ELMO: Emergent Representations In Imitation Learning Via Dynamic Hierarchical Action Chunking

Jasper Gerigk[*1,2] Yang Xu[*1,2] Nicholas Rhinehart[†1] Igor Gilitschenski[†1,2]

[1]University of Toronto [2]Vector Institute

[*]Equal contribution [†]Shared supervision

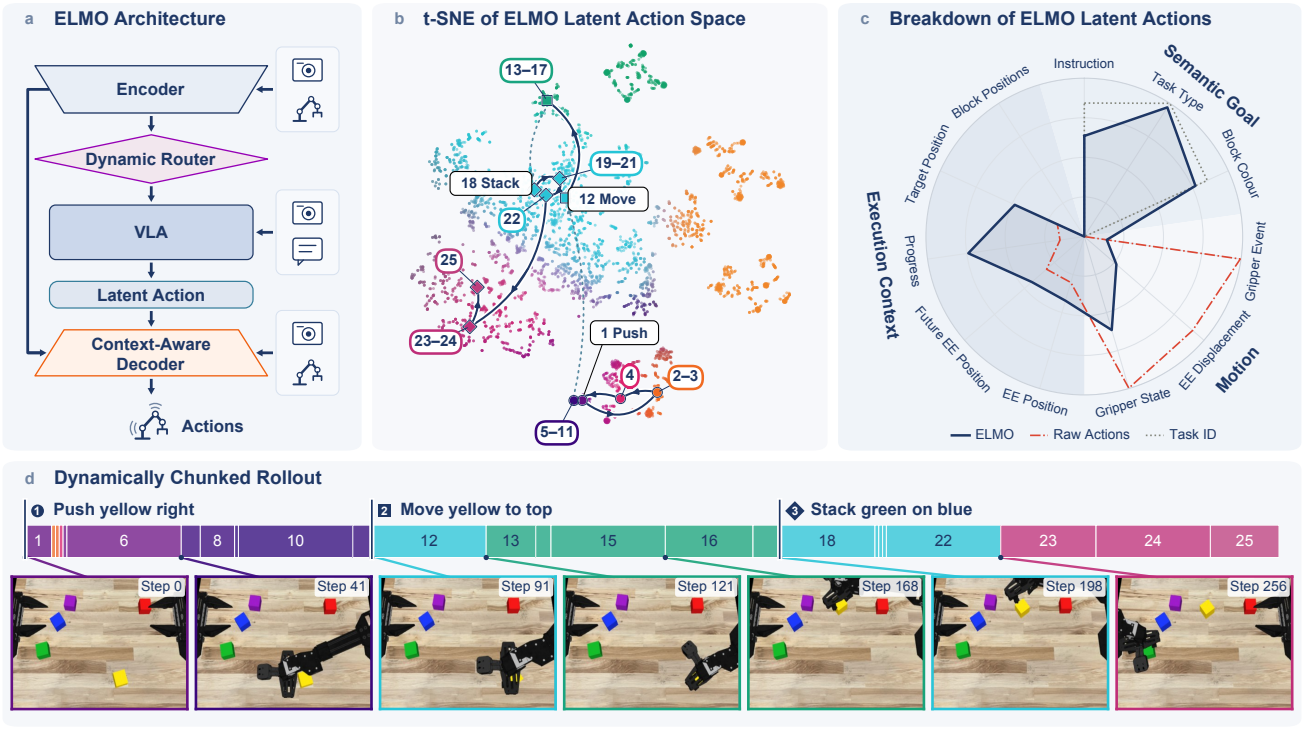

Figure 1. ELMO is a jointly-trained hierarchical policy (a) that learns semantically grounded latent actions (b,c) and segments trajectories into segments by shared intent (d). Chunk latents are traced in (b), with similar colours indicating proximity in latent space.

Abstract

Robot trajectories naturally decompose into variable-length phases with distinct semantic intents, such as reaching, grasping, and lifting. However, training policies to minimize MSE treats all actions as equal, obscuring their semantic structure. While latent action models attempt to create a structured prediction target by compressing action chunks into latent representations, existing approaches rely on two-stage training and context-blind decoders. Consequently, their representations are context-agnostic motion codes (e.g., "move down") which fail to capture behavioural intent (e.g., "pick up the block"). To overcome these bottlenecks, we introduce the Emergent Latent-action MOdel (ELMO), an end-to-end framework that encodes robot trajectories as grounded semantic codes—representations that ground task-level behavioural intent in the current environment. ELMO jointly trains a VLA planner and an observation-conditioned decoder over dynamically chunked trajectories, learning a semantically grounded action space without separate pretraining. Across standard benchmarks, ELMO outperforms comparable prior methods by 9.6 percentage points on LIBERO-Plus and +0.75 average task length on CALVIN. Furthermore, on Blocks—our shortcut-free simulation and real-world generalization benchmark—ELMO reaches a 38.6% success rate while all baselines remain below 5%, demonstrating its ability to capture the semantic structure of demonstrations. Website: https://uoft-isl.github.io/elmo.github.io/

1. Introduction

Vision-language-action (VLA) policies combine visual and language understanding with robot control by using pre-trained vision-language models (VLMs) [17]. While the concept of using a VLM backbone is well established, it is unclear how to best parametrise the action prediction task on top of it. Designs have tried predicting actions via discrete tokens [17, 31], diffusion [7, 41], or flow matching [3, 4, 39]. Further, it has been shown that making policies predict a chunk of actions rather than a single one improves performance [47], as these sequences capture the underlying non-Markovian behaviour and reduce compounding errors [18]. Going beyond the format of actions, training a VLM backbone to directly predict raw actions has been shown to hurt performance due to the large modality gap between language and raw actions [10]. Thus, research has explored alternative ways to represent actions [21, 34, 37].

At one extreme, hierarchical methods use language subtasks [1, 2, 35] as *semantic codes* to decompose instructions for a downstream controller. However, language remains too coarse, leaving excessive execution uncertainty to the action policy that lacks the full context required to keep trajectories synchronised with the original intent. At the other extreme, latent action models compress action chunks into continuous or discrete latents [19, 31, 49]. Yet, as we demonstrate in Section 4.2, these representations act merely as *motion codes*: context-agnostic encodings of kinematics without task-level intent. Ideally, an intermediate representation should combine the best of both as *grounded semantic codes*—representations that capture exact behavioural intent and environmental context while allowing an action policy to reactively adapt execution to the live scene.

Why do existing latent action models collapse into motion codes rather than learning grounded semantic codes? We argue that they are held back by two architectural bottlenecks: context-blind decoding and two-stage training. First, prior tokenizers and skill models (e.g., VQ-BeT, PRISE, and FAST [19, 31, 49]) reconstruct action sequences in isolation without an observation-conditioned decoder. Because the decoder is blind to the visual scene, the latent code is forced to memorize exact kinematics rather than abstract behavioural intent. Second, even when decoders utilize images [43], prior works adopt a hierarchical two-stage training paradigm: they first learn to compress actions into a fixed latent space in the first stage, and then train a planner to predict those frozen latents in the second step. Decoupling action compression from latent prediction prevents the latent space from being shaped by task-level semantics and visual-language context, optimizing the representation for isolated reconstruction rather than grounded planning.

Overcoming these two bottlenecks requires unifying latent prediction and observation-conditioned decoding into a single end-to-end training stage. To this end, we propose **ELMO**, an end-to-end hierarchical prediction framework that directly trains the VLA planner and the observation-conditioned decoder together, adapting the dynamic temporal chunking mechanism of H-Net [13] to continuous robot trajectories. Instead of relying on fixed temporal windows, ELMO dynamically allocates latent representations across trajectory boundaries and jointly optimizes the entire hierarchy. By training end-to-end in a single stage, the representation space is jointly shaped by vision-language planning and reactive visual control, allowing ELMO to learn a semantically grounded action space. During deployment, the VLA planner predicts the resulting *grounded semantic codes*, which the observation-conditioned action policy maps to raw controls. Grounded semantic codes serve as a compact, low-noise prediction target, improving VLM learning efficiency while preserving the reactivity and precision of receding-horizon control [7, 47].

Our contributions are threefold: (1) we introduce ELMO, an end-to-end hierarchical framework that unifies dynamic chunking with single-stage joint training of a VLA planner and an observation-conditioned decoder to learn grounded semantic codes; (2) we provide a comprehensive analysis of prior latent action methods, characterizing differences in their encoded information; and (3) we demonstrate strong performance on standard benchmarks (LIBERO, LIBERO-Plus, CALVIN) alongside superior generalization on Blocks, our specialized benchmark designed to eliminate evaluation shortcuts in existing suites.

2. Related Work

Vision-Language-Action Models. In recent years, VLA models have been extensively studied for generalist robotics. These policies leverage the general knowledge of VLMs by passing VLM embeddings of image and language to an action head for robot control [3, 17]. The action head predicts future robot actions from the current state. Following this structure, VLAs have been scaled to billions of parameters and trained on ever-increasing robot datasets [29]. Despite this scale, VLAs still struggle to generalize to unseen environments, tasks, and robot embodiments. Changes in camera viewpoint, initial robot state, or lighting can cause large drops in success rate [11]. We argue that this is partly a result of robot policies learning to predict high-frequency actions rather than the intent behind them. This results in brittle learning and poor generalization, motivating better action representations and training objectives.

Tokenization. In vision and language, models generally operate on a tokenized representation of the input rather than raw bytes directly. By grouping predictable local structures, tokenization shortens sequences and lets large models focus on long-range composition and semantic relationships rather than high-frequency noise. Vision models typically

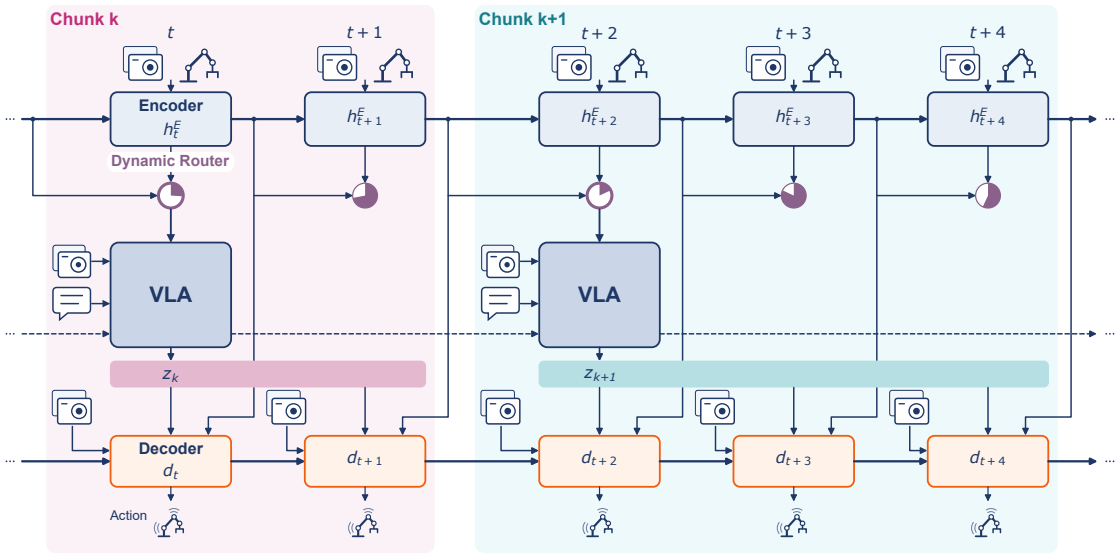

Figure 2. ELMO architecture. At each timestep, the encoder encodes the observations and robot state as h_t^E. Comparing it to h_{t-1}^E, the router determines whether a chunk should start and the VLA should predict a new latent action using the observations and language. The current latent action along with wrist image and encoder state is used by the recurrent flow-matching decoder to predict actions.

rely on learned patch embeddings trained end-to-end with the main network. Images are sliced into non-overlapping grid patches, and each flattened patch is mapped through trainable layers into a continuous latent vector [9]. Traditionally, language models have relied on separately trained tokenizers like byte-pair encoding (BPE) [33], which starts with a base alphabet and iteratively merges the most frequent adjacent symbol pairs into new tokens. More recently, learned tokenizers have been developed for language as well. H-Net [13] trains an encoder-backbone-decoder hierarchy end-to-end. Raw byte embeddings pass through an encoder that dynamically groups variable-length spans into learned tokens based on adjacent similarity. The backbone operates in this compressed token space and predictions are decoded back into raw bytes. At scales above 1B parameters, H-Net matches the training perplexity and downstream performance of BPE-tokenized transformers while removing the fixed external tokenizer [13].

Action Representations. Early VLAs discretized action dimensions into uniform bins to reframe continuous regression as autoregressive classification [17], while signal-compression methods like FAST apply the discrete cosine transform and BPE to tokenize short action chunks [31]. Alternative approaches decompose long-horizon tasks via language sub-goals [1], finer-grained language instructions [2], or fixed-hierarchy latent codes bridging VLMs and low-level controllers [34]. However, these rely on predefined linguistic abstractions and are thus limited in what they can express efficiently. Closer to our work, trajectory-level representation methods extract higher-level tokens via vector quantization and sequence compression [19, 20, 27, 45, 49], recently scaled up in VQ-VLA [40] and systematically benchmarked in ActionCodec [8]. Among these, only PRISE constructs variable-length segments via post-quantization BPE, focusing on raw compressibility rather than semantic or temporal structure [49]. Video-based latent action models (e.g., LAPA [44], HiLAM [16]) leverage unlabeled videos and optimize downstream policy integration [23], but learning from visual differences alone risks entangling robot motion with background change [22]. Most similar to our method are EDAR and SA-VLA [15, 43]. Although EDAR incorporates camera observations and predicts visual changes in its decoder, we find it still yields motion codes. Similarly, SA-VLA conditions its fixed-length decoder on proprioceptive state but without any information about the environment, it cannot encode anything besides a motion code [15]. More broadly, when decoders lack or underutilize live visual observations, latent tokens must independently reconstruct exact trajectories, collapsing into direct action reparameterizations. In contrast, we train a dynamic latent hierarchy end-to-end with a vision-aware decoder, which allows a latent representation carrying task-level intent along with its relevant context to emerge while delegating local execution to the decoder.

