

# BEAST: Efficient Tokenization of B-Splines Encoded Action Sequences for Imitation Learning

Hongyi Zhou<sup>†</sup>\* Weiran Liao<sup>†</sup> Xi Huang<sup>†</sup> Yucheng Tang<sup>†</sup> Fabian Otto<sup>§</sup>  
Xiaogang Jia<sup>†</sup> Xinkai Jiang<sup>†</sup> Simon Hilber<sup>†</sup> Ge Li<sup>†</sup> Qian Wang<sup>†</sup>  
Ömer Erdiñç Yağmurlu<sup>†</sup> Nils Blank<sup>†</sup> Moritz Reuss<sup>†</sup> Rudolf Lioutikov<sup>†</sup>  
<sup>†</sup> Karlsruhe Institute of Technology      <sup>§</sup> Microsoft Research

## Abstract:

We present the **B-spline Encoded Action Sequence Tokenizer (BEAST)**, a novel action tokenizer that encodes action sequences into compact discrete or continuous tokens using B-spline. In contrast to existing action tokenizers based on vector quantization or byte pair encoding, BEAST requires no separate tokenizer training and consistently produces tokens of uniform length, enabling fast action sequence generation via parallel decoding. Leveraging our B-spline formulation, BEAST inherently ensures generating smooth trajectories without discontinuities between adjacent segments. We extensively evaluate BEAST by integrating it with three distinct model architectures: a Variational Autoencoder (VAE) with continuous tokens, a decoder-only Transformer with discrete tokens, and Florence-2, a Vision-Language Model with an encoder-decoder architecture, demonstrating BEAST’s compatibility and scalability with large pretrained models. We evaluate BEAST across three established benchmarks consisting of **166 simulated tasks** and on three distinct robot settings with a total of **8 real-world tasks**. Experimental results demonstrate that BEAST (i) significantly reduces both training and inference computational costs, and (ii) consistently generates smooth, high-frequency control signals suitable for continuous control tasks while (iii) reliably achieves competitive task success rates compared to state-of-the-art methods. Videos and code are available at [the project page](#).

## 1 Introduction

Imitation learning has emerged as a powerful paradigm for training robots to perform complex tasks by learning from human demonstrations [1]. Early works [2, 3] in this field primarily focused on predicting single-step actions based on the current observation. However, recent research [4] highlights the importance of learning action sequences to capture the temporal coherence inherent in human demonstrations. Moreover, by modeling action sequences, we can reduce compounding errors [5] and create task demonstrations that more closely align with human methods [6]. Given the success of autoregressive next-token prediction models in natural language processing and other domains [7–9], it is compelling to explore similar techniques for modeling action sequences, leveraging their ability to predict and generate coherent sequences effectively.

Despite growing interest in this area, effective strategies to create action sequences of discrete tokens remain underexplored. Existing approaches often focus on single-step tokenization [10–12], vector quantization [13–15], or compression-based schemes [16]. However, they require training separate encoder-decoder networks for the tokenizer [13, 17] or produce variable-length token sequences for inputs of the same duration [16], which complicates applying fast token generation techniques such as parallel decoding [18]. Furthermore, existing action tokenizers do not consider the smooth transitions between subsequent action chunks, which could result in undesired jumps at transition.

\*Correspondence to hongyi.zhou@kit.edu

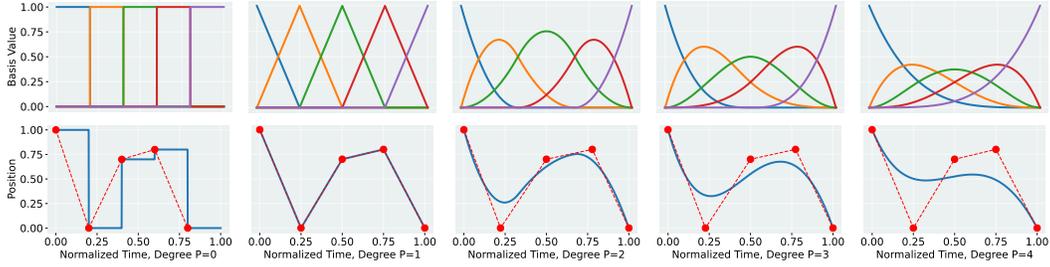


Figure 1: **From left to right:** Clamped B-Spline Basis  $P = 0, 1, 2, 3, 4$  (top) and their generated trajectories (Bottom). Given the same **control points**, a higher degree will lead to smoother trajectories. All generated trajectories start exactly from the first control point and end at the last control point. Notably, action chunk is conceptually equivalent to B-Splines of 0-th degree, i.e., split-wise constants, as shown in the leftmost subplots. This relation is explained in details later in Section 3.1.

To address these challenges, we propose the B-spline Encoded Action Sequence Tokenizer (**BEAST**), a novel tokenizer that represents continuous action sequences using B-splines [19]. BEAST offers versatility, allowing for effective integration with both discrete and continuous tokens. Different from tokenizers based on the vector quantization [13–15], it does not require additional tokenizer training. BEAST compresses action trajectories into fixed-length token sequences enabling efficient parallel decoding for faster token generation, requiring  $4 - 8 \times$  fewer tokens than binning-based tokenization. By using B-spline encoded control points as discrete tokens, BEAST ensures the generation of smooth action chunks, as well as the continuous connection between consecutive chunks.

**Our contributions are:** 1) We introduce BEAST, a novel B-spline-based tokenizer designed for modeling continuous action sequences. 2) We demonstrate the versatility of BEAST by integrating it into diverse model architectures that accommodate both continuous and discrete objectives. 3) We conduct extensive evaluations of simulated and real-world robotic tasks, showcasing its effectiveness. 4) We perform thorough ablation studies to assess the impact of various design choices.

## 2 Preliminaries

**Problem Formulation.** Our goal is to train a policy  $\pi(\mathbf{a}_{1:T} | \mathbf{s})$  that capable of mapping a given state  $\mathbf{s}$  to a corresponding sequence of actions  $\mathbf{a}_{1:T}$  which has  $T$  time steps and  $D$  Degrees of Freedom (DoF). To make this sequence prediction problem compatible with discrete generative models, we first transform the continuous action sequence into a sequence of discrete tokens. The goal of action sequence tokenization is to obtain a discrete token sequence  $\bar{\mathbf{v}}_{1:J}$ , where each token belongs to a vocabulary  $\bar{\mathcal{V}}$  with size  $|\bar{\mathcal{V}}|$ , by defining a transformation **tokenizer** :  $\mathbf{a}_{1:T} \rightarrow \bar{\mathbf{v}}_{1:J}$ ,

