of these trajectory sequences are inputted to the stage-1, a non-causal encoder will assign distinct sequences of z-tokens for both trajectories. But a causal encoder will assign same sequence of z-tokens for the first half of both trajectories and distinct to later parts. This allows the model to re-use the z-tokens corresponding to a semi-circle for creating other shaped-trajectories that has semi-circle in them for example C-shaped or infinity-shaped trajectories.

	Frozen ψ_θ	Finetuned ψ_θ	
		loss scale 10	loss scale 100
Few Shot	66.0 ± 3.6	68.8 ± 1.7	66.0 ± 1.0

Table 7: Success rates for decoder finetuning settings in few-shot IL.

- Table 7 illustrates the impact of decoder finetuning in LIBERO-LONG fewshot IL setting. QueST outperforms all baselines even without finetuning the decoder. Finetuning decoder should not be necessary in this setting, as LIBERO-LONG tasks are combination of two tasks from LIBERO-90 (pretraining set). This highlights QueST’s effectiveness in stitching trajectories using its learned skill-space. We report the finetuning results in the main paper, as they exhibit better performance.

D Additional Results

D.1 Fewshot IL

Fewshot Evaluation Protocol: In finetuning phase, we finetune ResNet-T, VQ-BeT & QueST for 100 epochs and ACT & Diffusion Policy for 200 epochs. For each task in MetaWorld, we evaluate

each method across 10 evenly spaced checkpoints for 5 seeds on 50 distinct initial states and report the results corresponding to the best performing checkpoint. For Libero, we found the final checkpoint to perform best for all methods and hence report results corresponding to it across 9 seeds.

Table 8: LIBERO 5-shot IL success rates across unseen 10 tasks. Results across 9 random seeds.

Task ID	ResNet-T	ACT	Diffusion Policy	PRISE	VQ-BeT	QueST
1	53.8 ± 11.6	20.0 ± 6.0	32.0 ± 5.5	26.7 ± 6.4	16.5 ± 11.9	56.4 ± 5.5
2	65.7 ± 16.9	33.3 ± 13.1	57.8 ± 5.9	48.3 ± 9.4	62.2 ± 15.8	82.9 ± 4.7
3	70.2 ± 7.8	67.7 ± 6.2	76.4 ± 7.2	70.0 ± 0.0	52.7 ± 7.4	66.7 ± 6.7
4	75.8 ± 7.6	70.3 ± 6.2	98.2 ± 1.7	78.3 ± 8.8	45.3 ± 7.7	88.0 ± 3.1
5	26.7 ± 11.9	35.0 ± 4.1	44.7 ± 6.0	45.0 ± 10.8	30.3 ± 13.0	42.0 ± 7.1
6	86.9 ± 4.9	68.3 ± 6.5	37.1 ± 3.7	90.0 ± 4.1	48.7 ± 17.2	92.4 ± 4.4
7	24.4 ± 8.4	15.0 ± 0.0	14.9 ± 6.7	25.0 ± 4.1	45.5 ± 8.2	58.2 ± 6.1
8	23.7 ± 12.1	26.7 ± 7.0	6.2 ± 3.2	45.0 ± 8.1	33.3 ± 9.3	47.1 ± 8.6
9	-	-	55.0 ± 4.0	-	15.0 ± 7.0	46.6 ± 6.2
10	-	-	68.3 ± 6.2	-	25.0 ± 0.0	65.0 ± 12.2

Table 9: MetaWorld 5-shot IL success rates across 5 unseen tasks. Results across 5 random seeds.

Task ID	ResNet-T	ACT	Diffusion Policy	PRISE	VQ-BeT	QueST
box-close-v2	63.2 ± 5.2	67.2 ± 5.2	68.0 ± 1.6	60.8 ± 6.6	75.3 ± 9.6	84.0 ± 7.3
disassemble-v2	68.8 ± 2.0	83.2 ± 3.2	81.3 ± 3.8	74.1 ± 7.3	92.7 ± 1.9	76.4 ± 26.0
hand-insert-v2	37.2 ± 4.1	53.2 ± 3.7	39.3 ± 1.9	60.0 ± 5.0	48.0 ± 6.5	49.6 ± 6.4
pick-place-wall-v2	42.8 ± 3.7	74.4 ± 6.9	70.7 ± 5.2	71.7 ± 5.7	65.3 ± 1.9	76.8 ± 11.4
stick-pull-v2	58.0 ± 8.8	76.0 ± 3.6	71.3 ± 1.9	67.5 ± 5.6	62.0 ± 11.4	72.8 ± 11.1

D.2 Multitask IL

Table 10: LIBERO-90 multitask IL success rates across 90 tasks. Results across 4 random seeds.