3. Methodology

We build **ELMO** (Figure 2) around our two core principles: dynamic chunking of the robot trajectory and a context-aware decoder. Since the trajectory chunks should reflect when the policy should make decisions, we train the hi-

erarchical VLA planner and observation-conditioned action policy together. ELMO adapts the dynamic chunking mechanism of H-Net [13] using a recurrent state encoder and a dynamic router to partition continuous robot trajectories into variable-length behavioural segments. Training all components together end-to-end via flow-matching action prediction and chunking losses allows us to learn *grounded semantic codes*. The structure emerges naturally during training and contains information about the current semantic task and execution context (Section 4.2), rather than raw motion, which is then utilised by the action policy.

3.1. Problem Setup and Architecture

Our problem setting is language-conditioned imitation learning, where policy learns to predict robot actions from an expert dataset $\mathcal{D} = \{(\ell_i, \{o_t, a_t\}_{t=0}^{T_i-1})\}_{i=1}^{N}$, where ℓ_i is a language instruction, $o_t = (I_t^{1:C}, s_t)$ denotes the RGB images from C camera views and robot proprioception s_t, and a_t is the demonstrated action. At each control step t, we aim to model a causal policy π_θ that predicts a distribution over an action horizon of length H:

$$\pi_\theta(A_t \mid o_{\le t}, \ell), \qquad A_t = (a_t, \ldots, a_{t+H-1}). \tag{1}$$

We formulate ELMO as a hierarchical encoder-decoder architecture that jointly trains a high-level VLA planner π_θ^{plan}, a low-level observation-conditioned action policy π_θ^{act}, and a dynamic routing mechanism $\pi_\theta^{\text{route}}$, forming a single end-to-end autoregressive model that uses the same causal flow during both training and inference.

Observations first pass through a recurrent state encoder, whose output is monitored by the router $\pi_\theta^{\text{route}}$ at every step t to predict a binary boundary decision $b_t \in \{0, 1\}$ (with $b_0 = 1$). Rather than querying the planner at a fixed stride, the steps where $b_t = 1$ define a sequence of chunk boundary times $\tau_k = \min\{t > \tau_{k-1} : b_t = 1\}$, partitioning the trajectory into variable-length behavioural segments. The inner VLA planner π_θ^{plan} is executed *only* at these selected boundary steps $t = \tau_k$, combining the encoder state with multi-view visual and language tokens to produce a new grounded semantic code z_k:

$$z_k \sim \pi_\theta^{\text{plan}}(\cdot \mid o_{\le \tau_k}, \ell) \quad \text{at } t = \tau_k. \tag{2}$$

At every control step t, let $k(t) = \max\{k : \tau_k \le t\}$ index the most recent boundary and $z_{k(t)}$ thus denote the active grounded semantic code. The action policy decoder π_θ^{act} combines $z_{k(t)}$ with the live encoder state and local visual feedback to predict an updated action window A_t in a receding-horizon manner, decoupling the variable chunk duration $\tau_{k+1} - \tau_k$ from the prediction horizon H:

$$A_t \sim \pi_\theta^{\text{act}}(\cdot \mid o_{\le t}, z_{k(t)}) \quad \forall t. \tag{3}$$

This design establishes two temporal pathways: a fast recurrent pathway running at the control rate and a slow causal transformer operating over dynamically learned chunks.

Recurrent State Encoder and Dynamic Router. Following H-Net [13], at each step t, wrist-camera tokens are spatially pooled and flattened into a compact visual summary w_t. An MLP embeds w_t alongside proprioception s_t. To decrease reliance on precise proprioception, we add $\sigma = 0.1$ noise to it. A recurrent state encoder E composed of residual GRU–MLP blocks updates the low-level hidden state:

$$e_t = f_{\text{emb}}(s_t, w_t), \qquad h_t^E = E(e_t, h_{t-1}^E). \tag{4}$$

To decide whether the current state warrants a latent plan update, a learned router measures the cosine dissimilarity between consecutive encoder states in a projected space:

$$p_t = \frac{1 - \cos(W_q h_t^E, W_k h_{t-1}^E)}{2}. \tag{5}$$

We set $p_0 = b_0 = 1$ and trigger a new chunk boundary ($b_t = 1$) whenever $p_t \ge 0.5$ or a maximum step timeout is reached. Retaining the encoder states at boundary steps $\{\tau_k\}$ yields a temporally compressed sequence.

VLA Planner. At each boundary τ_k, the selected encoder state $h_{\tau_k}^E$ is augmented with a learned vector of additional dimensions, widening it to the planner width, to form a query, following H-Net [13]. This query drives two cross-attention blocks that attend over the full set of current static- and wrist-camera tokens and the language instruction ℓ, all of which are jointly encoded by the SmolVLM2 [25] backbone. To preserve spatial structure during readout, visual tokens are enriched with 2D coordinate encodings. The resulting multimodal boundary token u_k is processed by a causal transformer T over preceding chunks in the episode:

$$z_k = \text{Proj}\left[\text{RMSNorm}\big(T(u_0, \ldots, u_k)_k\big)\right], \tag{6}$$

where Proj removes the dimensions added at the start of the planner, truncating the representation back to the encoder width. The vector z_k serves as a continuous grounded semantic code that summarizes behavioural intent and long-horizon task progress at a compressed temporal resolution.

Action Policy. The observation-conditioned action policy translates the active grounded semantic code $z_{k(t)}$ into raw actions. At each step t, the active $z_{k(t)}$ is combined with the live encoder state h_t^E via a residual skip connection [13] to form the decoder input d_t. A recurrent decoder processes d_t, and its output cross-attends to the wrist-camera to supply fine-grained, up-to-date local visual feedback for reactive manipulation while the high-level code remains fixed. The resulting feature h_t^D conditions a flow-matching action head $v_\theta(X_t^u, u, h_t^D)$ to generate the action window A_t. Specifically, for flow time $u \in [0, 1]$ and Gaussian noise $\epsilon_t \sim \mathcal{N}(0, I)$, the action head predicts the velocity field

transporting ϵ_t to A_t along the linear interpolation path $X_t^u = (1-u)\epsilon_t + uA_t$; during inference, A_t is sampled by integrating v_θ from $u = 0$ to $u = 1$ over five Euler steps.

3.2. Training and Inference

We jointly optimize the recurrent state encoder, router, VLA planner, action policy decoder, and action head end-to-end while keeping the visual backbone frozen.

Boundary-Weighted Flow Matching. During training, we sample $u \sim \mathcal{U}[0,1]$ and $\epsilon_t \sim \mathcal{N}(0, I)$ to train the action head $v_\theta(X_t^u, u, h_t^D)$ to predict the target velocity $A_t - \epsilon_t$. Let $l_{t,j}$ denote the mean squared velocity error at horizon offset j. We minimize the boundary-weighted flow-matching objective:

$$\mathcal{L}_{\text{act}} = \mathbb{E}_{u,\epsilon}\left[\frac{\sum_{t,j} m_{t,j} w_t l_{t,j}}{\sum_{t,j} m_{t,j} w_t}\right], \quad w_t = 1 + \alpha\ \text{sg}(b_t) \quad (7)$$

where $m_{t,j}$ masks invalid steps, $\text{sg}(\cdot)$ is the stop-gradient operator, and $\alpha = 4$ upweights action windows predicted immediately at chunk boundaries so the VLA planner learns informative updates. To allow action gradients to train the discrete router, a confidence-weighted straight-through estimator propagates gradients to p_t.

Router Regularisation. To prevent trivial routing, we use the H-Net approach [13] and regularise the router toward a target compression ratio of C_{ratio} control steps per chunk:

$$\mathcal{L}_{\text{ratio}} = \frac{C_{\text{ratio}}}{C_{\text{ratio}} - 1}\left[(C_{\text{ratio}} - 1)FG + (1-F)(1-G)\right], \quad (8)$$

where F is the empirical fraction of selected boundaries and G is the mean boundary probability across the batch. During early training, to improve routing stability, we also guide the router with a binary cross-entropy warm-start loss $\mathcal{L}_{\text{warm}}$ derived from heuristic kinematic boundaries, whose weight, w_{warm}, linearly decays over the first S_{warm} steps:

$$\mathcal{L} = \mathcal{L}_{\text{act}} + \lambda_{\text{ratio}}\mathcal{L}_{\text{ratio}} + w_{\text{warm}}\mathcal{L}_{\text{warm}}. \quad (9)$$

History Regularization. We incorporate two data-level mechanisms during training to address the drawbacks of history-conditioned policies. First, unlike language, robot behaviour is highly autocorrelated, which makes it easy to achieve low loss simply by extrapolating from recent history [31]. This is especially problematic for the VLA planner, as much of the plan can be inferred from history once the episode has started. We therefore do *onset-prioritised sampling* and focus training on episodes where the planner struggles to choose the correct initial plan, forcing it to rely more heavily on visual and task information. We track an exponential moving average L_i of each demonstration's unweighted loss over its first K_{init} steps and sample training trajectories with probability

$$q_i = \frac{1}{2N} + \frac{1}{2}\frac{L_i^2}{\sum_{r=1}^{N} L_r^2}, \quad (10)$$

balancing uniform dataset coverage with focused updates on difficult task openings.

Second, the policy can use its history to store object locations from earlier in the rollout, removing the need to re-observe the scene. While this shortcut succeeds during training, it makes the policy brittle during inference. We therefore break single-trajectory continuity by constructing training sequences from spliced fragments of different demonstrations that share the same instruction ℓ, without resetting hidden states or chunk history at the joins. Target masks $m_{t,j}$ prevent supervision windows from crossing artificial splice points. This limits the policy from relying on memory for long-term state tracking while still allowing it to use short-term history for smooth, precise execution.

Inference. At each step t, the encoder updates h_t^E and the router evaluates Eq. (5). The VLA planner predicts, only when the router places a new boundary. The flow head integrates noise over five Euler steps to produce A_t, and the robot executes the first action a_t before acquiring o_{t+1}.

4. Experiments

We evaluate ELMO across two standard manipulation benchmarks (LIBERO [24] and CALVIN [26]), a zero-shot robustness benchmark (LIBERO-Plus [11]), a new shortcut-free compositional benchmark (Blocks), and a physical bimanual robot. We first analyze the learned latent spaces (Section 4.2), showing that while prior action codecs primarily reparameterize low-level kinematics, ELMO learns grounded semantic codes that capture task intent, target identity, and execution context. We then demonstrate that this structured abstraction allows us to outperform alternative latent codes paired with a similar-sized backbone across standard and perturbed simulation benchmarks (Section 4.3), enables instruction following where existing methods fail (Section 4.4), and transfers effectively to real hardware (Section 4.5).

4.1. Benchmarks and Baselines

Benchmarks. We first evaluate on two established simulation benchmarks: LIBERO [24] and CALVIN [26] (see supplementary for details). LIBERO comprises four suites—Spatial, Object, Goal, and LIBERO-10 — targeting distinct facets of language-conditioned generalization, while CALVIN (ABCD $\rightarrow$ D) evaluates long-horizon execution and chaining of independent skills. However, studies show that models can solve these standard

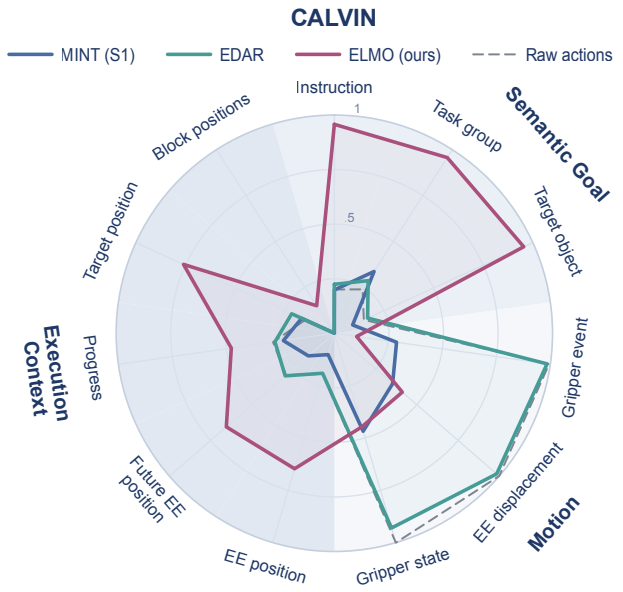

Figure 3. Linear probe performance across kinematic and semantic attributes. ELMO encodes more semantic and context information and less motion than all baselines.

benchmarks via spurious shortcuts—inferring task intent from initial scene layouts or gripper proximity rather than genuinely grounding the language instruction—leaving them brittle under distribution shifts [11, 14, 32, 42]. To rigorously test robustness and causal instruction following, we additionally evaluate zero-shot on LIBERO-Plus [11], which perturbations to LIBERO, and introduce our shortcut-free Blocks benchmark (Section 4.4).

Baselines. We compare ELMO against SmolVLA [36], which directly predicts continuous actions, and MINT-30M [12], a similarly sized hierarchical policy using latent actions. We also compare against alternative latent action tokenizers combined with SmolVLA [36]. Following the protocol of EDAR [43], we evaluate pretrained action codecs—FAST [31] and EDAR [43]—conditioned on the same SmolVLM2 [25] backbone (see supplementary for implementation details). All methods are evaluated closed-loop on identical test seeds across 3 training seeds.