**B-Splines** (Basis Splines) [20] are widely used in the field of computer graphics and computer-aided design. A B-Spline curve  $y$  is formulated through a linear basis function representation

$$\text{1-DoF B-Spline: } y(u) = \sum_{n=0}^{N-1} \Phi_n^P(u) c_n = \Phi^P(u) \mathbf{c}, \quad 0 \leq P < N, \quad u \in [k_0, k_M], \quad (1)$$

where  $\mathbf{c}$  are  $N$  control points and  $u$  is a continuous parameter, often interpreted as *normalized time*. The basis functions  $\Phi^P(u) = [\Phi_0^P(u), \dots, \Phi_{N-1}^P(u)]$  are  $N$  polynomial basis functions of  $P$ -th degree. These basis functions are defined over  $M$  intervals determined by  $M + 1$  knots in a vector  $[k_0, \dots, k_M]$ , and it satisfies  $M = N + P$  [20]. Typically, the knot vector is normalized such that  $k_0 = 0$  and  $k_M = 1$ . The basis functions  $\Phi_n^P(u)$  are recursively computed using the *Cox-de Boor recursion* [21]. We denote all recursive degrees<sup>2</sup> as  $q = 0 : P$ . For  $q = 0$ , the basis functions are

<sup>2</sup>The B-Spline degree  $P$  differs from the recursive degree  $q$ . Trajectories are represented by basis functions of degree  $P$ , while lower recursive degree  $q$  serve as intermediate representations in the recursive process.

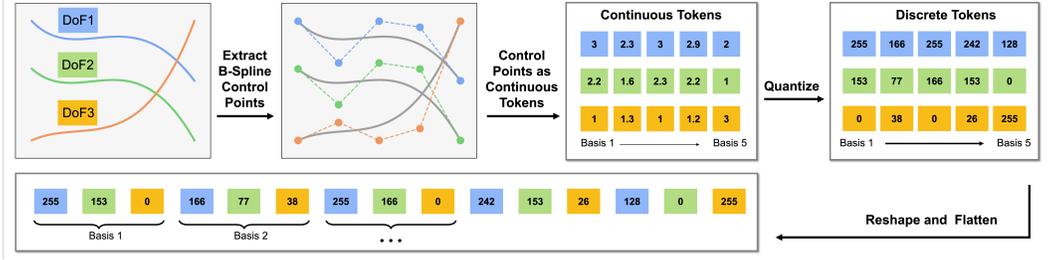


Figure 2: **Overview of the BEAST Encoding Pipeline:** Given a normalized action sequence, the BEAST pipeline first uses linear regression to extract continuous-valued control points, forming control point matrices that serve as intermediate continuous representations. These matrices are then quantized uniformly into discrete values within the range  $[0, 255]$  and subsequently flattened to produce discrete action tokens for auto-regressive next-token prediction or parallel prediction.

defined as piecewise constant and recursively using the  $(q - 1)$ -th degree basis for  $q > 0$  with

$$\text{piecewise constant: } \Phi_n^0(u) = \begin{cases} 1 & \text{if } k_n \leq u < k_{n+1}, \text{ and} \\ 0 & \text{otherwise.} \end{cases} \quad (2)$$

$$\text{recursive: } \Phi_n^q(u) = k_n^{q-1} \Phi_n^{q-1}(u) + (1 - k_{n+1}^{q-1}) \Phi_{n+1}^{q-1}(u), \quad (3)$$

where  $k_n^{q-1} = (u - k_n)/(k_{n+q} - k_n)$ .

**Clamped B-Spline.** In this work, we employ the *clamped uniform B-Spline*, where the first and last  $P + 1$  knots are repeated to ensure that the resulting curve starts at the first control point and ends at the last control point. In Figure 1, we demonstrate the resulting basis functions of degrees from  $P = 0$  to  $P = 4$ , together with their generated trajectories, given the same five control points. Clamped uniform B-splines are particularly suited for trajectory generation due to their smoothness, compact representation, and local support, where each control point only affects the curve locally.

**Parallel Decoding.** Unlike autoregressive generation, which predicts tokens sequentially and thus requires  $K$  forward passes for a sequence of length  $K$ , *parallel decoding* [18] enables the prediction of the entire output sequence in a single forward pass. This is achieved by feeding the model with  $K$  empty token embeddings and replacing the causal attention mask with a bidirectional mask, allowing the decoder to infer the entire sequence simultaneously. OpenVLA-OFT [18] leverages this approach for action sequence generation. In this work, we adopt the parallel decoding strategy to predict all BEAST tokens in a single pass, improving the inference efficiency without sacrificing accuracy.

### 3 B-Spline Encoded Action Sequence Tokenizer

In this section, we first describe how BEAST utilizes B-Spline to construct an efficient action sequence tokenizer that converts action sequences into either continuous or discrete action tokens. We then explain how smooth transitions between consecutive action sequences are achieved by enforcing the initial conditions of clamped B-splines. Finally, we discuss strategies for efficient integrating BEAST with various model architectures that predict discrete or continuous tokens.

#### 3.1 Action Sequence Tokenization with B-Spline Tokenizer

Following prior works in action tokenization [10, 16], we first normalize the input actions such that the 1st and 99th quantile value of each action dimension in the dataset maps to the range of  $[-1, 1]$ . Using quantiles makes the normalization robust against outlier data points.

Figure 2 presents an overview of the tokenization process. We begin by considering the tokenization of a 1-DoF trajectory. Given a normalized action sequence  $a_{1:T} = [a_1, a_2, \dots, a_T]$  of length  $T$ , our goal is to determine a set of  $N$  control points  $c$ , with  $N \leq T$ , that approximate the given action sequence at spline evaluations  $y(u)_{1:T}$ . The linear transformation  $u = t/T$  maps from action timestep to the parametric coordinate of the B-Spline. The spline evaluations  $y(u)_{1:T}$  are approximated by

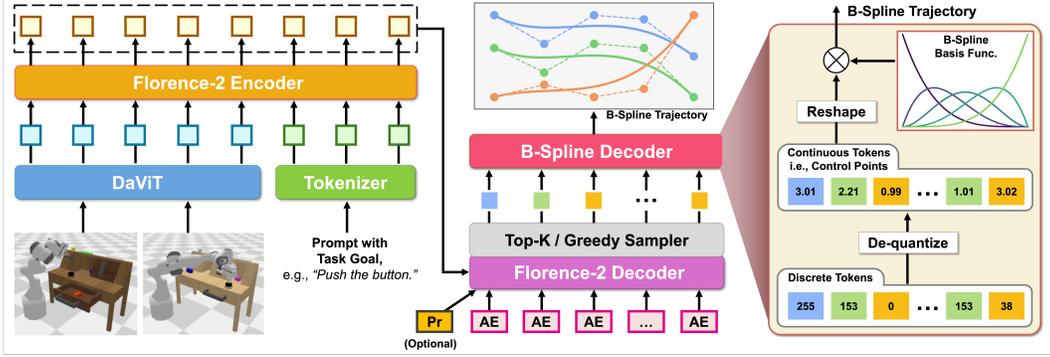


Figure 3: **BEAST-F** is a new VLA model that combines BEAST encoding with Florence-2 [22], a lightweight VLM with 0.77B parameters. BEAST produces uniform-length tokens, which allows BEAST-F to perform *parallel decoding* via learnable action embeddings (**AE**), instead of autoregressive next-token prediction. These discrete tokens are fed into the *B-Spline Decoder*, which first maps them to real-valued control points and then transforms those control points into continuous action sequences. The **Pr** token denotes an optional proprioceptive state.