Task ID	ResNet-T	ACT	Diffusion Policy	PRISE	VQ-BeT	QueST
1	1.00	0.90	0.99	0.80	1.00	1.00
2	0.96	0.30	0.98	0.35	0.94	0.97
3	0.96	0.50	0.99	0.70	0.97	0.91
4	0.74	0.22	0.91	0.50	0.99	0.94
5	0.95	0.58	0.93	0.45	0.95	0.98
6	0.91	0.39	0.99	0.65	0.98	0.98
7	0.95	0.29	0.94	0.50	0.86	0.93
8	0.96	0.72	0.90	0.95	0.80	0.99
9	0.74	0.41	0.93	0.60	0.75	0.93
10	0.97	0.65	0.91	0.35	0.83	0.90
11	0.97	0.82	0.98	0.95	0.96	0.97
12	0.90	0.73	0.94	0.95	0.80	0.94
13	0.82	0.62	0.81	0.20	0.87	0.76
14	0.86	0.72	0.94	0.40	0.49	0.71

Task ID	ResNet-T	ACT	Diffusion Policy	PRISE	VQ-BeT	QueST
15	0.87	0.49	0.95	0.35	0.46	0.58
16	0.97	0.86	0.99	0.75	0.98	0.96
17	0.72	0.40	0.89	0.40	0.53	0.80
18	0.79	0.20	0.76	0.15	0.80	0.67
19	0.93	0.75	0.99	0.30	0.91	1.00
20	0.87	0.41	0.98	0.65	0.68	0.92
21	1.00	0.82	0.99	1.00	0.96	1.00
22	0.90	0.44	0.97	0.30	0.91	0.93
23	0.97	0.75	0.99	0.85	0.95	0.92
24	0.75	0.11	0.85	0.05	0.69	0.82
25	0.97	0.44	0.99	0.95	0.94	1.00
26	0.97	0.85	1.00	0.90	0.85	0.99
27	0.72	0.14	0.88	0.55	0.50	0.52
28	0.72	0.20	0.86	0.05	0.45	0.68
29	1.00	0.68	1.00	1.00	0.97	1.00
30	1.00	0.19	1.00	1.00	0.92	0.97
31	0.91	0.83	0.96	0.50	0.85	0.90
32	0.99	0.90	1.00	0.85	0.88	0.99
33	0.57	0.20	0.58	0.20	0.37	0.67
34	0.85	0.56	0.84	0.30	0.87	0.98
35	0.93	0.52	0.97	0.80	0.98	0.92
36	0.97	0.67	0.99	0.75	0.98	0.97
37	0.85	0.24	0.97	0.25	0.73	0.74
38	0.78	0.41	0.91	0.30	0.90	0.62
39	0.86	0.32	0.90	0.20	0.90	0.88
40	0.96	0.35	0.98	0.85	0.90	0.93
41	0.90	0.27	0.79	0.50	0.91	0.92
42	1.00	0.74	1.00	0.55	0.89	1.00
43	0.98	0.41	0.99	0.80	0.97	0.98
44	0.80	0.39	0.89	0.40	0.83	0.93
45	0.99	0.83	1.00	0.85	0.99	0.98
46	0.97	0.60	1.00	0.55	0.91	0.99
47	0.75	0.37	0.31	0.35	0.65	0.91
48	0.87	0.27	0.53	0.25	0.88	0.98
49	0.90	0.55	0.96	0.65	0.48	0.95
50	0.88	0.54	0.82	0.65	0.60	0.99
51	0.80	0.33	0.28	0.40	0.87	0.88
52	0.79	0.28	0.00	0.10	0.91	0.75
53	0.74	0.33	0.34	0.30	0.84	0.82
54	0.88	0.64	0.73	0.60	0.79	0.87
55	0.83	0.51	0.77	0.50	0.95	0.93

Task ID	ResNet-T	ACT	Diffusion Policy	PRISE	VQ-BeT	QueST
56	0.85	0.62	0.49	0.35	0.92	0.83
57	0.99	0.64	1.00	0.80	1.00	0.97
58	0.95	0.57	1.00	0.50	1.00	0.99
59	0.84	0.56	0.78	0.20	0.98	0.95
60	0.94	0.68	0.89	0.65	0.91	1.00
61	0.91	0.95	0.90	0.80	0.98	1.00
62	0.96	0.75	0.58	0.85	0.99	0.81
63	0.70	0.43	0.38	0.40	0.84	0.78
64	0.73	0.04	0.41	0.40	0.38	0.78
65	0.73	0.16	0.75	0.15	0.68	0.85
66	0.76	0.45	0.65	0.15	0.84	0.79
67	0.84	0.72	0.66	0.30	0.87	0.93
68	0.78	0.73	0.44	0.55	0.74	0.85
69	0.83	0.68	0.59	0.85	0.90	0.93
70	0.88	0.56	0.57	0.90	0.93	0.89
71	0.90	0.52	0.92	0.55	0.97	0.92
72	0.85	0.52	0.98	0.35	0.85	0.94
73	0.89	0.59	0.86	0.60	0.84	0.98
74	0.72	0.18	0.61	0.30	0.33	0.70
75	0.77	0.45	0.38	0.45	0.95	0.95
76	0.64	0.22	0.21	0.25	0.30	0.61
77	0.89	0.70	0.35	0.65	0.70	0.89
78	0.57	0.46	0.14	0.80	0.85	0.97
79	0.63	0.28	0.06	0.45	0.68	0.86
80	0.73	0.59	0.01	0.30	0.87	0.98
81	0.65	0.53	0.08	0.30	0.44	0.70
82	0.63	0.24	0.54	0.35	0.61	0.70
83	0.80	0.56	0.49	0.80	0.89	0.94
84	0.55	0.35	0.47	0.55	0.43	0.75
85	0.70	0.74	0.79	0.75	0.93	0.92
86	0.69	0.53	0.13	0.75	0.47	0.89
87	0.84	0.65	0.98	0.95	0.86	0.92
88	0.82	0.54	0.96	0.65	0.87	0.97
89	0.91	0.77	0.70	0.55	0.96	0.97
90	0.80	0.29	0.91	0.85	0.89	0.56

