# DEEPTHINKVLA: ENHANCING REASONING CAPABILITY OF VISION-LANGUAGE-ACTION MODELS

Anonymous authors
Paper under double-blind review

# **ABSTRACT**

Enabling Vision-Language-Action (VLA) models to "think before acting" via Chain-of-Thought (CoT) is a promising path to overcoming the data-hungry nature of end-to-end robot policies. However, progress is stalled by a fundamental conflict: existing models use a single autoregressive decoder for both sequential CoT reasoning and high-dimensional, parallelizable robot actions. This architectural mismatch degrades motor control and fails to forge a strong causal link between thought and action. We introduce DeepThinkVLA, which resolves this conflict through a tightly integrated architecture and training strategy. Architecturally, our hybrid-attention decoder generates sequential CoT with causal attention and then switches to bidirectional attention for fast, parallel decoding of action vectors. This design is complemented by a two-stage training pipeline: we first use Supervised Fine-Tuning (SFT) to teach the model foundational reasoning, then apply Reinforcement Learning (RL) with task-success rewards to causally align the full reasoning-action sequence with desired outcomes. This synergy leads to state-of-the-art performance, achieving a 97.0% success rate on the LIBERO benchmark. Our ablations confirm the design's effectiveness: the hybrid architecture alone outperforms standard decoders by 15.5%, and the final RL stage provides a crucial 2% boost to secure top performance. <sup>1</sup>

# 1 Introduction

Vision–Language–Action (VLA) models have driven notable progress in robotic manipulation, enabling tasks like stacking blocks, opening drawers, and arranging household objects (Huang et al., 2023; Zitkovich et al., 2023; Yang et al., 2024). The dominant paradigm learns a reactive, end-to-end policy that directly maps high-level goals and sensory inputs to low-level motor commands (Chi et al., 2023; Kim et al., 2024; Bjorck et al., 2025). However, this approach faces a critical bottle-neck: learning this monolithic, perception-to-action mapping is inherently difficult and demands vast amounts of high-quality demonstration data (Ma et al., 2024; Liu et al., 2025b).

To address the challenges of direct perception-to-action mapping, one promising direction is to endow VLAs with the ability to "think before acting". Recent studies take initial steps in this direction by supervised fine-tuning (SFT) VLAs on embodied Chain-of-Thought (CoT) datasets (Zawalski et al., 2024; Chen et al., 2025). This approach forces the policy to generate CoT reasoning before actions, aiming to simplify the direct mapping problem.

However, a critical limitation of these methods lies in their underlying architecture. First, these studies typically employ a unified, purely autoregressive decoder to generate the entire reasoning and action sequence. This design creates a fundamental conflict: CoT, as natural language, is inherently sequential and well-suited for autoregressive modeling, whereas actions are often high-dimensional vectors where dimensions (e.g., end-effector translation and rotation) can be determined in parallel and are highly sensitive to latency (Liu et al., 2025d; Kim et al., 2025; Song et al., 2025). Forcing a single model to master both modalities can compromise motor control precision for the sake of language fluency. Second, relying solely on SFT often leads to the model merely "rote learning" the CoT annotations without establishing a strong causal link to subsequent actions. This results in limited improvement in task success, as the reasoning is not effectively utilized to guide behavior.

<sup>&</sup>lt;sup>1</sup>Code available at https://anonymous.4open.science/r/DeepThinkVLA

To the end, we propose DeepThinkVLA, an approach that synergistically co-designs the model architecture and the training strategy. From the architectural perspective, we design a hybrid-attention decoder. This decoder employs causal attention for generating the sequential COT and switches to bidirectional attention for parallel decoding of the high-dimensional action vector. This design not only respects the intrinsic properties of language and motor commands but also critically reduces the inference latency associated with action generation. From the training perspective, we introduce a Reinforcement Learning (RL) training pipeline: first, we perform SFT on an embodied CoT dataset to instill the basic "thinking" ability in the model. Subsequently, leveraging the high-speed rollouts enabled by our parallel action decoder, we apply outcome-based RL. This RL stage directly aligns the entire reasoning-action sequence with task success, optimizing the CoT to be not just plausible, but genuinely beneficial for solving the task.