minimizing the least-squares error

$$\mathbf{c} = \arg \min_{\mathbf{c}} \|\mathbf{y}_{1:T} - \mathbf{a}_{1:T}\|_2^2 = \arg \min_{\mathbf{c}} \|\Phi^P(u)\mathbf{c} - \mathbf{a}_{1:T}\|_2^2, \quad (4)$$

where  $\Phi^P(u) = [\Phi_1^P(u), \Phi_2^P(u), \dots, \Phi_N^P(u)]^\top$  represents precomputed B-spline basis functions defined over interval  $u \in [0, 1]$ . Ridge regression estimates the control points in closed form,  $\mathbf{c} = [c_0, c_1, \dots, c_{N-1}] = (\Phi^\top \Phi + \lambda \mathbf{I})^{-1} \Phi^\top \mathbf{a}_{1:T}$ , with  $\lambda$  acting as a regularization parameter. This efficient computation typically introduces only a small overhead, typically 3 to 5 milliseconds per batch. For a high-dimensional action sequence, i.e.  $D > 1$ , each DoF is encoded independently into  $\mathbf{c}_d$ , resulting in a matrix  $\mathbf{C}$  of shape  $D \times N$ , that stacks each DoF’s control points,  $\mathbf{C} = [c_1, c_2, \dots, c_D]^\top$ .

To form the final token sequence, this matrix is flattened by interleaving different action dimensions corresponding to the same basis functions, as illustrated in Figure 2. This flattening strategy preserves the temporal order inherent in the trajectory segments associated with each basis function.

**Remark 1: Connection to Action Chunking.** Action chunking, defined as a discrete sequence of actions  $a_0, a_1, \dots, a_T$ , is mathematically equivalent to a piecewise constant function generated by 0-th degree B-splines. As demonstrated in Figure 1 left most, each action step  $a_t$  can be identified as a control point  $c_n$  of 0-th degree B-Spline basis with  $t = n, T = N$ .

### 3.2 Enforcing Smooth Transition with Clamped B-Spline

Executing long-horizon tasks typically requires producing multiple small action sequences that connect seamlessly (replanning). While predicting action sequences effectively improves consistency within individual action chunks, a significant challenge lies in managing discontinuities at transitions between consecutive chunks, which often result in jerky motion during online execution. Common approaches to address this issue apply temporal ensembles of actions [4, 23], calculating moving averages over multiple predictions. However, these temporal ensembles require high-frequency replanning (typically every timestep) to generate sufficient chunks for effective ensemble averaging, which significantly constrains execution speed in online applications.

In contrast, BEAST employs clamped B-Spline to ensure smooth transitions between consecutive action chunks. As introduced in Section 2, clamped B-Spline is a specialized variant of B-Spline that guarantees to start from the first control point and end at the last control point, which is utilized to generate seamlessly connected action sequences, as illustrated in Figure 1. To ensure smooth transitions, we directly set the first control point  $c_0$  to the last action of the previous sequence. We then compute the residual trajectory  $\hat{\mathbf{a}}$  by subtracting the contribution of the first basis function:  $\hat{\mathbf{a}} = \mathbf{a} - c_0 \Phi_0^P$ . The remaining control points  $\hat{\mathbf{c}} = [c_1, c_2, \dots, c_{N-1}]$  are determined by solving the

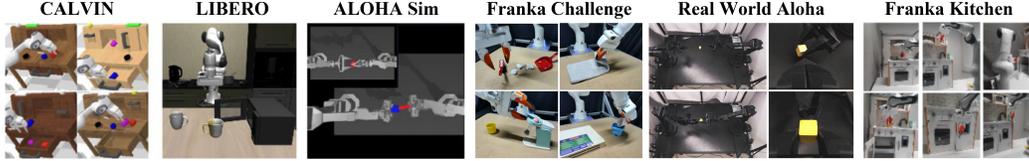


Figure 4: Simulation [4, 24, 25] and real world (Franka Challenge, Aloha, Franka Kitchen) tasks.

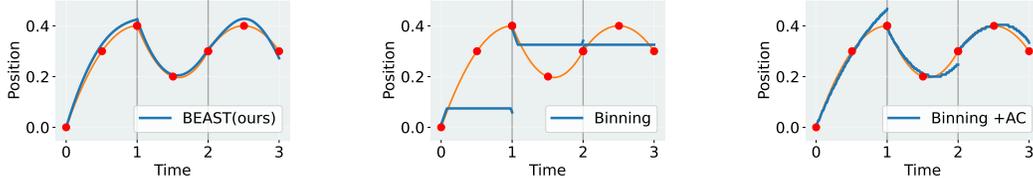


Figure 5: Comparison among BEAST, single-step binning tokenization and binning tokenization with action chunking (AC). The comparison is conducted through the same auto-regressive model with different tokenizers to fit the same **ground truth** cube splines given the same **context points**. BEAST is smooth within each sequence and continuous at the transitions between sequences.

linear regression problem similar to equation 4:  $\arg \min_c \|\hat{\Phi}^P(u)\hat{c} - \hat{a}\|_2$ . Through this approach, BEAST consistently generates action sequences with mathematically guaranteed smooth transitions between chunks. This will be further discussed in our toy task experiment in Section 4.1.

### 3.3 Combining BEAST tokens with different architectures

**Discrete Tokens.** We first evaluate BEAST in a simplified setting with a decode-only transformer (see Figure 9) with CLIP [26] for language encoding and Film-conditioned ResNet-18 [27, 28] as image encoder. Film-ResNets are used in many prior works given their high efficiency and strong performance [29, 30, 11]. The proprioceptive state of robot is projected to the embedding dimension with a two-layer MLP. We employ parallel decoding with bi-directional attention to accelerate the inference. To further demonstrate the scalability of BEAST with large-pretrained models, we combine BEAST with Florence-2, a small, pretrained VLM with Encoder-Decoder architecture (0.77B parameters). Following the previous works on autoregressive VLAs [10, 16], we overwrite the least used 256 tokens in the VLM vocabulary as our action tokens. We also employ a parallel decoding technique for the Florence variant, which significantly improves the throughput and reduces the latency for action generation. We provide an in-depth overview in Figure 3.

**Continuous Tokens.** We also explore the performance of combining BEAST with ACT [4]. ACT uses a conditional VAE (CVAE) with a Transformer Encoder-Decoder to predict a sequence of actions. We predict  $N$  BEAST continuous tokens, where each token has the dimension of  $D$ , instead of action sequences, this design choice keeps the temporal order inherent in the trajectory segments. Using the BEAST tokenization, we reduce the length of predicted token sequence by 6.67 times (from 100 to 15) without sacrificing the task performance. In addition, our method enables smooth trajectories without requiring temporal aggregation.