4.2. Latent Space Analysis

To understand what information ELMO, MINT, and EDAR capture in their latent representations across LIBERO, CALVIN, and Blocks, we fit ridge-regression linear probes to latent plans extracted from demonstration trajectories (for details see supplementary). As shown in Figure 3, MINT and EDAR latents largely encode the same information as raw action—concentrated in end-effector displacements and gripper states—functioning as *motion codes* that

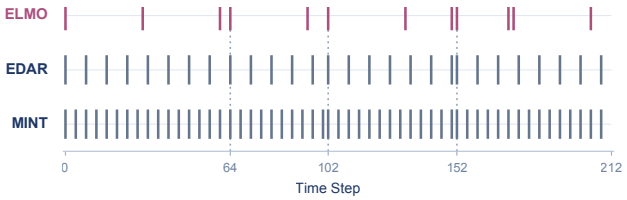

Figure 4. Dynamic chunk boundaries placed by ELMO compared to the static replan rate of MINT and EDAR along a CALVIN rollout (gray lines indicate new subtask).

reparameterize kinematics. In contrast, ELMO's latents are highly predictive of target objects, active skills, and task progress, while retaining contextual execution cues beyond the static task label. We therefore characterize ELMO's emergent representations as *grounded semantic codes*: they specify not only *what* semantic skill to execute, but *how* to ground it in the current scene. This distinction is echoed in the t-SNE plots (Figure 5). Whereas MINT and EDAR latents organize by movement direction and gripper state, ELMO's latents naturally cluster by skills and target tasks without explicit contrastive supervision.

Finally, visualizing ELMO's learned chunk boundaries (Figures 1 and 4) shows that plan updates naturally align with semantic sub-phases (e.g., reaching vs. grasping). ELMO dynamically allocates shorter chunks during contact-rich, high-precision phases like grasping and longer chunks during coarse transit, utilizing high-level queries more efficiently than fixed-stride methods.

Together, these properties yield substantial parameter efficiency without sacrificing performance. Because its latent space is both semantically grounded and adaptively segmented, ELMO outperforms baselines on manipulation benchmarks (Section 4.3) with only 22M unfrozen parameters across the encoder, router, VLA action head, and decoder—compared to 100M in SmolVLA's action head.

4.3. Main Simulation Results

Table 1 summarizes performance across LIBERO, LIBERO-Plus, and CALVIN. On LIBERO, ELMO achieves a 90.2% average success rate, outperforming SmolVLA (85.7%) and the best action codec (EDAR, 88.6%) while matching MINT-30M (90.2%) without relying on image augmentations (without which MINT-30M drops to 84.4%). The advantage of grounded semantic abstraction is most pronounced on long-horizon CALVIN (ABCD $\rightarrow$ D) chains, where ELMO completes an average of 4.08 consecutive subtasks—surpassing the strongest baseline (SmolVLA + FAST) by $+0.75$ subtasks.

On LIBERO-Plus, ELMO achieves 64.0% zero-shot success. While MINT-30M achieves 61.0%, it requires using image augmentation to do so. When comparing across methods without such augmentations, ELMO outperforms

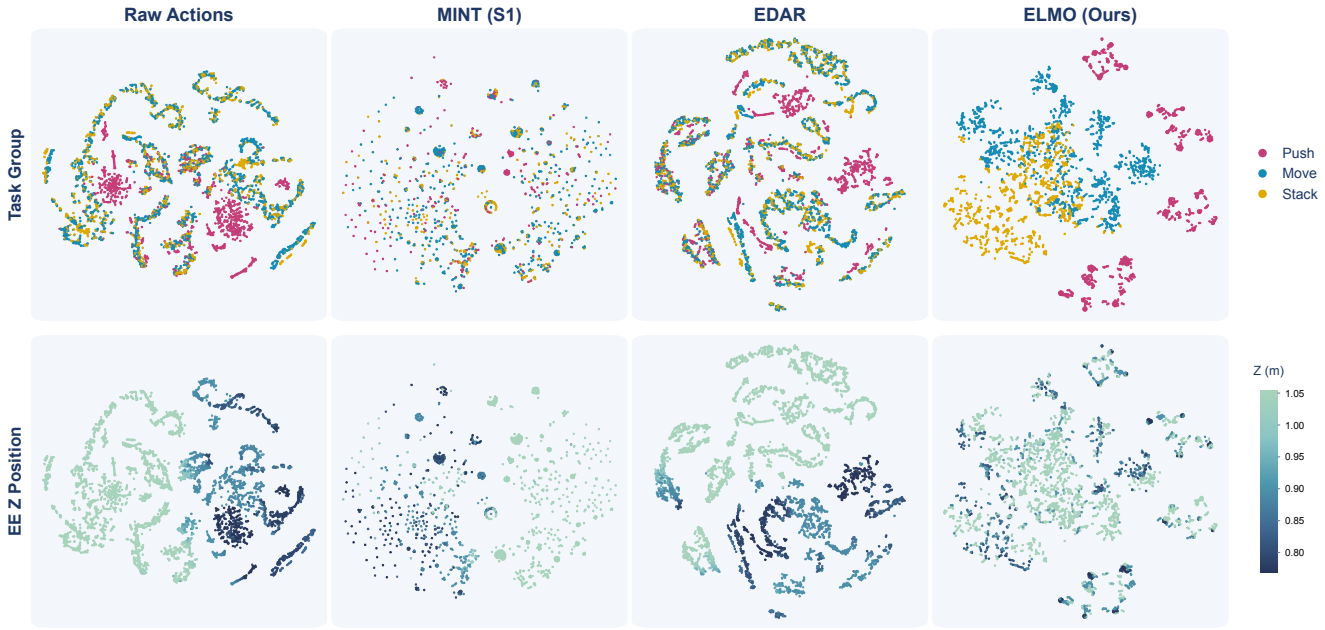

Figure 5. t-SNE projections comparing latent representations of MINT [12], EDAR [43], and ELMO on Blocks. For MINT, we plot the single intent tokens. EDAR's t-SNE structure is very close to that of raw actions and MINT's structure shows no organization by task. ELMO shows clear organization by task. Representations overlap as both push and stack trajectories share similar approach subsequences.

by 9.6 percentage points over the next best method, with the largest gains under Camera, Language, and Noise perturbations (see supplementary for details).

4.4. Compositional Grounding and Generalization

To isolate instruction grounding from scene-layout shortcuts in a real-world-transferable setting (see Table in supplementary), we introduce Blocks, a bimanual ALOHA MuJoCo benchmark [38, 48] featuring two opposing ViperX arms, three RGB cameras, and five distinctly coloured cubes placed randomly in every scene. Tasks span three families—directional pushing, stacking, and regional placement—with targets sampled independently of gripper proximity, occasionally requiring cross-table bimanual handovers. We generate 10,000 clean scripted demonstrations and augment them with 2,500 recovery trajectories in which an initially misaligned approach is actively corrected.

Together, these design choices force policies to rely on genuine language grounding rather than spatial shortcuts, providing a rigorous measure of generalization. Compounding this challenge, the 5 cm cubes afford only a 0.7 cm total tolerance within the 5.7 cm gripper opening, requiring fine-grained manipulation precision.

As shown in Table 2, all baseline policies collapse on Blocks ($\leq$ 5.0% ID, $\leq$ 2.0% OOD), failing to disambiguate which of the five cubes to manipulate when proximity shortcuts are removed. In contrast, ELMO achieves 38.6% ID and 31.7% OOD success by grounding the target colour in its high-level plan and purposefully approaching the instructed block. Remaining failures occur during fine-grained motion rather than target selection, reflecting the challenge of learning high-precision contact.

4.5. Real-World Evaluation

We deploy ELMO on a physical bimanual ALOHA setup mirroring our Blocks benchmark (Figure 6). Given the low success rates observed for stacking in simulation, we focus our real-world evaluation on the pushing and regional placement families. The two task families, five coloured blocks, and four directions yield 40 distinct task combinations. We fine-tune each policy from its simulation checkpoint for 10k steps using 30 human demonstrations per combination, and evaluate across 25 randomized pushing and placement trials (50 rollouts total; see supplementary for full training and deployment details). Overcoming substantial sim-to-real discrepancies in camera viewpoints, scene appearance, and physical dynamics, ELMO achieves 13/25 on push trials and 8/25 on regional placement trials, outperforming SmolVLA (2/25 push, 0/25 move).

4.6. Ablation Studies

Ablations train on CALVIN with 16,000 steps and 3 seeds.

Router Warm-Start. Removing the kinematic router warm-start ($\mathcal{L}_{\text{warm}}$) lowers average completed tasks only slightly, from 4.04 to 3.90, but it changes where the router

Table 1. **Main simulation benchmark results.** LIBERO and LIBERO-Plus report average episode success rates (%). CALVIN reports average consecutively completed subtasks (out of 5). Entries report mean ± standard error across 3 seeds.

	LIBERO ↑	LIBERO-Plus ↑					CALVIN ABCD→D ↑
Method	Avg.	Spatial	Object	Goal	Long	Avg.	Avg. Len.
SmolVLA [36]	85.7 ± 0.6	56.4 ± 0.9	56.1 ± 0.7	51.7 ± 1.1	36.9 ± 1.3	50.2 ± 0.8	3.07 ± 0.03
SmolVLA + FAST codec [31]	83.4 ± 1.2	55.5 ± 1.1	59.4 ± 1.7	48.1 ± 0.6	37.9 ± 1.6	50.1 ± 0.3	3.33 ± 0.05
SmolVLA + EDAR codec [43]	88.6 ± 0.8	63.0 ± 0.5	61.5 ± 0.7	52.8 ± 0.8	40.7 ± 0.9	54.4 ± 0.2	3.32 ± 0.02
MINT-30M [12]	$\mathbf{90.2 \pm 0.8}$	65.9 ± 0.9	70.8 ± 1.1	$\mathbf{61.1 \pm 0.9}$	46.5 ± 2.4	61.0 ± 1.3	2.68 ± 0.11
MINT-30M (w/o image aug.) [12]	84.4 ± 0.7	54.2 ± 1.9	53.0 ± 2.4	50.9 ± 1.4	33.4 ± 0.9	47.8 ± 1.3	–
ELMO (Ours)	$\mathbf{90.2 \pm 0.4}$	$\mathbf{72.6 \pm 1.0}$	$\mathbf{71.9 \pm 1.1}$	58.9 ± 0.1	$\mathbf{53.2 \pm 1.5}$	$\mathbf{64.0 \pm 0.8}$	$\mathbf{4.08 \pm 0.04}$

Table 2. **Blocks benchmark results.** Success rate (%) on in-distribution (ID) vs. held-out colour–position grid cells (OOD). Entries report mean ± std error across 3 seeds.

Method	ID (%) ↑	OOD (%) ↑
SmolVLA [36]	4.9 ± 0.7	0.7 ± 0.4
SmolVLA + FAST codec [31]	3.2 ± 0.4	0.5 ± 0.1
SmolVLA + EDAR codec [43]	4.7 ± 0.4	2.0 ± 0.5
MINT-30M [12]	2.0 ± 0.6	1.1 ± 0.3
ELMO (Ours)	$\mathbf{38.6 \pm 2.9}$	$\mathbf{31.7 \pm 5.1}$

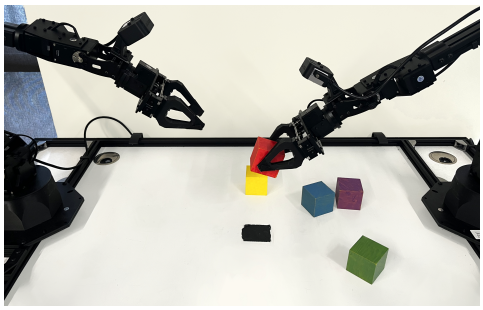

Figure 6. Real-world bimanual ALOHA evaluation setup [48].

places its boundaries. On held-out demonstrations, the warm-started router cuts within 2 steps of 82% of grasps, 3.8× the 21% expected from the same number of randomly placed boundaries, and the nearest cut usually falls exactly on the grasp. Without the warm-start this drops to 39%. The remaining cuts tend to come late, up to 11 steps after the grasp. The router thus discovers some event-aligned boundaries on its own, but the warm-start makes boundary placement at contact events consistent and precise.

Low-Level Feedback and Data Augmentation. Removing control-rate wrist cross-attention or the encoder skip connection reduces average completed tasks by on average 1.05 and 0.58, respectively, demonstrating that continuous low-level visual feedback is critical for the decoder. Disabling trajectory splicing and onset-prioritized sampling reduces average completed tasks by 0.79 as the model overly relies on its history and ignores the visual observations.

Compression and Dynamic Routing. Replacing the learned router with fixed-length chunks while maintaining $C_{\text{ratio}} = 12$ reduces average completed tasks from 4.04 to 3.89. This means dynamic routing does help align the chunks with current intentions but is not critical.

5. Conclusion and Limitations

We present **ELMO**, an end-to-end hierarchical VLA policy that avoids the collapse of prior latent action model representations into *motion codes*. By pairing dynamic temporal routing with an observation-conditioned flow-matching decoder, ELMO learns *grounded semantic codes* that capture task-level behavioural intent and environmental context while delegating reactive control to a lightweight action policy. On both LIBERO-Plus and CALVIN, ELMO achieves the best performance among the evaluated models with a fraction of the trainable parameters, and uniquely succeeds at instruction following on our shortcut-free Blocks benchmark in both simulation and the real world.