Based on the public weights of  $\pi_0$ -FAST (Pertsch et al., 2025), we build DeepThinkVLA by integrating a hybrid-attention architecture, an SFT cold-start on an embodied CoT dataset from our two-stage pipeline, and outcome-based RL. Our experiments provide clear evidence for the effectiveness of the proposed approach. After SFT, applying RL to align CoT reasoning with the generated actions improves task success rates by an additional 2% on the LIBERO-Long suite. The hybrid-attention architecture, designed to support the "think before acting" paradigm, proves critical: training without this modification on the embodied CoT dataset results in an average success rate 15.5% lower than with the hybrid design. Overall, DeepThinkVLA achieves an average success rate of 97.0% on the LIBERO benchmark under the single-model, multi-suite evaluation protocol.

# 2 RELATED WORK

Vision–Language–Action Models. Recent progress in VLAs has focused on architectural innovations built on the reactive perception-to-action paradigm. Early work such as RT-2 (Zitkovich et al., 2023) popularized this direction, leading to a wide range of models that adopt different VLM backbones (Black et al., 2025; Bjorck et al., 2025; Hung et al., 2025; Pertsch et al., 2025) and are trained on large-scale robotic datasets (Walke et al., 2023; Fang et al., 2024; O'Neill et al., 2024; Khazatsky et al., 2024; Wu et al., 2025). Variants include diffusion-based decoders that model complex action distributions through iterative refinement (Black et al., 2025; Liu et al., 2025a), block-parallel decoders that improve distributional modeling by applying bidirectional attention to predict multiple action tokens (Liu et al., 2025d; Kim et al., 2025; Song et al., 2025), and hierarchical structures that separate planning from execution (Belkhale et al., 2024; Cui et al., 2025; Team et al., 2025; Bu et al., 2025a). Yet across these designs, the action-generating policy remains reactive—mapping observations and goals directly to actions (or observations and subgoals in hierarchical variants). Consequently, the overall philosophy is still direct mapping, which becomes inadequate under a "think before acting" paradigm that demands both CoT-friendly text generation and high-performance action decoding.

Embodied Reasoning via Supervised Fine-Tuning. To move beyond purely reactive mapping, several recent studies have attempted to endow VLAs with embodied reasoning ability through SFT on CoT augmented data (Zawalski et al., 2024; Lin et al., 2025; Tan et al., 2025; Chen et al., 2025). This line of work typically constructs CoT annotations for existing embodied datasets (O'Neill et al., 2024), often by leveraging stronger cloud-based models (Team, 2023), and then fine-tunes open-source VLAs such as OpenVLA (Kim et al., 2024). While such methods provide a first step toward "think before acting," they face two persistent challenges. First, the availability of high-quality, CoT-annotated embodied data remains limited (Wang et al., 2024; Xu et al., 2024; Zhong et al., 2025). Second, SFT tends to treat reasoning and action generation as parallel objectives within a fixed dataset (Zhao et al., 2025a; Liu et al., 2025c), which can lead to shallow memorization of reasoning traces without strong alignment to the resulting actions. As a result, performance improvements from SFT alone are often modest, indicating the need for approaches that can more directly couple reasoning with successful task execution.

# 3 DEEPTHINKVLA

DeepThinkVLA is built on the synergistic co-design of a novel model architecture and an RL training pipeline. We first formalize the "think before acting" paradigm as a probabilistic decomposition

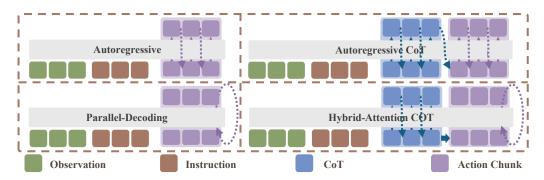


Figure 1: Comparison of VLA architectures. Existing designs adopt either fully autoregressive decoding or parallel bidirectional decoding. DeepThinkVLA introduces a hybrid-attention architecture, enabling autoregressive CoT reasoning alongside efficient parallel action generation.