## 4 Experiments

We conducted extensive evaluations in both simulated and real-world settings, targeting answering five key research questions (RQs): 1) What advantages does BEAST offer over commonly used binning-based tokenizers? 2) How does BEAST contribute to the performance on imitation learning benchmarks? 3) How does BEAST affect the training and inference efficiency? 4) Does BEAST generalize to real-world scenarios? BEAST is integrated into two different architectures: We combine BEAST and Florence-2 [22] and term this VLA variant as BEAST-F. In contrast to many baselines, we test BEAST-F **without second-stage pretraining on large-scale robot datasets**. We additionally

Train→Test	Method	PrT	Action Type	VLM	No. Instructions in a Row (1000 chains)					Avg. Len.
					1	2	3	4	5	
ABC→D	Diff-P-CNN [5]	×	Diffusion	×	63.5%	35.3%	19.4%	10.7%	6.4%	1.35
	MDT [31]	×	Diffusion	×	63.1%	42.9%	24.7%	15.1%	9.1%	1.55
	OpenVLA [10]	✓	Discrete	✓	91.3%	77.8%	62.0%	52.1%	43.5%	3.27
	3DDA [32]	×	Diffusion	×	93.8%	80.3%	66.2%	53.3%	41.2%	3.35
	MoDE [30]	✓	Diffusion	×	96.2%	88.9%	81.1%	71.8%	63.5%	4.01
	VPP [33]	✓	Diffusion	×	95.7%	91.2%	86.3%	81.0%	<b>75.0%</b>	4.29
	<b>BEAST-F (ours)</b>	×	Discrete	✓	<b>99.8%</b>	<b>96.5%</b>	<b>89.3%</b>	<b>82.7%</b>	74.4%	<b>4.42</b>
ABCD→D	Diff-P-CNN [5]	×	Diffusion	×	86.3%	72.7%	60.1%	51.2%	41.7%	3.16
	MoDE [30]	✓	Diffusion	×	97.1%	92.5%	87.9%	83.5%	77.9%	4.39
	MDT [31]	×	Diffusion	×	<b>98.6%</b>	95.8%	91.6%	86.2%	80.1%	4.52
	<b>BEAST-F (ours)</b>	×	Discrete	✓	98.1%	<b>96.2%</b>	<b>93.0%</b>	<b>89.3%</b>	<b>84.8%</b>	<b>4.61</b>

Table 1: **CALVIN Benchmark results for ABC and ABCD.** The table reports average success rates for individual tasks within instruction chains and the average rollout length (Avg. Len.) to complete 5 consecutive instructions, based on 1000 chains. Zero standard deviation indicates methods without reported standard deviations. BEAST-F achieves SoTA performance in both tasks.

	Spatial		Object		Goal		Long		Average	
	SR (↑)	Rank (↓)	SR (↑)	Rank (↓)						
Diff-P-CNN	78.3 ± 1.1%	6	92.5 ± 0.7%	4	68.3 ± 1.2%	6	50.5 ± 1.3%	6	72.4 ± 0.7%	6
Octo	78.9 ± 1.0%	5	85.7 ± 0.9%	6	84.6 ± 0.9%	4	51.1 ± 1.3%	5	75.1 ± 0.6%	5
OpenVLA	84.7 ± 0.9%	4	88.4 ± 0.8%	5	79.2 ± 1.0%	5	53.7 ± 1.3%	4	76.5 ± 0.6%	4
$\pi_0$	<b>96.8%</b>	1	<b>98.8%</b>	1	<b>95.8%</b>	1	85.2%	2	94.2%	1
$\pi_0$ -FAST	96.4%	2	96.8%	3	88.6%	3	60.2%	3	85.5%	3
<b>BEAST-F</b>	92.9%	3	97.5%	2	93.1%	2	<b>86.4%</b>	1	92.5%	2

Table 2: **Experimental Results for the LIBERO Benchmarks.** SR: Success Rate. Best results in each column are shown in bold. BEAST-F achieves comparable performance state-of-the-art VLA, despite with a much smaller model and without robot data pretraining.

apply BEAST to a small decoder-only transformer model (BEAST-D) and on top of a vanilla ACT [4] (BEAST-ACT).

#### 4.1 Comparing Against Binning-Based Tokenization

To answer **RQ1**, we begin with a 1D toy task to investigate the advantages of BEAST over binning-based tokenization. We follow the autoregressive prediction pipeline used in previous works [10, 16]. Note that BEAST can be used for both autoregressive prediction and parallel decoding. A small decoder-only transformer is trained to predict cubic splines from 3 control points. We compare against: 1) Single-step binning (denoted as **Binning**) [10], which discretizes each action into one of 256 bins, and 2) Chunk-level binning (denoted as **Binning+AC**), which discretizes entire action sequences of fixed length. We generate 2000 trajectories, 1s each at 100Hz resolution. Each model is trained for 8k steps and evaluated on 200 test sequences. BEAST achieves the lowest MSE ( $0.0004 \pm 0.0005$ ), outperforming chunked binning ( $0.0009 \pm 0.0013$ ) and single-step binning ( $0.0215 \pm 0.0216$ ), with the latter performing two orders of magnitude worse. To simulate real-world action chunking [5], we repeat the rollout prediction three times. As visualized in Figure 5, single-step binning fails to capture temporal structure and produces erratic outputs. Chunked binning captures some temporal coherence but results in jerky transitions due to discretization and a lack of continuity across chunks. In contrast, BEAST generates smooth trajectories with minimal error and requires only 5 tokens per 100-step sequence, resulting in an approximately 20x reduction in inference steps.

#### 4.2 Strong Performance on Established Simulation Benchmarks

To answer **RQ2**, we evaluate BEAST on established simulation benchmarks and compare with other SoTA imitation learning methods and VLAs.

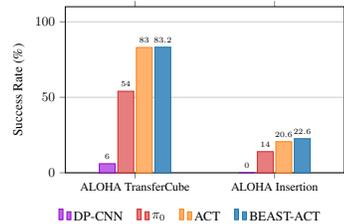


Figure 6: **ALOHA Benchmark results.** The success rate is reported over 500 episodes of evaluation.

**Simulation Benchmarks.** CALVIN [24] features 34 tabletop manipulation tasks with a Franka Panda robot using delta end-effector control across four scene configurations (splits A-D). The dataset contains 24,000 language-annotated demonstrations. We evaluate two settings: CALVIN ABC (zero-shot generalization) and CALVIN ABCD (scaling with more data). Performance is measured by success rates on sequential tasks and mean sequence length completion. All evaluations require policies to follow free-form language instructions and complete 5 tasks in sequence across 1,000 different instruction chains. LIBERO [25] tests a delta-EEF controlled Panda Robot across various scenes with **130** diverse tasks. We report results on four specialized benchmark settings with 10 tasks each (Long, Spatial, Object, and Goal). Success is measured as the percentage of successful task completions across 50 trials per task. ALOHA [4] tests an absolute joint position controlled ALOHA Robot in two challenging bi-manual manipulation tasks that require high precision.