Limitations. First, while dynamic routing emerges end-to-end, stable boundary alignment with contact events still benefits from an early kinematic warm-start ($\mathcal{L}_{\text{warm}}$) and a predefined target compression ratio C_{ratio}. Second, on Blocks, although ELMO resolves target disambiguation, it still struggles with sub-centimeter grasping and stacking precision—suggesting the need for closed-loop RL fine-tuning or higher-resolution local sensing. Finally, extending ELMO's recurrent state encoder to cross-embodiment pretraining and unlabeled video data remains an important direction for future work.

6. LLM Usage Statement

LLM agents were used in the research process to write code, create visualizations, and monitor experiments. All results were manually verified by the authors. LLMs were also used to refine the paper submission to improve writing quality.

References

[1] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, et al. Do as I Can, Not as I Say: Grounding Language in Robotic Affordances. *arXiv preprint arXiv:2204.01691*, 2022. 2, 3

[2] Suneel Belkhale, Tianli Ding, Ted Xiao, Pierre Sermanet, Quan Vuong, Jonathan Tompson, Yevgen Chebotar, Debidatta Dwibedi, and Dorsa Sadigh. RT-H: Action Hierarchies Using Language. In *Robotics: Science and Systems (RSS)*, 2024. 2, 3

[3] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π_0: A Vision-Language-Action Flow Model for General Robot Control. *arXiv preprint arXiv:2410.24164*, 2024. 2

[4] Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Manuel Y. Galliker, Dibya Ghosh, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Devin LeBlanc, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Allen Z. Ren, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz, James Tanner, Quan Vuong, Homer Walke, Anna Walling, Haohuan Wang, Lili Yu, and Ury Zhilinsky. $\pi_{0.5}$: A Vision-Language-Action Model with Open-World Generalization. In *Conference on Robot Learning (CoRL)*. PMLR, 2025. 2

[5] Remi Cadene, Simon Alibert, Alexander Soare, Quentin Gallouedec, Adil Zouitine, Steven Palma, Pepijn Kooijmans, Michel Aractingi, Mustafa Shukor, Dana Aubakirova, Martino Russi, Francesco Capuano, Caroline Pascal, Jade Choghari, Khalil Meftah, Maxime Ellerbach, Jess Moss, and Thomas Wolf. Lerobot: State-of-the-art machine learning for real-world robotics in pytorch. https://github.com/huggingface/lerobot, 2024. 1, 2

[6] Tianxing Chen, Zanxin Chen, Baijun Chen, Zijian Cai, Yibin Liu, Zixuan Li, Qiwei Liang, Xianliang Lin, Yiheng Ge, Zhenyu Gu, Weiliang Deng, Yubin Guo, Tian Nian, Xuanbing Xie, Qiangyu Chen, Kailun Su, Tianling Xu, Guodong Liu, Mengkang Hu, Huan-Ang Gao, Kaixuan Wang, Zhixuan Liang, Yusen Qin, Xiaokang Yang, Ping Luo, and Yao Mu. RoboTwin 2.0: A Scalable Data Generator and Benchmark with Strong Domain Randomization for Robust Bimanual Robotic Manipulation. In *International Conference on Machine Learning (ICML)*. PMLR, 2026. 6

[7] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion Policy: Visuomotor Policy Learning via Action Diffusion. *The International Journal of Robotics Research*, 2025. 2

[8] Zibin Dong, Yicheng Liu, Shiduo Zhang, Baijun Ye, Yifu Yuan, Fei Ni, Jingjing Gong, Xipeng Qiu, Hang Zhao, Yinchuan Li, et al. ActionCodec: What Makes for Good Action Tokenizers. In *Conference on Robot Learning (CoRL)*, 2026. 3

[9] Alexey Dosovitskiy. An Image Is Worth 16x16 Words: Transformers for Image Recognition at Scale. *arXiv preprint arXiv:2010.11929*, 2020. 3

[10] Danny Driess, Jost Springenberg, Brian Ichter, Lili Yu, Adrian Li-Bell, Karl Pertsch, Allen Ren, Homer Walke, Quan Vuong, Lucy Xiaoyang Shi, et al. Knowledge Insulating Vision-Language-Action Models: Train Fast, Run Fast, Generalize Better. *Advances in Neural Information Processing Systems*, 2026. 2

[11] Senyu Fei, Siyin Wang, Junhao Shi, Zihao Dai, Jikun Cai, Pengfang Qian, Li Ji, Xinzhe He, Shiduo Zhang, Zhaoye Fei, et al. LIBERO-Plus: In-Depth Robustness Analysis of Vision-Language-Action Models. *arXiv preprint arXiv:2510.13626*, 2025. 2, 5, 6, 1

[12] Renming Huang, Chendong Zeng, Wenjing Tang, Jintian Cai, Cewu Lu, and Panpan Cai. Mimic Intent, Not Just Trajectories. *arXiv preprint arXiv:2602.08602*, 2026. 6, 7, 8, 4

[13] Sukjun Hwang, Brandon Wang, and Albert Gu. Dynamic Chunking for End-to-End Hierarchical Sequence Modeling. *arXiv preprint arXiv:2507.07955*, 2025. 2, 3, 4, 5

[14] Tianchong Jiang, Xiangshan Tan, Samuel Wheeler, Luzhe Sun, Tewodros W Ayalew, and Matthew Walter. What Are We Actually Benchmarking in Robot Manipulation? In *Conference on Robot Learning (CoRL)*, 2026. 6

[15] Tengyue Jiang, Chunpu Xu, Jiayue Kang, and Yao Mu. SA-VLA: State-Aware Tokenizer for Improving Vision-Language-Action Models' Performance. *arXiv preprint arXiv:2606.30113*, 2026. 3

[16] Hanjung Kim, Lerrel Pinto, and Seon Joo Kim. Hierarchical Latent Action Model. *arXiv preprint arXiv:2603.05815*, 2026. 3

[17] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. OpenVLA: An Open-Source Vision-Language-Action Model. *arXiv preprint arXiv:2406.09246*, 2024. 2, 3

[18] Filippo Lazzati, Kyle Stachowicz, William Chen, Alberto Maria Metelli, Andrew Wagenmaker, and Sergey Levine. Why Does Action Chunking Improve Behavioral Cloning Performance in Robotic Control? In *Conference on Robot Learning (CoRL)*, 2026. 2

[19] Seungjae Lee, Yibin Wang, Haritheja Etukuru, H. Jin Kim, Nur Muhammad Mahi Shafiullah, and Lerrel Pinto. Behavior Generation with Latent Actions. *arXiv preprint arXiv:2403.03181*, 2024. 2, 3

[20] Hao Li, Qi Lv, Rui Shao, Xiang Deng, Yinchuan Li, Jianye Hao, and Liqiang Nie. STAR: Learning Diverse Robot Skill

Abstractions through Rotation-Augmented Vector Quantization. In *International Conference on Machine Learning (ICML)*, 2025. 3

[21] Yi Li, Yuquan Deng, Jesse Zhang, Joel Jang, Marius Memmel, Caelan Garrett, Fabio Ramos, Dieter Fox, Anqi Li, Abhishek Gupta, et al. HAMSTER: Hierarchical Action Models for Open-World Robot Manipulation. In *International Conference on Learning Representations (ICLR)*, 2025. 2

[22] Zuolei Li, Xingyu Gao, Xiaofan Wang, and Jianlong Fu. LatBot: Distilling Universal Latent Actions for Vision-Language-Action Models. *arXiv preprint arXiv:2511.23034*, 2025. 3

[23] Yihan Lin, Haoyang Li, Yang Li, Haitao Shen, Yihan Zhao, Chao Shao, and Jing Zhang. From Pixels to Tokens: A Systematic Study of Latent Action Supervision for Vision-Language-Action Models. In *International Conference on Machine Learning (ICML)*, 2026. 3

[24] Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. LIBERO: Benchmarking Knowledge Transfer for Lifelong Robot Learning. *arXiv preprint arXiv:2306.03310*, 2023. 5, 1, 6

[25] Andrés Marafioti, Orr Zohar, Miquel Farré, Merve Noyan, Elie Bakouch, Pedro Cuenca, Cyril Zakka, Loubna Ben Allal, Anton Lozhkov, Nouamane Tazi, Vaibhav Srivastav, Joshua Lochner, Hugo Larcher, Mathieu Morlon, Lewis Tunstall, Leandro von Werra, and Thomas Wolf. SmolVLM: Redefining Small and Efficient Multimodal Models. *arXiv preprint arXiv:2504.05299*, 2025. 4, 6, 2

[26] Oier Mees, Lukas Hermann, Erick Rosete-Beas, and Wolfram Burgard. CALVIN: A Benchmark for Language-Conditioned Policy Learning for Long-Horizon Robot Manipulation Tasks. *arXiv preprint arXiv:2112.03227*, 2021. 5, 1, 2, 6

[27] Atharva Mete, Haotian Xue, Albert Wilcox, Yongxin Chen, and Animesh Garg. QueST: Self-Supervised Skill Abstractions for Learning Continuous Control. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2024. 3

[28] Jeremy Morgan, Hyeonho Oh, Prajwal Vijay, Jincen Song, Ashvin Arora, Hojung Lim, Alina Du, Jesse Thomason, Gaurav S. Sukhatme, and Ishika Singh. Colosseum V2: Benchmarking Generalization for Vision-Language-Action Models. *IEEE Robotics and Automation Letters*, 2026. 6

[29] Abby O'Neill, Abdul Rehman, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, et al. Open X-Embodiment: Robotic Learning Datasets and RT-X Models: Open X-Embodiment Collaboration 0. In *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024. 2

[30] Seohong Park and Sergey Levine. Behavioral Cloning Mystery. *arXiv preprint arXiv:2610.07056*, 2026. 6

[31] Karl Pertsch, Kyle Stachowicz, Brian Ichter, Danny Driess, Suraj Nair, Quan Vuong, Oier Mees, Chelsea Finn, and Sergey Levine. FAST: Efficient Action Tokenization for Vision-Language-Action Models. *arXiv preprint arXiv:2501.09747*, 2025. 2, 3, 5, 6, 8

[32] Kohei Sendai, Tatsuya Matsushima, and Yusuke Iwasawa. MINERVA: How Small Can a Manipulation Policy Be and Still Solve LIBERO?, 2026. 6

[33] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural Machine Translation of Rare Words with Subword Units. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2016. 3

[34] Yide Shentu, Philipp Wu, Aravind Rajeswaran, and Pieter Abbeel. From LLMs to Actions: Latent Codes as Bridges in Hierarchical Robot Control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024. 2, 3

[35] Lucy Xiaoyang Shi, Brian Ichter, Michael Robert Equi, Liyiming Ke, Karl Pertsch, Quan Vuong, James Tanner, Anna Walling, Haohuan Wang, Niccolo Fusai, Adrian Li-Bell, Danny Driess, Lachy Groom, Sergey Levine, and Chelsea Finn. Hi Robot: Open-Ended Instruction Following with Hierarchical Vision-Language-Action Models. In *International Conference on Machine Learning (ICML)*, 2025. 2

[36] Mustafa Shukor, Dana Aubakirova, Francesco Capuano, Pepijn Kooijmans, Steven Palma, Adil Zouitine, Michel Aractingi, Caroline Pascal, Martino Russi, Andres Marafioti, et al. SmolVLA: A Vision-Language-Action Model for Affordable and Efficient Robotics. *arXiv preprint arXiv:2506.01844*, 2025. 6, 8, 2, 3

[37] Richard S Sutton, Doina Precup, and Satinder Singh. Between MDPs and Semi-MDPs: A Framework for Temporal Abstraction in Reinforcement Learning. *Artificial Intelligence*, 1999. 2

[38] Emanuel Todorov, Tom Erez, and Yuval Tassa. MuJoCo: A Physics Engine for Model-Based Control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2012. 7, 5

[39] Qiuyue Wang, Mingsheng Li, Jian Guan, Jinhui Ye, Sicheng Xie, Yitao Liu, Junhao Chen, Zhixuan Liang, Jie Zhang, Xintong Hu, et al. Qwen-VLA: Unifying Vision-Language-Action Modeling across Tasks, Environments, and Robot Embodiments. *arXiv preprint arXiv:2605.30280*, 2026. 2

[40] Yating Wang, Haoyi Zhu, Mingyu Liu, Jiange Yang, Hao-Shu Fang, and Tong He. VQ-VLA: Improving Vision-Language-Action Models via Scaling Vector-Quantized Action Tokenizers. In *IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE, 2025. 3

[41] Junjie Wen, Yichen Zhu, Minjie Zhu, Zhibin Tang, Jinming Li, Zhongyi Zhou, Xiaoyu Liu, Chaomin Shen, Yaxin Peng, and Feifei Feng. DiffusionVLA: Scaling Robot Foundation Models via Unified Diffusion and Autoregression. In *International Conference on Machine Learning (ICML)*, 2025. 2

[42] Youguang Xing, Xu Luo, Junlin Xie, Lianli Gao, Heng Tao Shen, and Jingkuan Song. Shortcut Learning in Generalist Robot Policies: The Role of Dataset Diversity and Fragmentation. In *Conference on Robot Learning (CoRL)*. PMLR, 2025. 6

[43] Yuecheng Xu, Tong Yang, Jingkai Jia, Chi Zhang, Xuelong Li, and Wenqiang Zhang. EDAR: Learning Environment-Dependent Action Representations for Robotic Manipulation. *arXiv preprint arXiv:2607.11427*, 2026. 2, 3, 6, 7, 8

[44] Seonghyeon Ye, Joel Jang, Byeongguk Jeon, Sejune Joo, Jianwei Yang, Baolin Peng, Ajay Mandlekar, Reuben Tan, Yu-Wei Chao, Bill Yuchen Lin, Lars Liden, Kimin Lee, Jianfeng Gao, Luke Zettlemoyer, Dieter Fox, and Minjoon Seo. Latent Action Pretraining from Videos. *arXiv preprint arXiv:2410.11758*, 2024. 3