(Section 3.1). We then introduce our core architectural innovation, the hybrid-attention decoder, which is specifically designed to implement this decomposition efficiently (Section 3.2). Finally, we detail our RL training strategy, where this architecture first enables foundational reasoning via SFT and then makes large-scale RL practical for aligning this reasoning with task success (Section 3.3).

#### 3.1 PROBLEM FORMULATION

Standard VLA policies learn a direct mapping from visual observations (V) and language instructions (L) to a sequence of actions (A). Instead, our work adopts the principle of "think before acting". We implement this principle by introducing a latent reasoning variable, the Chain-of-Thought (R). This approach decomposes the problem. Rather than modeling the direct policy, we model the joint probability of reasoning and then acting:

$$P(A, R|V, L) = P(A|V, L, R) P(R|V, L).$$
 (1)

The advantage of this decomposition is twofold. First, learning P(R|V,L) is highly efficient. Most VLAs are built upon large VLM backbones, which already contain rich semantic and reasoning knowledge. Fine-tuning such models on a relatively small set of synthetically generated embodied CoT data is often sufficient to adapt their reasoning capability to the robotics domain. Second, learning P(A|V,L,R) becomes significantly simpler compared to directly modeling P(A|V,L). The CoT R serves as an explicit, step-by-step plan that disambiguates high-level instructions L, turning the original ill-posed, one-to-many mapping into a constrained and well-specified mapping from a reasoning step to its corresponding motor action. This principled factorization also enables emergent self-correction behaviors, as illustrated in Figure 5.

#### 3.2 Hybrid-Attention Architecture

To implement the factorized policy from Eq. 1, we introduce a novel hybrid-attention architecture. Its primary motivation is to resolve the fundamental conflict between the modalities of reasoning and action within a unified autoregressive decoder.

We propose a novel hybrid-attention architecture that aligns the decoding mechanism with the intrinsic properties of each modality. The core of our design is a dynamic attention mode within a single decoder: For CoT Generation (P(R|V,L)): The decoder employs standard autoregressive causal attention. This respects the sequential nature of language, where each reasoning token is generated based on its predecessors. For Action Generation (P(A|V,L,R)): After generating the CoT, the attention mechanism switches to bidirectional (non-causal) attention. This allows the model to process the entire action specification jointly and decode the action vector in parallel, acknowledging that different dimensions of a motor command (e.g., translation and rotation) are often determined concurrently.

Beyond resolving this core modality mismatch, the parallel decoding of actions yields a critical practical advantage: a significant reduction in inference latency. This speedup is the key enabler

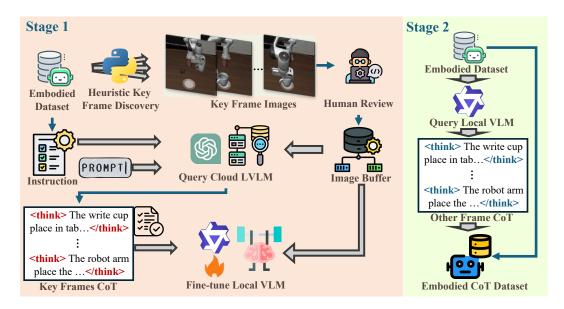


Figure 2: Pipeline for constructing an embodied CoT dataset. Stage 1 extracts keyframes via gripper state signals and queries a cloud LVLM to generate CoT for those frames. Stage 2 fine-tunes a local vision–language model on the keyframe CoT and uses it to annotate the remaining frames.

for our subsequent training stage. While standard autoregressive models are prohibitively slow for the massive number of rollouts required by on-policy Reinforcement Learning, our architecture's high-throughput action generation makes large-scale online fine-tuning computationally tractable. Further architectural details are summarized in Figure 1.

#### 3.3 RL Training Pipeline

With our efficient architecture in place, we train DeepThinkVLA using a two-stage pipeline designed to first instill foundational reasoning and then align it with optimal task outcomes.