**Baselines.** We compare our Vision-Language-Action Model (VLA) against SOTA VLA policies and specialized approaches, using results reported in prior publications for fair comparison. Our primary baselines are OpenVLA [10] (7.7B parameters),  $\pi_0$  [34] (3.3B parameters),  $\pi_0$ -FAST[16] (3.3B parameters), and a standard Diffusion Policy using a CNN [5]. For the bi-manual manipulation tasks, we compare the BEAST-ACT variant with small action chunking models to a vanilla ACT [4],  $\pi_0$ , and a standard Diffusion Policy using a CNN.

**Results.** Table 1 summarizes the performance of all policies on the CALVIN benchmark, where BEAST-F outperforms a diverse set of baselines across two settings, establishing a new state of the art. Unlike the most competitive baselines, BEAST-F achieves these results without relying on additional pretraining. On the various LIBERO benchmarks, our tokenizer achieves strong performance, being surpassed only by  $\pi_0$ -VLA. However,  $\pi_0$  relies on large-scale pretraining to reach its performance, whereas BEAST-F remains competitive without it. In the most challenging long-horizon task setting, LIBERO-LONG, BEAST-F outperforms all baselines. See Table 2 for detailed results. For the bi-manual tasks (Figure 6), BEAST-ACT and ACT demonstrate significantly better performance than  $\pi_0$ . BEAST-ACT achieved a higher success rate than vanilla ACT in both tasks.

### 4.3 Advantages in Training and Inference Speed

Next, we verify the inference and training efficiency of BEAST to answer **RQ3**. Specifically, we consider the VLA variant **BEAST-F** and compare it against several recent VLAs [10, 34, 16], as well as a standard CNN-based Diffusion Policy[5]. We measure the inference efficiency on an RTX 4090 GPU. As shown in Table 3, BEAST-F demonstrates clear computational advantages. It achieves a throughput of 617.3 Hz (e.g., generates approximately 617 actions per second), which is  $2.14\times$  faster than  $\pi_0$ ,  $4.72\times$  faster than Diffusion Policy, and  $101.4\times$  faster than OpenVLA. In addition, BEAST-F achieves the lowest latency at just 19 milliseconds, where latency refers to the time taken to generate one action chunk. These gains are due to the parallel decoding, which enables generating the action sequence in a single forward pass.

Method	Throughputs (Hz) $\uparrow$	Latency (s) $\downarrow$
DP (0.26B)	130.67	0.341
OpenVLA (7B)	6.09	0.164
$\pi_0$ (3.3B)	288.11	0.104
<b>BEAST-F (0.77B)</b>	<b>617.3</b>	<b>0.019</b>

Table 3: **Mean inference efficiency** (1000 steps in Bf16). All policies except OpenVLA use chunking length 50 (48 for DP).

We further evaluate the training efficiency by comparing BEAST-F against  $\pi_0$  and  $\pi_0$ -FAST. To exclude the bias introduced by the pretraining datasets, we trained all models **without robot dataset pretraining**. We report the success rate on LIBERO-LONG benchmark every 10k training steps in Figure 7. BEAST-F reaches a approximate 80% success rate at just 20k steps, whereas  $\pi_0$  reaches only around 20% at the same point. Notably,  $\pi_0$ -FAST shows no success till 30k steps.  $\pi_0$ -FAST’s poor performance indicates a heavy reliance on robot dataset pretraining, which further underscores the training efficiency of our method.

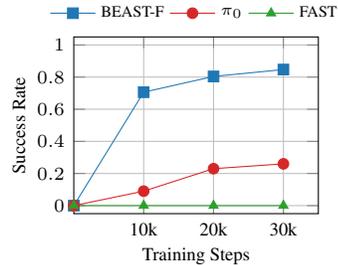


Figure 7: LIBERO-LONG.

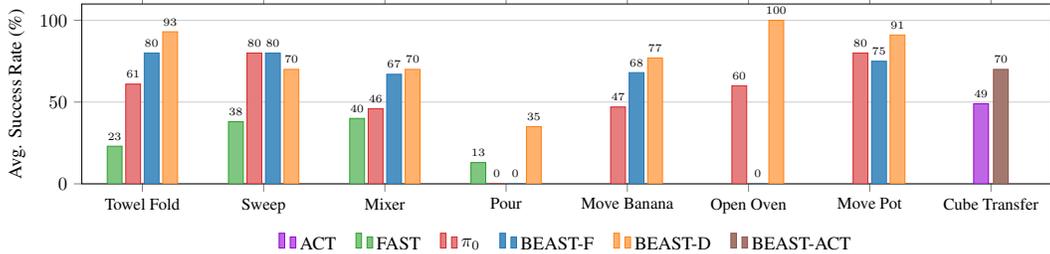


Figure 8: **Experimental Results on Real-World Robot Tasks.** This figure shows the average task success rate across eight real-world tasks. Each task and method was evaluated over 10 runs (30 runs for Cube Transfer). Success rates are measured at the sub-task level. Detailed descriptions of all sub-tasks are provided in Appendix D. BEAST variants achieve strong performance in real world.

#### 4.4 Real-World Evaluation with 3 Different Robot Setups

To answer **RQ4**, we assess the effectiveness of BEAST across diverse real-world scenarios with varying data collection frequencies. We evaluate BEAST on 8 challenging manipulation tasks across 3 different experimental setups: 1) **Franka Challenge**: Four tabletop manipulation tasks (Towel Fold, Sweep, Mixer, Pour) using a joint position-controlled Franka robot with data collected at 20Hz, 2) **Real Kitchen**: Three manipulation tasks on a toy kitchen setup (Move Banana, Open Oven, Move Pot) with data collected at 35Hz, 3) **Bi-manual ALOHA**: A cube transfer task using a bi-manual ALOHA robot with data recorded at 60Hz. For each task in the **Franka Challenge** and **Franka Kitchen** setups, we conduct 10 evaluation runs per method, while for the **ALOHA cube transfer** task, we performed 30 runs. The average success rate for each task is reported in Figure 8. For tasks comprising multiple stages, we track intermediate milestones to better evaluate the completion of each sub-task. Appendix D provides a detailed description of all setups and tasks. We compare BEAST against  $\pi_0$ [34],  $\pi_0$ -FAST[16], and ACT [4]. We finetune  $\pi_0$  and  $\pi_0$ -FAST from the official pretrained checkpoints for an additional 60k and 40k steps, respectively. For each method, we train one multitask model for all four tasks, Real Franka tasks, and another for the Real Kitchen tasks. The results demonstrate that BEAST-F achieves 52.86% success rate and BEAST-D achieves 76.57%. In contrast  $\pi_0$  achieves 53.43% and FAST only 28.5%. Interestingly, the smaller model (BEAST-D) outperforms all the VLAs, including the Florence variant with BEAST. We attribute this effect to the relatively small real-world dataset of only 50 demonstrations for each task. For the Aloha Cube Transfer task, we compare BEAST-ACT against the base ACT that directly predicts action sequences in the joint space. BEAST-ACT achieves 70% success, which is 21% higher than the base ACT.