[45] David Yunis, Justin Jung, Falcon Z Dai, and Matthew R Walter. Subwords as Skills: Tokenization for Sparse-Reward Reinforcement Learning. *Advances in Neural Information Processing Systems*, 2024. 3

[46] Mengao Zhao, Ziang Li, Chaodong Huang, Mengchen Ma, Haoyi Jiang, Yiwei Jin, Xinjie Wang, Yun Du, Xuewu Lin, Taojun Ding, et al. InstructMove: A Text-Indispensable Benchmark for Instruction-Following Manipulation. *arXiv preprint arXiv:2608.22990*, 2026. 6

[47] Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. *arXiv preprint arXiv:2304.13705*, 2023. 2, 7

[48] Tony Z. Zhao, Jonathan Tompson, Danny Driess, Pete Florence, Seyed Kamyar Seyed Ghasemipour, Chelsea Finn, and Ayzaan Wahid. ALOHA Unleashed: A Simple Recipe for Robot Dexterity. In *Conference on Robot Learning (CoRL)*, 2024. 7, 8, 5

[49] Ruijie Zheng, Ching-An Cheng, Hal Daumé, III, Furong Huang, and Andrey Kolobov. PRISE: LLM-Style Sequence Compression for Learning Temporal Action Abstractions in Control. *arXiv preprint arXiv:2402.10450*, 2024. 2, 3

ELMO: Emergent Representations In Imitation Learning Via Dynamic Hierarchical Action Chunking

Supplementary Material

A. ELMO Implementation & Training Details

In CALVIN and LIBERO, o_t pairs one scene and one wrist camera with an 8D state s_t, and a_t is 7D (relative end-effector pose and binary gripper); in Blocks, o_t pairs a cropped overhead and two wrist cameras with a 20D state s_t, and a_t is 14D (bimanual joint-position deltas and absolute gripper targets). Both s_t and a_t are z-score normalized. Images are encoded by the frozen SmolVLM2 vision tower into an 8×8 token grid per camera. In the control-rate encoder E, wrist tokens are pooled to 4×4, projected to 32D, flattened into w_t, and embedded with s_t via a 2-layer MLP f_{emb}. At each boundary τ_k, the frozen first 16 SmolVLM2 language layers process the boundary camera tokens (with start/end delimiters) and instruction ℓ, projecting the 960D output to $D_{\text{main}} = 384$ and adding linear camera-ID and sinusoidal 2D grid encodings. The extended boundary state $h^E_{\tau_k}$ queries these tokens via pre-norm cross-attention; the resulting u_k adds a learned chunk-position embedding and passes through the causal transformer T. The decoder input is $d_t = \bar{z}_{k(t)} + W_{\text{skip}} h^E_t$ (W_{skip} zero-initialized, W_q, W_k identity-initialized). Decoder wrist cross-attention keys add learned positional embeddings and zero-initialized camera-ID and 2D coordinate projections. All remaining architectural and optimization hyperparameters are listed in Table 3. Ablations on CALVIN and Blocks were trained for 16,000 steps with otherwise default parameters.

Heuristic Router Warm-Start. For $\mathcal{L}_{\text{warm}}$ (w_{warm} decaying linearly from 1 to 0 over S_{warm}), target boundaries are extracted causally from raw proprioception s_t. A step t is a candidate boundary if any gripper state crosses its open/close threshold between $t-1$ and t, or if the arm speed $v_t = \|s^{\text{arm}}_t - s^{\text{arm}}_{t-1}\|$ (computed over end-effector position on CALVIN/LIBERO and arm joint positions on Blocks) reaches a local minimum at $t-1$. A local minimum requires $v_{t-1} < v_{t-2}$, $v_{t-1} \le v_t$, and v_{t-1} to fall below the 25th percentile of training per-step speeds (computed over moving steps only on Blocks). Candidates within 4 steps of a prior heuristic boundary are discarded.

B. Benchmark Protocols

We evaluate on LIBERO [24], LIBERO-Plus [11], CALVIN [26], and our Blocks benchmark. We describe the training data, policy interfaces, evaluation procedures, and baseline implementations below. Unless otherwise specified, variation across seeds refers to independently trained policies. For baselines, which have a separate action tokenizer training step, representation models are shared across policy seeds.

Table 3. **ELMO architecture and training hyperparameters.**

Hyperparameter	Value
Action horizon H	16
Encoder / decoder width D_{enc}	256 (pre-norm residual GRU–MLP, RMSNorm)
Planner width D_{main}	384
Boundary cross-attention	6 heads
Planner causal transformer T	4 blocks, width 384, 6 heads
Decoder wrist cross-attention	4 heads (8×8 grid), RMSNorm
Flow-matching head v_θ	3-layer MLP, width 1024
Target compression ratio C_{ratio}	12 (max timeout: 30 steps)
Ratio loss weight λ_{ratio}	0.03
Router warm-start S_{warm}	5,000 updates ($w_{\text{warm}} : 1 \to 0$)
Onset window K_{init} / EMA decay	16 steps / 0.7 (unseen init. to mean)
Random sequence start prob.	0.5
Spliced fragments per sequence	2–3
Max spliced length / min fragment	192 / 8 (CALVIN); 256 / 64 (LIBERO, Blocks)
Optimizer	AdamW ($\beta = (0.9, 0.95)$, WD 0.1, grad clip 1.0)
Batch size	32
Base learning rate	$2 \times 10^{-4} \to 2 \times 10^{-5}$ (cosine w/ restarts)
Control-rate LR multiplier	$\sqrt{C_{\text{ratio}} D_{\text{main}} / D_{\text{enc}}} \approx 4.24$
Cycle length	8,000 updates
Cycle warm-up	1,000 steps (cycle 1); AdamW state reset + 500 steps
Total training updates	16,000 (LIBERO, CALVIN); 40,000 (Blocks)

B.1. LIBERO

Training data. We train jointly on the four standard suites: LIBERO-Spatial, LIBERO-Object, LIBERO-Goal, and LIBERO-10, also referred to as LIBERO-Long. Each suite contains ten tasks. The baseline dataset is `lerobot/libero` [5], containing 1,693 demonstrations across the four suites. We do not use LIBERO-90 demonstrations for any training run on LIBERO.

Environment initialization and termination. Both cameras render at 256×256 resolution. Images are rotated by 180° to match the training and LeRobot convention. Each episode resets to its specified benchmark initial state, followed by ten settling steps with zero arm motion and an open-gripper command. The policy state and action queue are reset between episodes. An episode succeeds when the simulator's task predicate becomes true within the task's step budget.

Evaluation uses the default simulator control frequency of 20 Hz.

Evaluation sample counts. We evaluate indices 0–24 for each model, totaling $4 \times 10 \times 25 = 1{,}000$ episodes per seed.

We report success percentages for each suite and their

Suite	Maximum policy steps
LIBERO-Spatial	280
LIBERO-Object	280
LIBERO-Goal	300
LIBERO-10 / Long	520

Table 4. Evaluation step budgets for LIBERO.

Perturbation category	Variants
Camera viewpoints	1,599
Robot initial states	1,550
Language instructions	1,537
Lighting	1,142
Background textures	1,076
Sensor noise	1,601
Object layouts	1,525
Total	10,030

Table 5. Composition of the full LIBERO-Plus evaluation.

mean. Uncertainty across policies is computed across training seeds.

B.2. LIBERO-Plus

Zero-shot robustness evaluation. We evaluate the clean-LIBERO policies on LIBERO-Plus [11] without additional training. The full benchmark contains 10,030 perturbed task variants. We execute one trial per variant using initial-state index zero, with 256×256 camera rendering. Each method retains its clean-LIBERO action-execution horizon and uses the corresponding base suite's step budget.

Instruction handling. For non-language perturbations, the upstream task-language field can include metadata derived from the perturbation filename. For example, an instruction may contain camera or initialization identifiers following the actual task description. We remove this metadata and recover the canonical training instruction. For language-perturbation variants, we preserve the intended rewritten instruction.

Aggregation. We report success within each perturbation category and overall success pooled across task variants. The overall score is task-weighted:

$$\mathrm{SR}_{\mathrm{Plus}} = \frac{1}{10{,}030} \sum_{i=1}^{10{,}030} \mathbb{I}[\text{variant } i \text{ succeeds}]. \tag{11}$$

It is not the unweighted mean of the seven category percentages. For multiple training seeds, we report the mean and standard error across seeds.

B.3. CALVIN

Training split. We evaluate on the CALVIN [26] ABCD→D split.

We use `fywang/calvin-task-ABCD-D-lerobot` [5]. The combined dataset contains 24,053 episodes. Training is restricted to episodes 0–22,965, comprising 22,966 episodes and 1,357,600 frames. Episodes 22,966–24,052 belong to the validation split and are excluded from training.

Observations and actions. The dataset provides a 200×200 static-camera image, an 84×84 gripper-camera image, a 15-dimensional robot observation, and a seven-dimensional relative action. The baseline policies use the dataset's 15-dimensional robot state.

ELMO transforms this state to eight dimensions: three TCP position values, three axis-angle orientation values, gripper opening width, and gripper command. Orientation is expressed relative to a fixed reference rotation, avoiding the discontinuity of an absolute axis-angle representation near π. Both baselines and ELMO output CALVIN's seven-dimensional relative action, with a binary gripper command.

Chained-task evaluation. Each evaluation chain contains five subtasks. The environment is initialized once to the chain's specified initial condition. For each subtask, the policy receives its validation instruction and has a budget of 360 control steps. Policy memory and action queues are reset at each subtask transition, while the physical environment continues from the preceding successful subtask. Evaluation stops at the first failed subtask or after all five subtasks succeed.

Let $L_i \in \{0, \ldots, 5\}$ denote the number of consecutively completed subtasks in chain i. We report

$$\mathrm{SR}_k = \frac{1}{N} \sum_{i=1}^{N} \mathbb{I}[L_i \geq k], \qquad k \in \{1, \ldots, 5\}, \tag{12}$$

$$\mathrm{AvgLen} = \frac{1}{N} \sum_{i=1}^{N} L_i = \sum_{k=1}^{5} \mathrm{SR}_k. \tag{13}$$

Chain selection. The evaluator uses the official 1,000-chain list to select the first 500 entries for evaluation.

B.4. Shared SmolVLA Backbone

SmolVLA [36], FAST [31], and EDAR [43] use the LeRobot 0.6.1 SmolVLA implementation. They load the pretrained `SmolVLM2-500M-Video-Instruct` [25] vision-language backbone and initialize a fresh action expert. The vision-language backbone is frozen. The action

expert, state projection, and representation-specific output components are trained.

The backbone uses the first 16 VLM layers. This initialization supplies pretrained vision-language weights rather than a robot-action-pretrained SmolVLA policy. State and action preprocessing uses mean/std normalization unless a representation performs an additional transformation internally. These three baseline families train without image augmentation.

B.5. SmolVLA

SmolVLA directly predicts a 50-step action chunk using its flow-matching action expert. At deployment, it executes the first ten actions before querying the policy again. Inference uses ten flow-integration steps.

The training recipe [36] uses AdamW with betas $(0.9, 0.95)$, $\epsilon = 10^{-8}$, weight decay 10^{-10}, and gradient-norm clipping at ten. A 1,000-update warmup is followed by cosine decay from 10^{-4} to 2.5×10^{-6} over the 100k-update schedule. Training uses bf16 autocast and `torch.compile`.

The principal Blocks baseline runs use 40k-update schedules to match ELMO.

B.6. SmolVLA + FAST

Action representation. FAST [31] replaces the continuous action output with autoregressive prediction of discrete DCT/BPE tokens.

The BPE vocabulary contains 1,024 entries, with one additional EOS token and one padding token, yielding 1,026 token IDs. The tokenizer is fitted separately on each benchmark's training actions, with minimum BPE frequency two. The fitted sequence capacities are 51 slots for LIBERO and 61 slots for CALVIN.

Policy training and inference. The SmolVLA expert predicts tokens with causal attention and cross-entropy supervision. Training uses teacher forcing: each output position receives the preceding target token. Inference uses greedy autoregressive decoding. The frozen codec reconstructs actions through inverse BPE, inverse quantization, and inverse DCT.

On LIBERO and CALVIN, FAST represents ten actions and executes all ten. The final Blocks codec represents 17 actions, approximately one second, and uses 107 token slots. Blocks evaluation executes ten actions per query.

The policy recipe (taken from EDAR [43]) uses a 1k-update warmup and a 100k-update cosine schedule from 1.5×10^{-4} to 5×10^{-6}. Training uses fp32 with compilation.

B.7. EDAR

Two-stage training. EDAR first learns an action/effect representation. The representation model is then frozen, and SmolVLA is trained to predict its latent registers. The final implementation matches the paper [43].

Representation architecture. The encoder receives a 16-step action sequence and frozen DINOv3-Base features from the current camera images. It contains eight encoder layers and eight decoder layers, with width 512, eight attention heads, feed-forward expansion four, and RMSNorm. Four learned registers produce the continuous action representation. Each register has seven dimensions on LIBERO and CALVIN.

The encoder jointly processes action tokens, registers, and current visual features. The decoder uses shared attention with separate action and visual feed-forward paths. Action queries are causal and read the latent and action context. The visual branch additionally reads current image features and predicts future visual features. The action decoding path does not require a live observation.

Representation objective. The representation objective combines action reconstruction with future-effect prediction:

$$\mathcal{L}_{\text{codec}} = \mathcal{L}_{\text{act}} + \lambda_{\text{effect}} \mathcal{L}_{\text{effect}}. \tag{14}$$

Action reconstruction sums squared errors over action coordinates and averages over valid time steps. The effect loss computes cosine distance after flattening each view's patch-feature tensor, then averages across views and samples. We use $\lambda_{\text{effect}} = 0.2$ for LIBERO and 0.1 for CALVIN.