**SFT Cold-Start for Foundational Reasoning** SFT cold-start is designed to equip the model with a foundational CoT reasoning capability. This phase requires a specific supervision format comprising complete (V, L, R, A) sequences. However, most existing large-scale embodied datasets lack explicit CoT annotations and instead provide only (V, L, A) tuples. To address this critical data gap, we developed a scalable, two-stage data augmentation pipeline that generates high-quality CoT annotations, as illustrated in Figure 2.

Our pipeline is optimized for both annotation quality and cost-efficiency. In stage 1, we identify semantically significant keyframes within each trajectory by detecting changes in the gripper state, which often indicate subtask boundaries. For these keyframes, CoT annotations are obtained by querying a powerful, general-purpose cloud-based Vision-Language Model (VLM) (see prompt in Appendix Fig. 6). In stage 2, to efficiently annotate the numerous intermediate frames, we fine-tune a smaller, locally-deployed VLM on the high-quality keyframe annotations obtained in Stage 1. This specialized model then automatically generates CoT annotations for the transitional frames. To ensure data fidelity, we apply schema checks to filter malformed outputs and enforce temporal consistency, resulting in a uniform embodied CoT dataset suitable for SFT.

**Learning Reasoning and Action via RL** While SFT provides a strong imitation learning foundation, it cannot adapt to novel scenarios or optimize beyond the sub-optimal trajectories present in the static dataset. To overcome these limitations and truly learn a policy that maximizes task success, we introduce a second training stage based on online Reinforcement Learning. We use an outcome-based reward to jointly optimize the entire reasoning-action sequence towards the singular goal of task completion (Figure 3).

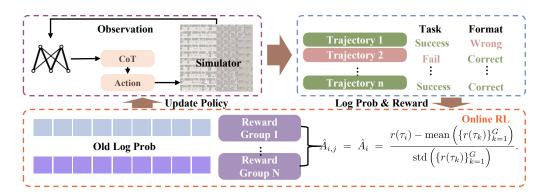


Figure 3: Reinforcement learning stage with grouped credit assignment. The model generates CoT and action sequences that are executed in the simulator to produce trajectories with verifiable rewards. Rewards are grouped and standardized to compute token-level advantages, which update the policy via a clipped surrogate objective with KL regularization to the SFT reference.

We cast the online RL stage as a policy optimization problem: starting from the initial SFT policy, we aim to maximize the expected outcome-based reward through interactive learning. We adopt an on-policy policy gradient algorithm with a clipped surrogate objective akin to PPO. The VLA collects trajectories from the environment, computes advantages, and updates its policy accordingly.

Formally, during the RL rollout, the state at each step t is defined as  $s_t = [o_t^{\mathrm{vis}}, \ell_{\mathrm{task}}]$ , where  $o_t^{\mathrm{vis}}$  denotes the visual observation and  $\ell_{\mathrm{task}}$  is the task instruction. Given the state input, the VLA outputs  $\mathcal{A}_t = [a_t^{\mathrm{cot}}, a_t^{\mathrm{robot}}]$ , where the reasoning tokens  $a_t^{\mathrm{cot}}$  are generated autoregressively, and the action tokens  $a_t^{\mathrm{robot}} \in \mathbb{R}^{h \times d}$  are decoded in parallel. Here, h denotes the action chunk size and d corresponds to the robot's control dimension (e.g., d=7 for a 6-DoF manipulator plus gripper control).

Let  $\pi_{\theta}$  denote the current policy and a trajectory  $\tau$  sampled from the old policy  $\pi_{\theta_{\text{old}}}$  is defined as  $\tau = [(s_0, \mathcal{A}_0), (s_1, \mathcal{A}_1), \dots, (s_T, \mathcal{A}_T)]$ . The reward function  $\mathcal{R}(\tau)$  is sparse, awarded only at the end of the trajectory based on a verifiable task completion signal  $\mathcal{I}_{\text{success}}$ . No intermediate reward is given for the semantics of the reasoning trace. A small format-regularization reward  $\mathcal{I}_{\text{format}}$  is added to prevent stylistic drift. Hence, the reward function is defined as