Table 11: MetaWorld multitask IL success rates across 45 tasks. Results across 5 random seeds.

Task ID	ResNet-T	ACT	Diffusion Policy	VQBeT	QueST
assembly-v2	0.73	0.97	0.88	0.82	1.00
basketball-v2	0.76	0.80	0.78	0.82	0.68

Task ID	ResNet-T	ACT	Diffusion Policy	VQBeT	QueST
bin-picking-v2	0.89	1.00	0.96	0.20	0.94
button-press-topdown-v2	1.00	1.00	1.00	1.00	1.00
button-press-topdown-wall-v2	1.00	1.00	1.00	1.00	1.00
button-press-v2	1.00	1.00	1.00	1.00	1.00
button-press-wall-v2	1.00	1.00	0.98	0.98	0.98
coffee-button-v2	1.00	1.00	1.00	1.00	1.00
coffee-pull-v2	0.90	0.92	0.96	0.82	0.98
coffee-push-v2	0.89	0.96	0.86	0.94	0.90
dial-turn-v2	0.98	0.99	1.00	1.00	1.00
door-close-v2	1.00	1.00	1.00	1.00	1.00
door-lock-v2	1.00	0.99	1.00	1.00	1.00
door-open-v2	0.96	0.95	0.96	0.94	0.94
door-unlock-v2	1.00	1.00	1.00	1.00	1.00
drawer-close-v2	1.00	1.00	1.00	1.00	1.00
drawer-open-v2	1.00	1.00	1.00	1.00	1.00
faucet-close-v2	1.00	1.00	1.00	1.00	1.00
faucet-open-v2	1.00	1.00	1.00	1.00	1.00
hammer-v2	0.95	1.00	0.98	1.00	0.94
handle-press-side-v2	1.00	1.00	1.00	1.00	1.00
handle-press-v2	1.00	1.00	1.00	1.00	1.00
handle-pull-side-v2	0.69	0.94	0.78	0.74	0.98
handle-pull-v2	1.00	1.00	1.00	1.00	1.00
lever-pull-v2	0.94	0.93	0.84	0.80	0.92
peg-insert-side-v2	0.81	0.94	0.90	0.76	0.86
peg-unplug-side-v2	0.88	0.91	0.88	0.92	0.90
pick-out-of-hole-v2	0.62	0.89	0.74	0.34	0.76
pick-place-v2	0.67	0.71	0.76	0.74	0.78
plate-slide-back-side-v2	1.00	1.00	1.00	1.00	1.00
plate-slide-back-v2	1.00	1.00	1.00	1.00	1.00
plate-slide-side-v2	0.98	1.00	0.98	0.98	1.00
plate-slide-v2	1.00	1.00	1.00	1.00	1.00
push-back-v2	0.72	0.64	0.76	0.64	0.80
push-v2	0.84	0.90	0.84	0.76	0.92
push-wall-v2	0.92	0.98	0.94	0.94	1.00
reach-v2	0.39	0.37	0.32	0.28	0.36
reach-wall-v2	0.49	0.47	0.52	0.36	0.42
shelf-place-v2	0.65	0.85	0.66	0.76	0.88
soccer-v2	0.42	0.25	0.42	0.36	0.52
stick-push-v2	0.75	1.00	0.96	0.94	0.96
sweep-into-v2	0.90	0.92	0.88	0.90	0.84
sweep-v2	0.98	1.00	0.98	1.00	1.00

Task ID	ResNet-T	ACT	Diffusion Policy	VQBeT	QueST
window-close-v2	1.00	1.00	1.00	1.00	1.00
window-open-v2	1.00	1.00	1.00	1.00	1.00

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The main claims made in the abstract and introduction are that we achieve stronger performance than several state of the art baselines. We show several results justifying this claim in Figures 3 and 2.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We stated our limitation in the last section of the main paper.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: We don't have theoretical results in the paper.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We show experimental details in the paper and describe our architecture in great detail. In the future we will release code.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We include codes in the supplementary materials.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We have the training details in the appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We report the averaged metrics on different datasets, and the difference is significant.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: All the experiments can be run on one A6000 GPU.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: This paper conforms with the ethics code.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: We discuss the broader impact in the main paper.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: all the models and datasets used in this paper are properly credited and licenses are respected.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new datasets/models are proposed in this paper.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer:
answerNA

Justification: [NA]

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.