## 5 Conclusion

We present BEAST, a B-spline-based tokenizer for continuous robot actions that compresses arbitrary trajectories into fixed-length token sequences while preserving smooth transitions between segments. BEAST supports discrete and continuous outputs and integrates seamlessly with various model architectures. By exploiting parallel decoding, it delivers fast inference and high compression rates without sacrificing performance. In extensive experiments—both in simulation and on real robots—BEAST consistently achieves strong results, demonstrating the effectiveness of our tokenization strategy. **Limitations:** Although BEAST delivers strong performance, it is sensitive to the choice of the number of B-spline basis functions, which can markedly affect task outcomes (Section ??). The optimal count depends on the smoothness and sampling frequency of the trajectory; our experiments indicate that using 5–10 bases works well for one-second robot trajectories. **Future Work:** We plan to extend BEAST to large-scale robot pretraining and to integrate continuous token representations with diffusion- and flow-matching objectives, aiming to further boost downstream task performance.

## References

- [1] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, J. Peters, et al. An algorithmic perspective on imitation learning. *Foundations and Trends® in Robotics*, 7(1-2):1–179, 2018.
- [2] Y. Ding, C. Florensa, P. Abbeel, and M. Phielipp. Goal-conditioned imitation learning. *Advances in neural information processing systems*, 32, 2019.
- [3] C. Lynch and P. Sermanet. Language conditioned imitation learning over unstructured data. *Robotics: Science and Systems XVII*, 2020. URL <https://api.semanticscholar.org/CorpusID:235657751>.
- [4] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.
- [5] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song. Diffusion policy: Visuomotor policy learning via action diffusion. In *Proceedings of Robotics: Science and Systems (RSS)*, 2023.
- [6] L. Lai, A. Z. Huang, and S. J. Gershman. Action chunking as policy compression. 2022.
- [7] A. Van Den Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. In *International conference on machine learning*, pages 1747–1756. PMLR, 2016.
- [8] T. Mikolov, M. Karafiát, L. Burget, J. Cernocký, and S. Khudanpur. Recurrent neural network based language model. In *Interspeech*, volume 2, pages 1045–1048. Makuhari, 2010.
- [9] A. Radford, K. Narasimhan, T. Salimans, I. Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- [10] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi, et al. Openvla: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.
- [11] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. Ryoo, G. Salazar, P. Sanketi, K. Sayed, J. Singh, S. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich. Rt-1: Robotics transformer for real-world control at scale. In *arXiv preprint arXiv:2212.06817*, 2022.
- [12] 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.
- [13] S. Lee, Y. Wang, H. Etukuru, H. J. Kim, N. M. M. Shafiullah, and L. Pinto. Behavior generation with latent actions. In *International Conference on Machine Learning*, pages 26991–27008. PMLR, 2024.
- [14] A. Szot, B. Mazouze, H. Agrawal, R. D. Hjelm, Z. Kira, and A. Toshev. Grounding multimodal large language models in actions. *Advances in Neural Information Processing Systems*, 37: 20198–20224, 2024.
- [15] A. Szot, B. Mazouze, O. Attia, A. Timofeev, H. Agrawal, D. Hjelm, Z. Gan, Z. Kira, and A. Toshev. From multimodal llms to generalist embodied agents: Methods and lessons. *arXiv preprint arXiv:2412.08442*, 2024.
- [16] K. Pertsch, K. Stachowicz, B. Ichter, D. Driess, S. Nair, Q. Vuong, O. Mees, C. Finn, and S. Levine. Fast: Efficient action tokenization for vision-language-action models. *arXiv preprint arXiv:2501.09747*, 2025.

- [17] 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.
- [18] M. J. Kim, C. Finn, and P. Liang. Fine-tuning vision-language-action models: Optimizing speed and success. *arXiv preprint arXiv:2502.19645*, 2025.
- [19] W. J. Gordon and R. F. Riesenfeld. B-spline curves and surfaces. *Computer Aided Geometric Design*, pages 95–126, 1974. URL <https://api.semanticscholar.org/CorpusID:118698454>.
- [20] H. Prautzsch, W. Boehm, and M. Paluszny. *Bézier and B-spline techniques*. Springer Science & Business Media, 2002.
- [21] C. d. Boor. Subroutine package for calculating with b-splines. In *Los Alamos Sci. Lab*. Los Alamos, NM, USA, 1971.
- [22] B. Xiao, H. Wu, W. Xu, X. Dai, H. Hu, Y. Lu, M. Zeng, C. Liu, and L. Yuan. Florence-2: Advancing a unified representation for a variety of vision tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4818–4829, 2024.
- [23] Q. Li, Y. Liang, Z. Wang, L. Luo, X. Chen, M. Liao, F. Wei, Y. Deng, S. Xu, Y. Zhang, et al. Cogact: A foundational vision-language-action model for synergizing cognition and action in robotic manipulation. *arXiv preprint arXiv:2411.19650*, 2024.
- [24] O. Mees, L. Hermann, E. Rosete-Beas, and W. Burgard. Calvin: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks. *IEEE Robotics and Automation Letters*, 7(3):7327–7334, 2022.
- [25] 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.
- [26] 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.
- [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] 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.
- [29] H. Bharadhwaj, J. Vakil, M. Sharma, A. Gupta, S. Tulsiani, and V. Kumar. Roboagent: Generalization and efficiency in robot manipulation via semantic augmentations and action chunking, 2023.
- [30] M. Reuss, J. Pari, P. Agrawal, and R. Lioutikov. Efficient diffusion transformer policies with mixture of expert denoisers for multitask learning. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=nDmw1oE13N>.
- [31] M. Reuss, Ö. E. Yağmurlu, F. Wenzel, and R. Lioutikov. Multimodal diffusion transformer: Learning versatile behavior from multimodal goals. In *Robotics: Science and Systems*, 2024.
- [32] T.-W. Ke, N. Gkanatsios, and K. Fragkiadaki. 3d diffuser actor: Policy diffusion with 3d scene representations. In *8th Annual Conference on Robot Learning*, 2024. URL <https://openreview.net/forum?id=gqCQx0bVz2>.

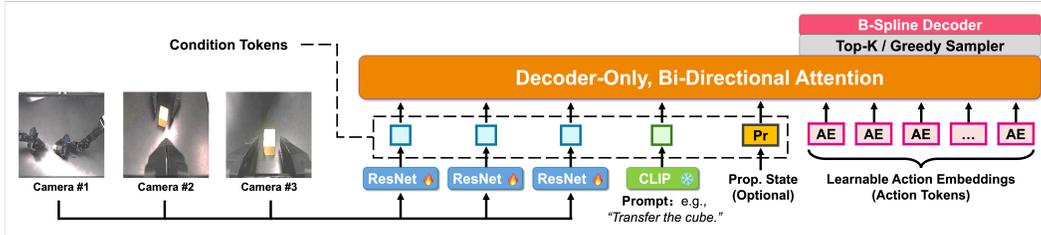


Figure 9: **Overview of BEAST-D:** BEAST-D is a small transformer model that integrates BEAST. It replaces the causal attention in the decoder-only transformer with bidirectional attention to enable fast parallel decoding. BEAST-D integrates ResNet as image encoder and CLIP as language encoder.