$$\mathcal{R}(\tau) = \alpha_s \cdot \mathcal{I}_{\text{success}} + \alpha_f \cdot \mathcal{I}_{\text{format}},$$

$$\mathcal{I}_{\text{success}} = \begin{cases} 1, & \text{if task success,} \\ 0, & \text{otherwise,} \end{cases} \qquad \mathcal{I}_{\text{format}} = \begin{cases} 1, & \text{if CoT format correct,} \\ 0, & \text{otherwise.} \end{cases}$$

$$(2)$$

where  $\alpha_s$  and  $\alpha_f$  are weighting coefficients. Then the token-level clipped surrogate objective is:

$$\mathcal{J}(\theta) = \mathbb{E}_{\tau \sim \pi_{\theta_{\text{old}}}} \left[ \sum_{j=1}^{N} \min \left( \omega_j(\theta) \, \hat{A}_j, \, \operatorname{clip} \left( \omega_j(\theta), 1 - \epsilon, 1 + \epsilon \right) \, \hat{A}_j \right) \right], \tag{3}$$

where  $N = |\mathcal{A}_t| \times T$  denotes the total number of tokens in a trajectory  $\tau$ .  $\omega_j(\theta) = \frac{\pi_{\theta}(a_j|s_t, a_{< j})}{\pi_{\theta_{\text{old}}}(a_j|s_t, a_{< j})}$  and  $\hat{A}_j$  are the importance ratio and the advantage for token  $a_j \in [\mathcal{A}_0, \dots, \mathcal{A}_T]$  within trajectory  $\tau$ , respectively.

To propagate the sparse, outcome-based reward  $\mathcal{R}(\tau)$  to each token prediction, we adopt the credit assignment strategy from GRPO (Shao et al., 2024). A group of G trajectories is collected for each task prompt and their rewards are standardized to compute a shared advantage value for all tokens within a given trajectory. For token j in trajectory i, the advantage is:

$$\hat{A}_{i,j} = \frac{\mathcal{R}(\tau_i) - \operatorname{mean}\left(\{\mathcal{R}(\tau_k)\}_{k=1}^G\right)}{\operatorname{std}\left(\{\mathcal{R}(\tau_k)\}_{k=1}^G\right)}.$$
(4)

This relative credit assignment encourages the model to prefer reasoning and action sequences that lead to better-than-average outcomes, effectively selecting for more functional thought processes.

Combining the clipped surrogate objective with the GRPO-style advantage and a KL-divergence penalty to the original SFT policy  $\pi_{ref}$ , which prevents catastrophic forgetting, our final objective is:

$$\mathcal{J}_{\text{final}}(\theta) = \mathbb{E}_{s \sim \text{env}, \{\tau_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}} \left[ \frac{1}{G} \sum_{i=1}^G \frac{1}{N} \sum_{j=1}^N \min \left( \omega_{i,j}(\theta) \, \hat{A}_{i,j}, \right. \right. \\
\left. \text{clip} \left( \omega_{i,j}(\theta), 1 - \epsilon, 1 + \epsilon \right) \, \hat{A}_{i,j} \right) - \beta \, \text{KL} \left( \pi_{\theta}(\cdot \mid s) \, \| \, \pi_{\text{ref}}(\cdot \mid s) \right) \right], \tag{5}$$

where  $\omega_{i,j}(\theta)$  is the the importance ratio for token  $a_j \in [A_0, \dots, A_T]$  in trajectory i By maximizing  $\mathcal{J}_{\text{final}}$ , the VLA simultaneously refines its reasoning and action abilities, with both aligned toward the singular goal of maximizing the final task success rate.

### 4 EXPERIMENTS

# 4.1 EXPERIMENTAL SETUP

Implementation Details. DeepThinkVLA is initialized from the public  $\pi_0$ -FAST weights (Pertsch et al., 2025). We refactor the baseline policy with our hybrid-attention decoder (Sec. 3.2), yielding a 2.9B parameter model. Training proceeds in two stages. First, an SFT cold-start is performed on our embodied CoT dataset (Sec. 3.3). A hybrid mask supervises CoT tokens with causal attention and action tokens with bidirectional attention in a single forward pass, optimized by token-level cross-entropy loss. Second, we apply online RL (Sec. 3.3) to align generated CoT with action execution, using outcome-based rewards and GRPO-style grouped credit assignment. Additional training details, including dataset statistics, hyperparameters, and inference settings, are provided in Appendix A.2.