DINOv3-B uses 224-pixel inputs and retains its native 14×14 patch grid. Inputs use ImageNet normalization and bilinear resizing. Flattening each view for the cosine objective is the implementation's explicit interpretation of the effect-loss reduction.

Representation optimization. Codec training follows the paper [43] and uses Adam with betas $(0.9, 0.99)$, batch size 64, 100k updates, and cosine learning-rate decay from 10^{-4} to 10^{-5}. Gradient norm is clipped at one. The final codec uses all training frames without an additional codec holdout.

Policy optimization and deployment. After freezing the codec, the SmolVLA flow expert predicts its four register tokens directly. The register values are not standardized using fitted latent statistics. Policy training [43] uses batch size 64 and a 100k-update schedule with 1k warmup and cosine decay from 1.5×10^{-4} to 5×10^{-6}. Inference retains the stock ten-step flow sampler. The decoded action

Setting	SmolVLA	FAST	EDAR
Policy batch size	64	64	64
Output representation	Raw action chunk	Autoregressive discrete tokens	Four continuous registers
Predicted action horizon, LIBERO/CALVIN	50	10	16
Executed actions per query	10	10	8
Peak policy learning rate	10^{-4}	1.5×10^{-4}	1.5×10^{-4}
Final schedule learning rate	2.5×10^{-6}	5×10^{-6}	5×10^{-6}
Schedule length	100k	100k	100k
Warmup updates	1k	1k	1k

Table 6. Principal SmolVLA-based policy settings on LIBERO and CALVIN. Values were taken from papers wherever possible.

Setting	LIBERO	CALVIN
Action horizon	16	32
Scales	$(1, 2, 4)$	$(1, 2, 3, 4)$
Total code positions	7	10
Codebook size	512	512
Code dimension	32	32

Table 7. SDAT tokenizer configurations.

horizon is 16, of which eight actions are executed before replanning.

On LIBERO and CALVIN, the principal recipe uses 100k codec updates and 100k policy updates. We train one codec for all three policy seeds.

On Blocks, the codec is trained for 20k updates, followed by the principal 40k-update policy runs. The main LIBERO/CALVIN policies use bf16 and compilation. Blocks use fp32 and compilation.

B.8. MINT-30M

Policy architecture. MINT [12] follows its own implementation rather than the shared SmolVLA backbone. Our implementation matches the MINT-30M configuration specified in the paper rather than the released MINT-Light.

Visual features concatenate outputs from frozen DINOv2 ViT-L/14 with register tokens and frozen SigLIP SO400M/14, using 224-pixel inputs. Language is encoded by frozen BERT-base. Its pooled representation conditions the policy through FiLM/AdaLN modulation, without language-token cross-attention.

The trainable policy has eight layers, hidden width 256, twelve attention heads with head dimension 64, and feed-forward width 1,024.

SDAT action tokenizer. MINT represents actions with a multiscale residual vector-quantized tokenizer, SDAT. The policy predicts coarse-to-fine codes, and a frozen SDAT decoder reconstructs the action window.

Tokenizers are retrained on each benchmark's training actions. The final LIBERO baseline does not use the released LIBERO-130 tokenizer. Tokenizer training uses AdamW, batch size 1,024, learning rate 3×10^{-5}, weight decay 0.01, betas $(0.9, 0.95)$, and parameter EMA 0.99. It retains the implementation's reconstruction, vector-quantization, and DCT objectives. LIBERO trained for 25,300 updates and CALVIN for 26,700 tokenizer updates. Both reached similar held-out validation MSE.

Policy optimization. Policy training uses batch size 128 and 30k updates. AdamW uses learning rate 3×10^{-4}, betas $(0.9, 0.95)$, and weight decay 0.01. The schedule uses 1k warmup followed by cosine decay to 9×10^{-5}. Feed-forward dropout is 0.05, and policy parameter EMA is 0.999.

Deployment. Inference uses EMA parameters, top-three code sampling, and intent-weighted temporal ensembling of overlapping action chunks. The ensemble combines intent similarity, prediction confidence, and temporal decay. The policy executes four actions per query.

Image augmentation. The published version uses image augmentations. Since other baselines and ELMO do not use image augmentations, we train a version without to be more directly comparable. The augmented configuration applies up to two image transformations selected from brightness, contrast, saturation, hue, sharpness, and affine transformations. These are reported separately, since augmentation substantially affects LIBERO-Plus robustness.

Bimanual Blocks adaptation. For Blocks, the model receives three cameras and uses a shared per-arm SDAT tokenizer over seven-dimensional action streams. The tokenizer uses horizon 32, scales $(1, 2, 3, 5, 8)$, and a $1{,}024 \times 32$ codebook. Codes for both arms are predicted at each scale. The principal Blocks policies train for 40k updates without image augmentation because color is essential for instruction following.

B.9. Detailed Simulation Results

LIBERO-Plus Perturbation Breakdown. For LIBERO-Plus [11], we evaluate the LIBERO-trained checkpoints zero-shot across all seven perturbation families: object layout, camera viewpoint, robot initial state, language paraphrasing, lighting, background texture, and sensor noise. Table 8 reports the per-perturbation success rates alongside the overall average across all evaluated episodes.

CALVIN Subtask Completion. On the CALVIN (ABCD $\rightarrow$ D) benchmark [26], policies execute 500 chains of five sequential instructions in the held-out environment D. Table 9 reports the success rates for completing at least $k \in \{1, 2, 3, 4, 5\}$ consecutive subtasks alongside the mean chain length.

C. Blocks Benchmark Details

C.1. Motivation, Scene, and Observation Space

Blocks is a language-conditioned tabletop manipulation benchmark built in MuJoCo [38] on the bimanual ALOHA platform [48] (`google-deepmind/aloha_sim`) with two opposing 6-DoF ViperX 300s arms, one overhead RGB camera, and two wrist-mounted RGB cameras. It evaluates object selection, spatial generalization, bimanual manipulation, and instruction following while avoiding the spatial/proximity shortcuts, missing language conditioning, or complex real-world transfer barriers of prior benchmarks (Table 10). Every scene contains the same five persistent 5 cm cubes (0.05 kg each) colored red, green, blue, yellow, and purple, ensuring that object identity cannot be inferred merely from which objects are present in the scene.

Standard simulation parameters are summarized in Table 11. The tabletop workspace spans

$$x \in [-0.32, 0.32]\ \mathrm{m}, \qquad y \in [-0.22, 0.18]\ \mathrm{m} \tag{15}$$

at height $z = 0.74\,\mathrm{m}$. Cubes are initialized with pairwise center-to-center separations of at least 10 cm and initial yaw angles sampled uniformly from $[-\pi/4, +\pi/4]$. The simulator steps physics at 2 ms (0.002 s) and applies control actions every 30 physics substeps, yielding a 60 ms control interval (16.7 Hz). The simulator accepts 14-dimensional commands:

$$q = \left[q^L_{1:6}, g^L, q^R_{1:6}, g^R\right], \tag{16}$$

where $g = 1$ denotes an open gripper and $g = 0$ a closed gripper. Policies receive only RGB images, state, and the current language instruction.

Cameras and visual preprocessing. The overhead camera renders at 1280×720; the simulator crops the rectangle from pixel $(330, 291)$ to $(950, 663)$ and resizes it to 480×288. The left and right wrist cameras each render at 384×216. ELMO feature construction applies an additional overhead-image crop from $(48, 0)$ to $(432, 264)$, followed by bilinear resizing to 512×512.

C.2. Tasks, Sampling, and Generalization Splits

Task families and sampling. Language instructions are instantiated from four templates across three task families:

- `stack the A block on the B block`;
- `move the A block to the SIDE of the table`;
- `move the A and B blocks to the SIDE of the table`;
- `push the A block to the DIRECTION`.

For move tasks, destination regions on the tabletop are defined by

$$\text{left}: \ x < -0.15\ \mathrm{m}, \qquad \text{right}: \ x > 0.15\ \mathrm{m}, \tag{17}$$

$$\text{top}: \ y > 0.06\ \mathrm{m}, \qquad \text{bottom}: \ y < -0.10\ \mathrm{m}, \tag{18}$$

which may overlap at the corners; expert placement targets lie at least 2 cm inside the requested boundary.

To prevent frequent argument rejection from skewing the task-kind distribution, the task sampler first randomizes the task-kind order and then attempts up to 60 feasible argument draws for the selected kind before advancing to another kind. For move tasks, a second block is proposed with probability 0.4 when available. Blocks covered by another block are excluded from the movable set, and a move instruction is rejected if any named block already lies inside the requested region. Push tasks require the target block to rest on the table with a clear, reachable motion corridor.

Color–position generalization split. To evaluate compositional spatial generalization, the workspace is partitioned into a 4×2 grid defined by x-boundaries $(-0.32, -0.16, 0, 0.16, 0.32)\,\mathrm{m}$ and y-boundaries $(-0.22, -0.02, 0.18)\,\mathrm{m}$. Using zero-based grid indices, each cube color is excluded from resting or sliding at tabletop height inside one assigned cell during training: red $(0, 1)$, green $(3, 0)$, blue $(0, 0)$, yellow $(1, 0)$, and purple $(3, 1)$. Carrying a grasped cube in free space above its held-out cell remains permitted. All evaluation seeds are strictly disjoint from training seeds. For position out-of-distribution (position-OOD) evaluation, a designated block is initialized inside its held-out cell and the instruction is forced to name that block.

C.3. Scripted Demonstration Expert and Dataset

Training demonstrations are collected by a scripted inverse-kinematics expert following Cartesian waypoints. To induce trajectory diversity, each task receives a speed multiplier sampled from $[0.85, 1.15]$, grasp positions receive up to ± 2 mm XY jitter, and approach heights receive up to ± 1 cm jitter.

Table 8. **LIBERO-Plus success rate (%) by perturbation type.** Mean ± std error across 3 seeds.

Method	Camera	Robot Init.	Language	Light	Background	Noise	Layout	Total
SmolVLA [36]	53.3 ± 3.2	26.7 ± 1.3	11.9 ± 0.3	81.3 ± 1.6	76.6 ± 0.8	55.9 ± 0.9	61.6 ± 0.9	50.2 ± 0.8
SmolVLA + FAST [31]	54.5 ± 1.9	24.0 ± 0.9	15.9 ± 0.7	78.5 ± 2.4	71.2 ± 1.3	57.7 ± 2.0	62.1 ± 0.6	50.1 ± 0.3
SmolVLA + EDAR [43]	55.3 ± 0.5	28.7 ± 0.2	19.2 ± 0.7	$\mathbf{86.7 \pm 1.1}$	80.4 ± 1.4	62.2 ± 2.0	64.2 ± 0.8	54.4 ± 0.2
MINT-30M [12]	58.7 ± 1.3	$\mathbf{49.9 \pm 1.8}$	20.3 ± 0.7	86.3 ± 1.1	$\mathbf{81.3 \pm 2.2}$	72.9 ± 2.2	$\mathbf{70.1 \pm 0.9}$	61.0 ± 1.3
MINT-30M (w/o image aug.)	35.0 ± 2.8	45.5 ± 1.6	18.8 ± 1.2	75.5 ± 0.9	64.9 ± 2.0	46.7 ± 2.7	61.3 ± 1.0	47.8 ± 1.3
ELMO (Ours)	$\mathbf{88.3 \pm 1.0}$	36.7 ± 1.4	$\mathbf{46.0 \pm 1.4}$	72.5 ± 2.1	47.5 ± 2.8	$\mathbf{88.3 \pm 1.1}$	64.2 ± 0.2	$\mathbf{64.0 \pm 0.8}$

Table 9. **CALVIN (ABCD→D) sequence completion.** Columns 1–5 report the percentage of chains completing at least k consecutive subtasks, and Length is the sum of these five completion probabilities. Mean over three seeds.

Method	1	2	3	4	5	Length ↑
SmolVLA [36]	90.1	74.3	59.6	47.3	35.8	3.07
SmolVLA + FAST [31]	89.9	77.5	65.4	55.1	44.9	3.33
SmolVLA + EDAR [43]	90.5	75.9	65.0	56.1	45.1	3.32
MINT-30M [12]	86.1	68.8	51.6	36.5	25.4	2.68
ELMO (Ours)	**97.0**	**90.5**	**82.0**	**73.3**	**64.8**	**4.08**

Table 10. **Comparison of manipulation benchmarks.** Existing benchmarks either permit spatial/proximity shortcuts [14] and thus lack fully random object locations, lack language conditioning, or involve complex setups that hinder direct sim-to-real transfer.

Benchmark	**Random Locations**	**Language Instruction**	**Easy Setup Transfer To Real**
LIBERO [24]	✗	✓	✗
CALVIN [26]	✗	✓	✗
RoboTwin 2.0 [6]	✗	✓	✗
Colosseum V2 [28]	✗	✓	✗
OCBench [30]	✓	✗	✗
InstructMove [46]	✓	✓	✗
Blocks (Ours)	✓	✓	✓

Table 11. **Standard Blocks simulation settings.**

Component	Setting
Arm-base positions	$x = \pm 0.469$ m
Cube edge length	0.05 m
Cube mass	0.05 kg
Tabletop height	0.74 m
Physics timestep	0.002 s
Physics substeps per action	30
Control interval	0.06 s
Control frequency	Approximately 16.7 Hz
Cube/table sliding friction	0.5
Arm gravity compensation	Enabled
Initial arm pose	Neutral, grippers open

Arm parking and settling. Before an active arm enters the shared central workspace, the inactive arm returns to its neutral parked pose. The active arm decelerates during final pre-contact descent. After each motion segment, the expert holds its target until the measured grasp center settles within 4 mm of the commanded center, for at most 10 additional control steps.