- [33] Y. Hu, Y. Guo, P. Wang, X. Chen, Y.-J. Wang, J. Zhang, K. Sreenath, C. Lu, and J. Chen. Video prediction policy: A generalist robot policy with predictive visual representations, 2024. URL <https://arxiv.org/abs/2412.14803>.
- [34] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter, et al. *pi\_0*: A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.

## A Architectures

## B Baselines Implementation

$\pi_0$ :  $\pi_0$  is a generalist robot policy that combines a pre-trained VLM backbone with a lightweight *action expert* module trained from scratch to generate continuous actions using *flow matching*. A key innovation of  $\pi_0$  is its cross-embodiment training strategy, which integrates over 900M timesteps of data from 7 distinct robot embodiments and 68 manipulation tasks, enabling generalization across heterogeneous hardware platforms. The model is trained using a two-phase pipeline: a large-scale pre-training stage leveraging Internet-scale semantic priors, followed by post-training on curated task-specific data to enhance performance on complex, dexterous tasks.

**FAST**: FAST introduces a novel compression-based tokenization method, named *Frequency-space Action Sequence Tokenization*, for training autoregressive VLA models on high-frequency, dexterous robot control tasks. Unlike prior VLAs that struggle with discretizing continuous actions at high frequencies, FAST leverages the *Discrete Cosine Transform (DCT)* and *Byte-Pair Encoding (BPE)* to produce compact, information-rich action tokens, marking a significant advance in training efficiency.

## C Hyperparameters

Hyperparameter	LIBERO				CALVIN	
	SPATIAL	OBJECT	GOAL	LONG	ABCD→D	ABC→D
Action Sequence Length	20	20	20	20	20	20
Number of Basis	10	10	10	10	10	10
Vocabulary Size	256	256	256	256	256	256
Optimizer	AdamW	AdamW	AdamW	AdamW	AdamW	AdamW
Betas	[0.9, 0.95]	[0.9, 0.95]	[0.9, 0.95]	[0.9, 0.95]	[0.9, 0.95]	[0.9, 0.95]
Learning Rate	2e-5	2e-5	2e-5	2e-5	2e-5	2e-5
Batch Size	128	128	128	128	32	32
Train Steps (k)	35	35	50	70	30	30

Table 4: Summary of BEAST-F hyperparameters for all simulation experiments.

Hyperparameter	REAL KITCHEN	REAL FRANKA
Action Sequence Length	80	20
Number of Basis	15	5
Vocabulary Size	256	256
Optimizer	AdamW	AdamW
Betas	[0.9, 0.95]	[0.9, 0.95]
Learning Rate	2e-5	2e-5
Batch Size	96	96
Train Steps (k)	60	60

Table 5: BEAST-F hyperparameters for real robot experiments.

Hyperparameter	REAL KITCHEN	REAL FRANKA
Action Sequence Length	80	20
Number of Basis	10	5
Vocabulary Size	256	256
Transformer Layers	6	6
Attention Heads	8	8
Embedding Dim	256	256
Image Encoder	FiLM-ResNet18	FiLM-ResNet18
Goal Lang Encoder	CLIP ViT-B/32	CLIP ViT-B/32
Attn Dropout	0.1	0.1
Residual Dropout	0.1	0.1
MLP Dropout	0.1	0.1
Optimizer	AdamW	AdamW
Betas	[0.9, 0.999]	[0.9, 0.999]
Learning Rate	3e-4	3e-4
Weight Decay (Trans/Other)	0.05 / 0.05	0.05 / 0.05
Batch Size	384	256
Train Steps (k)	60	60
EMA	False	False

Table 6: BEAST-D hyperparameters for real robot experiments.

## D Real Robots Setup & Tasks

### D.1 Robot System Details

**Real Kitchen.** This setup consists of a single Franka Emika robot operating within a simulated kitchen environment. It is equipped with two OAK-D Lite cameras providing top-down and side perspectives, each delivering visual input at a resolution of 250×250 pixels. The robot has an 8-dimensional configuration and action space, which includes seven joint and one gripper states.

**Real Franka.** This configuration features a single Franka Emika robot situated in a general-purpose tabletop environment designed for more challenging manipulation tasks. Visual observations are obtained from two Orbbec Femto Bolt cameras, positioned to capture left and right perspectives. The input images are resized to a resolution of 180×320 pixels. The robot configuration and action space remain the same as the Franka Kitchen setup.

**ALOHA.** Based on the ALOHA setup [4], the system incorporates two 6-DoF Trossen ViperX robotic arms. The environment includes two wrist-mounted and an additional top-mounted Logitech C920 camera. The combined system operates in a 14-dimensional configuration and action space, accounting for both arms’ joint and gripper states.

### D.2 Tasks Description and Evaluation Metrics

In the Real Kitchen setup, the robot performs pick-and-place tasks, whereas in the Real Franka setup, the robot is required to execute more diverse manipulation behaviors, such as sweeping or pouring. For each task performed by the Franka Emika robot, a scoring rubric is defined to quantitatively evaluate task progression. The specific evaluation criteria for each task are detailed below.

- **Open the door:** The task begins under one of two initial conditions: with or without an object placed on the stove. The objective is for the robot to open the oven door. Task completion is evaluated as a success or a failure. Although the task involves three key motion phases, as shown in Figure 10 (Open the door), all the policies under evaluation are capable of completing the task in its entirety once the robot successfully grasps the handle. A trial is considered successful if the robot fully opens the oven door by first grasping the handle and then using its fingers to push the opposite side of the door, ensuring that it is completely open. We conducted four evaluation trials with no object on the stove and one with a randomly placed object.