**Benchmarks.** We evaluate DeepThinkVLA on the LIBERO benchmark (Liu et al., 2023), which provides four suites (Object, Spatial, Goal, and Long) of language-conditioned robotic manipulation tasks. Each suite contains 10 tasks. Evaluation is conducted under 50 randomized initial conditions for each task and performance is measured by the average Success Rate (SR) under the standard single-model, multi-suite evaluation protocol.

**Baselines.** We evaluate DeepThinkVLA against a wide spectrum of recent VLA systems, covering autoregressive SFT models, diffusion-based approaches, parallel-decoding methods, RL-augmented policies and commercial baselines such as  $\pi_0$  and  $\pi_0$ -FAST. A detailed description of all baselines is provided in Appendix A.3.

#### 4.2 Main Results

The main results are shown in Table 1. Our proposed model, DeepThinkVLA, establishes a new state-of-the-art by achieving the highest average success rate (SR) of 97.0%. This result validates the superior overall performance of our Hybrid-Attention mechanism compared to a suite of leading methods across autoregressive, diffusion, and parallel-decoding architectures. DeepThinkVLA shows exceptional proficiency across all categories, securing the top performance in Object manipulation (99.0%), Goal completion (96.4%), and long-horizon tasks (Long, 74.6%). When compared against top-performing autoregressive models like UniVLA (96.2% average), our model demonstrates greater robustness by outperforming it in the Object, Spatial, and Goal categories. Similarly, against the leading diffusion model,  $\pi_0$  (94.2% average), which excels in Spatial tasks (96.8%), DeepThinkVLA shows a clear advantage in the Object, Goal, and overall average metrics. The substantial performance gap over parallel-decoding methods further underscores the architectural benefits of our approach. In conclusion, by achieving the highest average success rate through a balanced and consistently superior performance across a wide spectrum of tasks, DeepThinkVLA not only sets a new benchmark standard on the LIBERO suite but also validates our Hybrid-Attention design as a promising direction for developing more capable and generalist robotic agents.

Table 1: Main results on the LIBERO Simulation Benchmark. All reported values denote the task success rate (SR, %) evaluated under 50 randomized initial conditions per task, averaged within each suite and across all suites. **Wrist** indicates the use of an additional wrist-mounted camera.

Category	Method	Object	Spatial	Goal	Long	Average
	TraceVLA (Zheng et al., 2024)	85.2	84.6	75.1	54.1	74.8
	Octo (Team et al., 2024)	85.7	78.9	84.6	51.1	75.1
	OpenVLA (Kim et al., 2024)	88.4	84.7	79.2	53.7	76.5
Autoregressive	SpatialVLA (Qu et al., 2025)	89.9	88.2	78.6	55.5	78.1
J	GRAPE (Zhang et al., 2024)	91.2	87.6	82.2	55.8	79.2
	NORA (Hung et al., 2025)	95.4	92.2	89.4	74.6	87.9
	VLA-RL (Lu et al., 2025)	91.8	90.2	82.2	59.8	81.0
	$\pi_0$ -FAST (Pertsch et al., 2025)	96.8	96.4	88.6	60.2	85.5
	UniVLA (Bu et al., 2025b)	96.8	96.5	95.6	92.0	95.2
Difference on	Diffusion Policy (Chi et al., 2023)	92.5	78.3	68.3	50.5	72.4
Diffusion	$\pi_0$ (Black et al., 2025)	98.8	96.8	95.8	85.2	94.2
D	CoT-VLA-7B (Zhao et al., 2025b)	91.6	87.5	87.6	69.0	81.1
Parallel-Decoding	OpenVLA-OFT (Kim et al., 2025)	92.7	91.3	90.5	86.5	90.3
Hybrid-Attention   DeepThinkVLA (Ours)		99.0	96.6	96.4	96.2	97.0