Cross-table relay. When a task requires transporting a cube across single-arm reach limits, the expert executes a bimanual handoff via a central relay zone. The source arm places the block at a candidate relay location ($x \in \{-3, 0, 3\}$ cm combined with candidate y coordinates and small random offsets) and retreats to its parked pose before the receiving arm retrieves the block.

Pushing. Pushes use a half-open gripper and low contact height to prevent tipping. Long pushes are executed via at most four short corrective strokes (each at most 12 cm) aligned with the cube's live position.

Pose-recovery demonstrations. To expose the policy to off-nominal states and the corrections that follow them, we generate an additional set of demonstrations in which the expert acts on a deliberately wrong estimate of its target. For each pick, place, or push approach, with probability 0.7 the expert's estimate of the target pose is offset by 1.5–4.5 cm in a uniformly random horizontal direction, with an additional yaw error of up to ±0.15 rad. The expert plans its approach on this estimate: it moves above the wrong position and descends to 10 cm above the target while still centred on it. The estimate then converges linearly with height to the true pose while the gripper continues to descend slowly, so the arm visibly corrects towards the cube before the normal grasp, place, or push begins. Every frame is labelled with its phase. Actions of the approach towards the wrong estimate are masked from the action loss, so the policy is never trained to imitate the error itself.

Filtering and dataset composition. During data generation, a demonstration episode is retained only if the expert completes the task without disturbing uninstructed cubes,

violating spatial holdouts, or incurring inter-arm collisions. We generated 10,000 clean tasks and 2,500 pose-recovery tasks. After filtering for collisions and invalid block locations, the collection consists of 8,441 clean tasks and 2,097 pose-recovery tasks, in which the pick, place, or push point is randomly perturbed so the expert must execute corrective adjustments before manipulating the block.

C.4. Scoring and Evaluation

Goal predicates and failure conditions. Let p_b denote the 3D center of block b and p_b^0 its position at the start of the task. A task succeeds only when for ten consecutive frames its geometric goal predicate is satisfied and both grippers are open:

- **Stacking (A on B):** The instructed upper cube A must lie within 2.5 cm horizontally of the lower cube B with a vertical center separation of 5.0 ± 1.2 cm:

$$|p_{A,x} - p_{B,x}| < 0.025 \text{ m}, \tag{19}$$
$$|p_{A,y} - p_{B,y}| < 0.025 \text{ m}, \tag{20}$$
$$|p_{A,z} - p_{B,z} - 0.05| < 0.012 \text{ m}. \tag{21}$$

- **Moving:** Every instructed cube must rest within 1.2 cm of tabletop height inside the target region: left ($x < -0.15$ m), right ($x > 0.15$ m), top ($y > 0.06$ m), or bottom ($y < -0.10$ m).
- **Pushing:** Let u be the commanded cardinal unit vector in the tabletop plane, $u_\perp$ its perpendicular unit vector, and $d = p_{A,xy} - p^0_{A,xy}$. The instructed cube A must translate by at least 7 cm along u, deviate by at most 5 cm laterally, and remain within 1.2 cm of tabletop resting height:

$$d^\top u \geq 0.07 \text{ m}, \qquad |d^\top u_\perp| \leq 0.05 \text{ m}. \tag{22}$$

If any uninstructed cube $b \notin \mathcal{B}_{\text{named}}$ (where both blocks in a stacking instruction count as named) moves by more than 3 cm in 3D Euclidean distance from its initial position,

$$\exists b \notin \mathcal{B}_{\text{named}} : \|p_b - p_b^0\|_2 > 0.03 \text{ m}, \tag{23}$$

the rollout terminates immediately as a wrong-object failure. This wrong-block check takes precedence over goal success. Any task exceeding 450 control steps times out.

Evaluation tracks and metrics. Standard evaluation covers 500 in-distribution (ID) tasks sampled from the training initial-position distribution and 250 position-OOD tasks where the named target starts in its held-out color–position cell.

C.5. Task Breakdown and Diagnostic Comparisons

Table 12 disaggregates Blocks performance by task family (Pushing, Moving, and Stacking) on a shared evaluation manifest alongside episode-level failure modes (wrong-object disturbance and timeout rates).

D. Real-World Evaluation Details

Evaluation Setup. We conduct our real-world evaluation on a standard bimanual ALOHA setup, consisting of two parallel-jaw follower arms—each equipped with a wrist-mounted camera—alongside an overhead scene camera and a lower desk camera (unused). The setup also includes two leader arms for collecting training demonstrations via teleoperation. This physical environment does not identically replicate the simulated Blocks environment, differing primarily in joint dynamics (e.g., controller gains), camera extrinsics (e.g., wrist camera placement and viewpoints), and visual appearance (e.g., lighting and tabletop textures). The blocks are 3D-printed from PLA with a side length of 5 cm and painted to approximately match the colors used in simulation.

Training Procedure. For both the SmolVLA baseline and our model (ELMO), we initialize from the best simulation checkpoints and fine-tune on real-world teleoperated demonstrations for 10k steps, reinitializing the optimizer and learning rate schedule. We select 10k steps based on the convergence of the training and validation curves. Rather than co-training on combined simulation and real-world data, we fine-tune exclusively on real-world data, as this yields stronger performance for both models—likely due to the substantial sim-to-real gap between the simulated and physical setups.

We collect 30 demonstrations for each combination of task family (push and move), block color, and direction, totaling 1,200 demonstrations across 40 distinct task combinations. Before each demonstration, the blocks are reset to randomized positions. Additionally, we crop the wrist and scene camera feeds to better align their fields of view with the simulation environment.

Inference Procedure. Both models follow the same inference procedure used in simulation. Because deploying the policies directly on the physical robot causes excessive jitter, we smooth action execution using temporal ensembling from ACT [47]. Specifically, at each control step, we compute the executed action as an exponentially weighted average over all overlapping action predictions for the current timestep, weighting the i-th prediction by $w_i = \exp(-m \cdot i)$ with $m = 0.05$. While querying the policy at every timestep ($k = 1$) maximizes smoothing, the higher computational overhead of SmolVLA prevents real-time execution at that rate. Thus, we query SmolVLA every $k = 10$ steps, whereas ELMO is queried at every step ($k = 1$).

Evaluation Procedure. Before each trial, we place the blocks at randomized positions on the table, and an auto-

Table 12. **Blocks task-family and failure-mode breakdown.** Results on the in-distribution (ID) and out-of-distribution (OOD) tasks, mean over three seeds. Push, Move and Stack report task success rates (%). Wrong object and Timeout report the percentages of all evaluation rollouts terminating with each failure condition.

	ID					OOD				
Method	Push ↑	Move ↑	Stack ↑	Wrong obj. ↓	Timeout ↓	Push ↑	Move ↑	Stack ↑	Wrong obj. ↓	Timeout ↓
SmolVLA [36]	12.9	0.2	0.0	23.7	71.4	2.2	0.0	0.0	24.1	75.2
SmolVLA + FAST [31]	8.1	0.4	0.2	33.5	63.3	1.8	0.0	0.0	37.2	62.3
SmolVLA + EDAR [43]	12.5	0.2	0.0	24.7	70.5	5.7	0.7	0.0	20.9	77.1
MINT-30M [12]	5.4	0.0	0.0	15.4	82.6	3.5	0.0	0.0	14.8	84.1
ELMO (Ours)	**83.0**	**17.2**	**7.2**	**10.5**	**50.9**	**74.6**	**15.2**	**10.4**	**13.1**	**54.8**

mated script samples the target color and direction. For the push task, a trial is marked successful if the robot executes a clear pushing motion (rather than grasping) that moves the target block along an approximately straight path in the specified direction. For the move task, success requires picking up the target block and placing it in the location corresponding to the specified direction. Each policy is allotted $20\,\mathrm{s}$ per trial; timeouts and excessive contact with non-target blocks are counted as failures.

E. Additional Latent Space Analysis Details

To compare what each method's latent encodes, we encode the same demonstration chunks with every model and fit linear probes and t-SNE embeddings on the result. For each benchmark we sample 6,000 chunk start points from the training demonstrations, on a stride of 8 steps and only where a full action horizon remains, using a fixed random seed. Each chunk is labelled with the instruction, the task group, the target object, whether the gripper opens or closes within the next 16 steps, the episode progress, the end-effector position at the start and 16 steps later, the end-effector displacement over those 16 steps, the commanded gripper state, and (CALVIN, Blocks) the target block position and the positions of all blocks. Every model encodes the chunk with its own frozen representation, using the inputs it receives during policy training:

- **MINT (all scales):** MINT's policy predicts the tokens of a frozen multi-scale VQ-VAE over action chunks. We take this VQ-VAE from the final MINT policy and tokenise each chunk's ground-truth actions over its native horizon. The latent is the quantised representation $\hat{f}$ the VQ-VAE decoder reads (the sum of the upsampled codebook embeddings over all scales that the MINT policy learns to predict). For bimanual tasks the two arms are tokenised as separate streams after per-dimension quantile normalisation, and their latents are concatenated.
- **MINT (scale 1):** the codebook embedding of the single coarsest-scale token of the same tokenisation (per arm), which is where the MINT policy places its intention.
- **EDAR:** EDAR's policy is trained to predict the latent code of a frozen action autoencoder, which encodes a chunk of 16 actions together with the camera images at the chunk's first step into a small set of register tokens. We take this autoencoder from the final EDAR policy checkpoint and encode each demonstration chunk with it, giving it the ground-truth actions (normalised with the autoencoder's own statistics) and the camera images at the chunk's first step. The latent is the encoder's register tokens, flattened into a single vector.
- **ELMO:** the latent plan vector the decoder receives at the chunk's first step, taken from a teacher-forced forward pass over the whole demonstration.

Probes. All latents are standardised using statistics from the training folds. We fit logistic regression (L2, $C{=}1$) for categorical labels and ridge regression ($\alpha{=}10$) for continuous ones, with 5-fold cross-validation grouped by episode so that no episode appears in both training and test folds. Categorical probes are scored as accuracy above the majority-class rate, $(\mathrm{acc} - p_{\mathrm{maj}})/(1 - p_{\mathrm{maj}})$, and continuous probes as R^2 averaged over output dimensions; negative scores are shown as 0 in the radar charts. As a reference we probe the raw 16-step action chunk.

t-SNE. Each latent is standardised and, if its dimension d exceeds 50, projected onto its first 50 principal components (PCA to $\min(50, d)$ dimensions). The result is embedded in 2D with t-SNE (perplexity 30, PCA initialisation). Linear probes are fit on the full standardised latents, without PCA.

F. Additional Figures

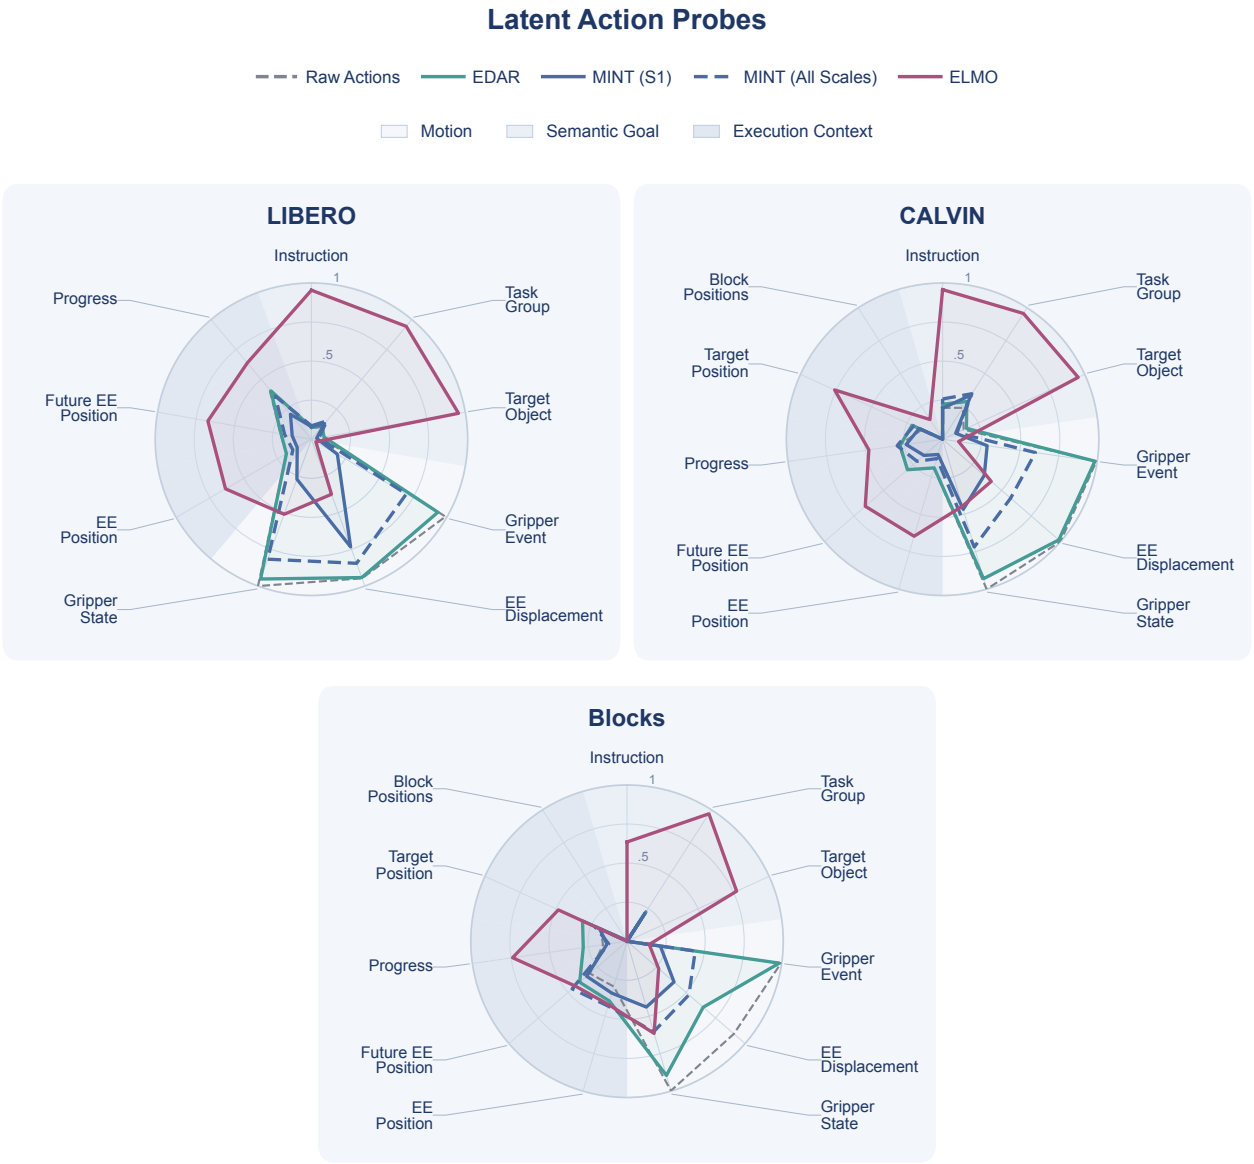

Figure 7. Extended linear probe performance on all benchmarks.