- **Banana into the pot:** In this task, the robot aims to grasp a banana and place it on or into a pot. The initial conditions are categorized based on the relative positions of the pot and the banana, as well as the position of the banana relative to the corresponding stove. The pot is assumed to be correctly positioned on the stove. The banana, however, may be placed directly on the stove in one of three orientations or slightly offset to the left or right. A total of 10 trials are conducted across these different initial configurations. Task performance is scored on a scale of 0 to 3, with one point awarded for each of the following criteria: (1) successfully positioning the banana between the robot’s two fingers, (2) lifting the banana off the surface, and (3) placing the banana onto or into the pot. If the robot attempts to grasp more than three times, exhibits jerky hand movements, or significantly displaces the pot from its original position, the subtask is awarded 0.5 points to reflect partial completion.
- **Pot into the sink:** This task is similar to the one described previously. The initial conditions in this task differ based on the relative position between the pot and the two stoves. Task performance is evaluated using three criteria: (1) successfully positioning the pot between the robot’s two fingers - note that directly grasping the pot with its handle is considered an unstable grasp and is awarded 0.5 points, (2) lifting the pot off the surface, and (3) placing the pot in the sink. In this task, penalties for jerky hand movements are still applied.
- **Towel folding:** The objective of this task is to neatly fold a towel that is randomly oriented at the beginning of each trial. One point is awarded for each of the following: lifting a corner of the towel, completing the fold, and achieving accurate alignment of the folded towel.
- **Sweep:** In the Sweep task, the positions of the broom, dustpan, and trash vary across trials. Four pieces of trash are placed on the table for the robot to clean. A maximum of four points can be awarded, based on the following criteria: successfully grasping the broom, performing a single sweeping motion, demonstrating the ability to execute multiple sweeping motions, and sweeping all trash into the dustpan.
- **Mixer:** In this task, a cup and a mixer are placed on the table. The robot’s objective is to sequentially (1) open the mixer, (2) grasp the cup, (3) place the cup on the mixer’s platform, and (4) close the mixer. Task performance is evaluated based on the successful completion of these four subtasks, with one point awarded for each. Notably, unlike previous tasks, the language instructions provided to the robot consist of three separate sentences corresponding to the actions of opening/closing the mixer and placing the cup onto the platform.
- **Pour:** In the Pour task, the source cup is initially placed on a platform and contains plastic pellets that simulate liquid. The objective is to pour the pellets into a designated target cup. Task performance is evaluated out of a maximum of 4 points, awarded based on the following criteria: (1) successfully grasping the source cup, (2) pouring the pellets into the target cup, (3) ensuring that all pellets are poured into the target cup, and (4) placing the source cup back on the platform.
- **ALOHA cube transfer:** In the cube transfer task, the ALOHA robot is designed to pick up a randomly placed cube using its right arm and then transfer the cube to its left arm. The performance of the task is evaluated by assigning scores to three specific steps: (1) successfully picking up the cube, (2) successfully initiating the transfer with the left arm making contact with the cube, and (3) the left arm successfully taking possession of the cube while the right arm releases it.

## E Compute Resources

We train and evaluate all the models based on our private clusters. Each node contains 4 NVIDIA A100, for BEAST-F we use 4 GPUs for training. For BEAST-D and BEAST-ACT, we use one GPU for training. We report the average training cost in Table 7.

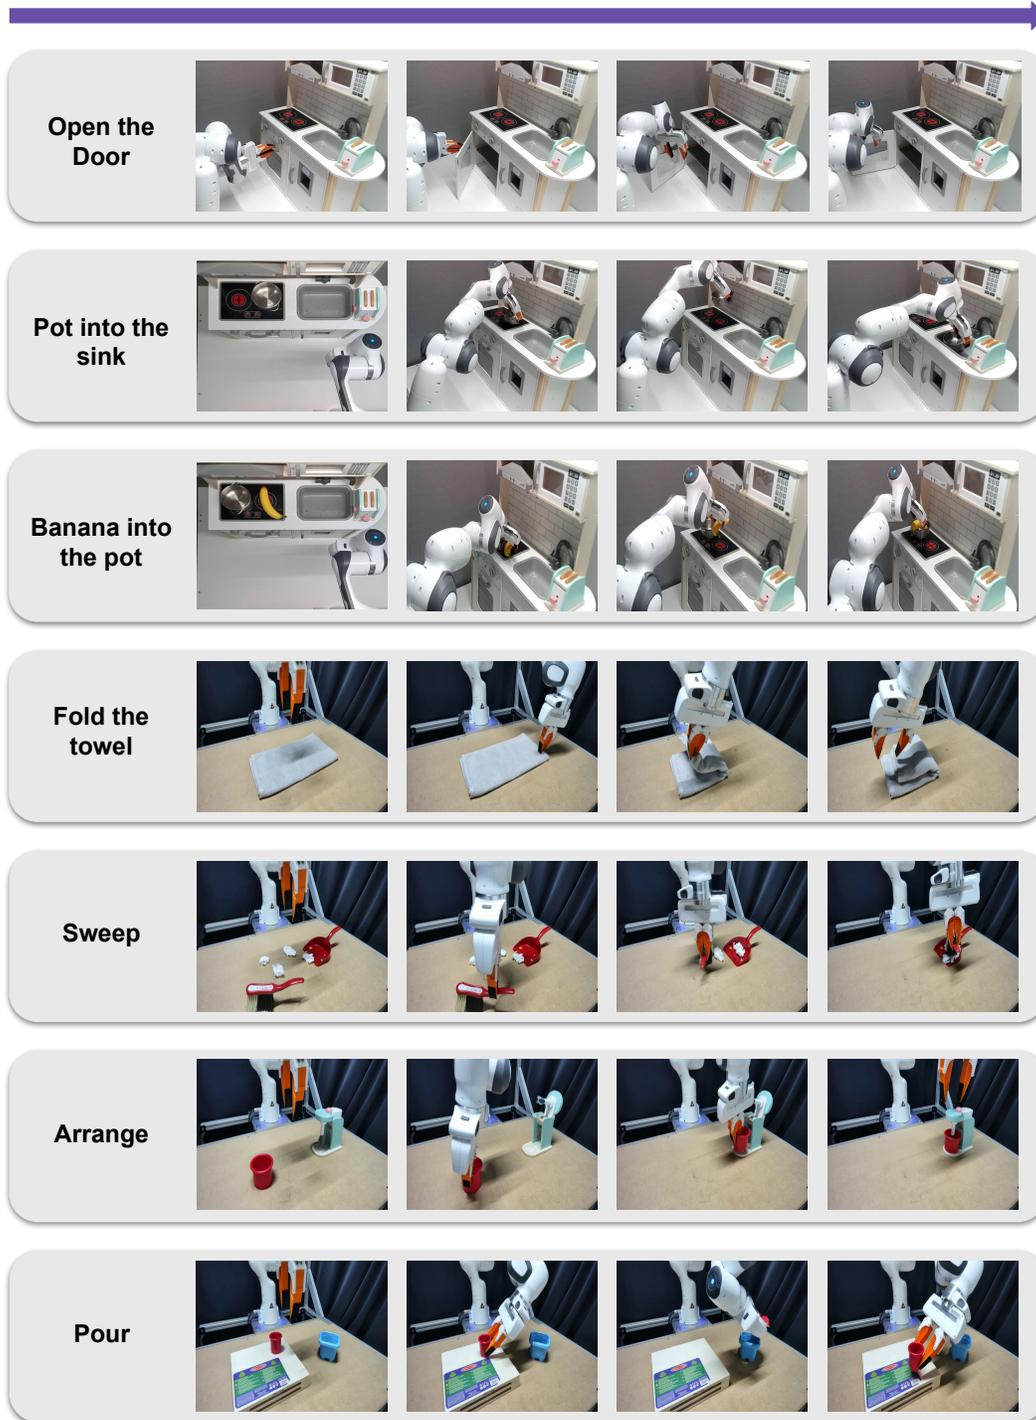


Figure 10: Key frames for real world different tasks

	BEAST-F	BEAST-D	BEAST-ACT
vRAM	64GB	8GB	15GB
steps/hour	6000	10000	11000

Table 7: Training time for each variant.