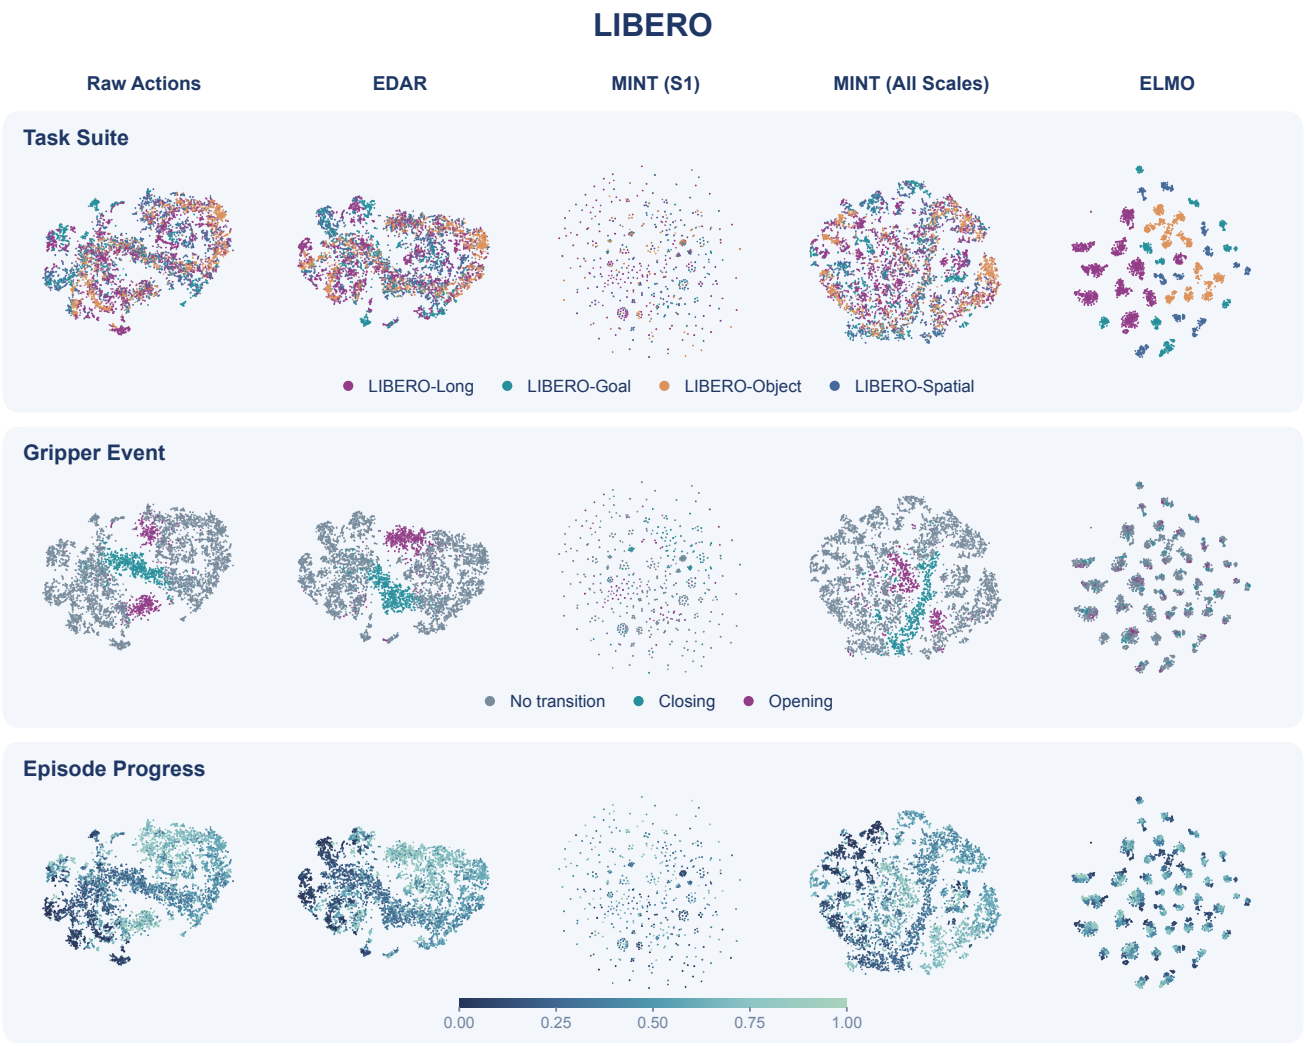

Figure 8. Additional t-SNE plots on the LIBERO benchmark.

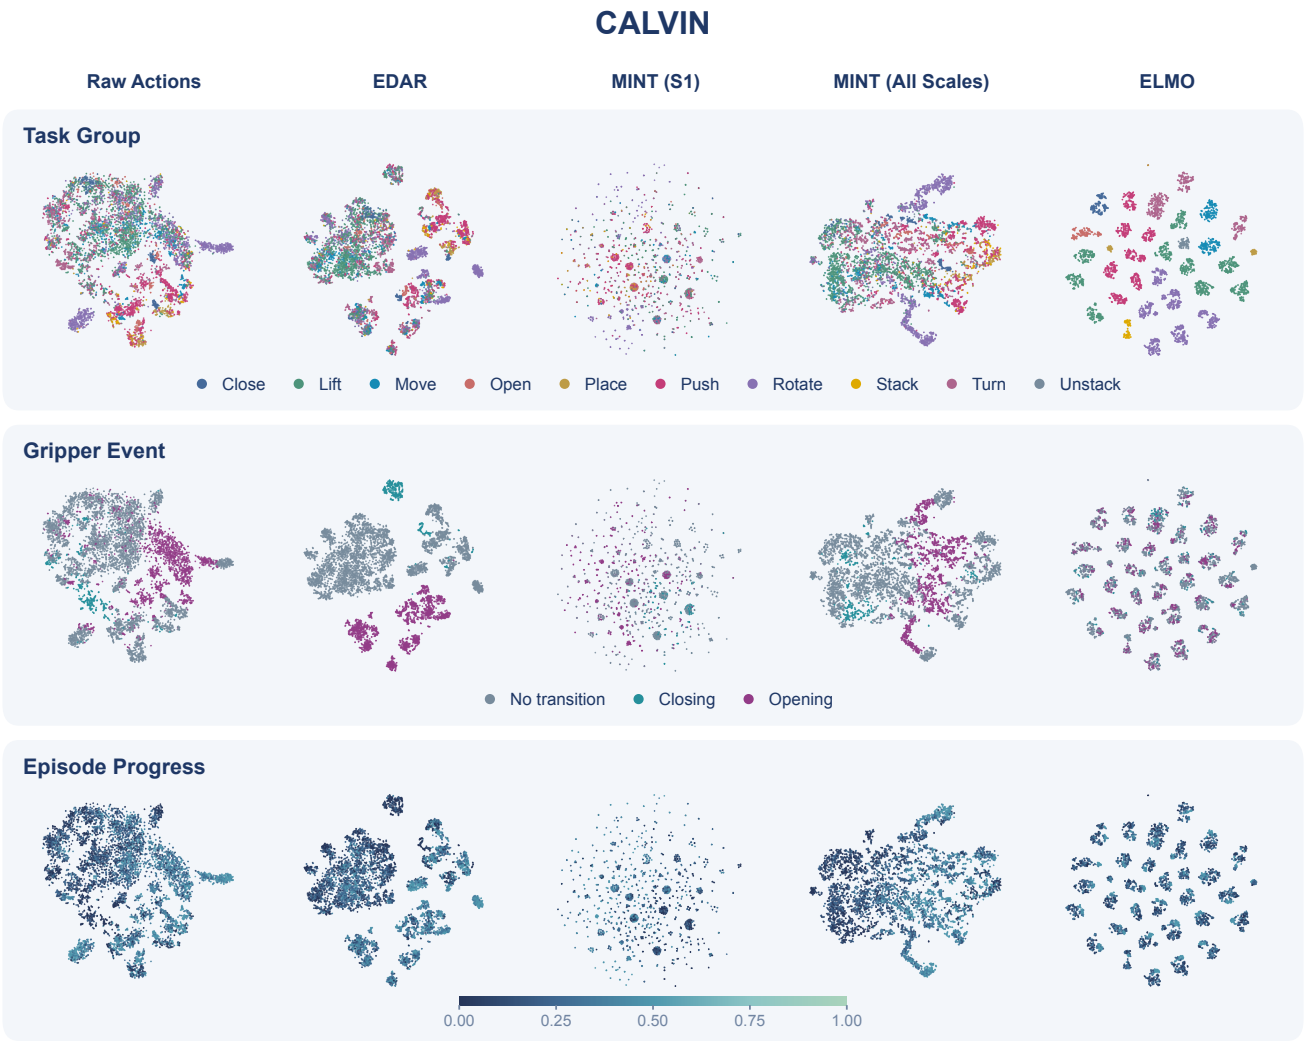

Figure 9. Additional t-SNE plots on the CALVIN benchmark.

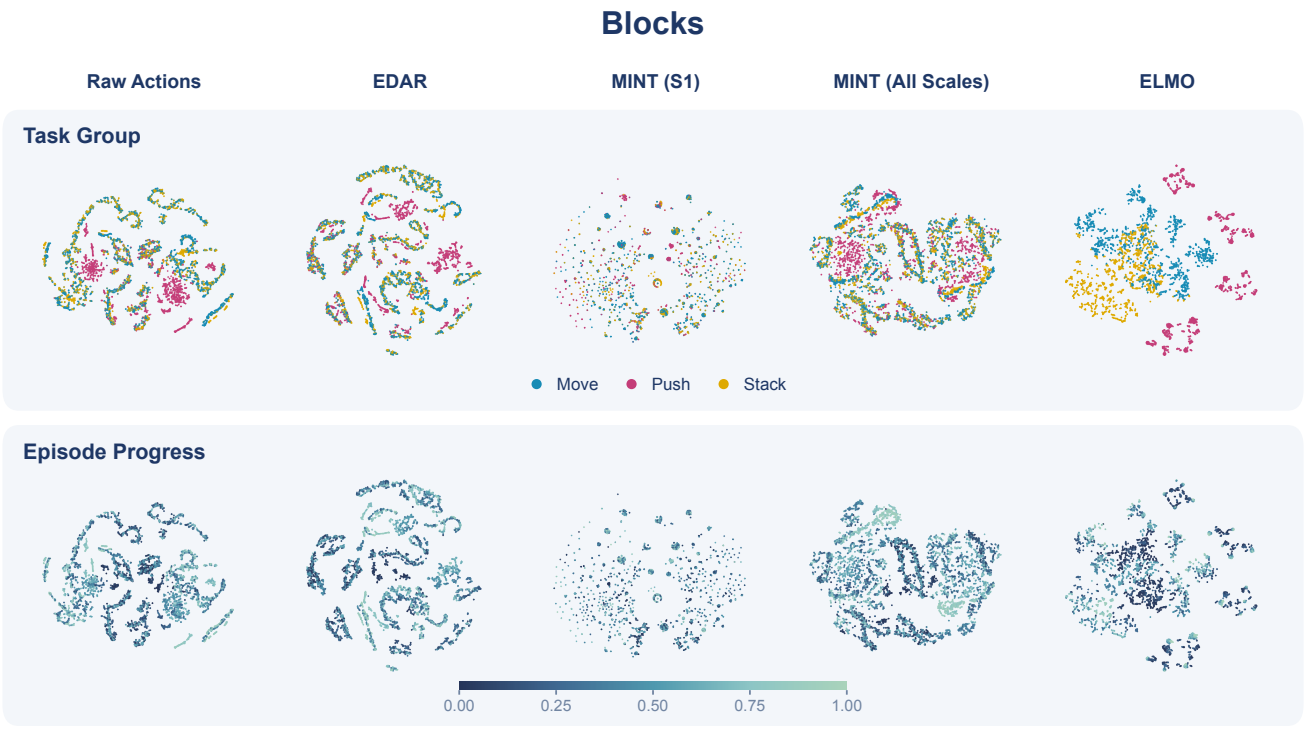

Figure 10. Additional t-SNE plots on the Blocks benchmark.