#### 4.3 ANALYSIS EXPERIMENT

Ablation on the Role of CoT. To clarify the role of embodied CoT in our framework, we conduct experiments to examine two potential contributions: (i) whether CoT primarily facilitates representation learning during training, and (ii) whether it also serves as an explicit reasoning guidance mechanism at inference time. To this end, we evaluate two intervention settings. In the Mask CoT condition, all CoT tokens at inference are replaced with a fixed placeholder sequence, <think></think>, effectively removing their semantic content while preserving structural alignment. In the Random CoT condition, the CoT tokens are substituted with a randomly generated, fixed-length sequence of semantically meaningless tokens, thereby disrupting both their informational and structural coherence.

As reported in Table 2, the Mask CoT variant shows almost no performance loss compared to the full model (Average: 96.5% vs. 96.8%), suggesting that training with embodied CoT enriches internal representations even when reasoning is not explicitly used at test time. In contrast, the Random CoT setting leads to a sharp decline in success (Average: 85.1%; Long-horizon: 87.8% vs. 94.2%), indicating that once reasoning traces are utilized during inference, their semantic coherence is essential for successful execution. These results demonstrate that embodied CoT benefits the policy in two complementary ways: it strengthens representation learning during training and, when preserved, provides actionable guidance during inference.

Table 2: **Ablation on CoT and hybrid attention.** All values except the last column denote success rate (%). The last column reports relative inference time. Methods include the baseline  $\pi_0$ -FAST, a naive AR-CoT variant, and our hybrid-attention DeepThinkVLA under different inference settings.

Category	Method	Object (%)	Spatial (%)	Goal (%)	Long (%)	Average (%)	Rel. Inference Time
Baseline	$\pi_0$ -FAST	96.8	96.4	88.6	60.2	85.5	1.0×
AR-CoT	π <sub>0</sub> -FAST (Full CoT)	95.8	93.8	74.6	61.0	81.3	4.0×
Hybrid CoT	DeepThinkVLA (Mask CoT)	99.0	97.2	96.0	93.6	96.5	0.175×
	DeepThinkVLA (Random CoT)	97.8	94.4	60.2	87.8	85.1	$0.175 \times$
	DeepThinkVLA (Full CoT)	99.0	97.2	96.8	94.2	96.8	1.4×

**Architecture Matters for the Think-Before-Acting Paradigm.** To examine how architectural choices influence the effectiveness of the "think before acting" paradigm, we conducted experiments comparing different model structures under CoT supervision. In particular, we tested whether a standard autoregressive decoder can effectively leverage CoT, or if a specialized hybrid design is neces-

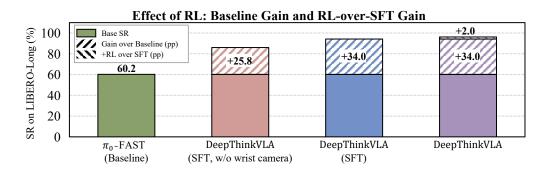


Figure 4: **Effect of RL on long-horizon task performance (LIBERO-Long)**. Bars show base SR for each model, while lighter shaded segments indicate gains over the baseline. For DeepThinkVLA, the additional teal segment highlights the extra improvement from RL over SFT (+2 pp). The figure illustrates that all DeepThinkVLA variants outperform the baseline, and RL further aligns CoT reasoning with action generation to boost success rate.

sary. Specifically, we applied CoT supervision directly to the autoregressive baseline  $\pi_0$ -FAST and compared its performance with our proposed hybrid-attention architecture.

As shown in Table 2, the autoregressive baseline with CoT supervision ( $\pi_0$ -FAST (CoT)) underperforms the original model (Average: 81.3% vs. 85.5%) and suffers a 4× increase in inference latency. This indicates that forcing a single autoregressive decoder to jointly generate reasoning and actions introduces interference, weakening both learning and efficiency. By contrast, our hybridattention architecture, which decouples reasoning and action through autoregressive causal attention for CoT and block-parallel decoding for actions, achieves markedly better results (Average: 96.8%). This represents a +15.5 percentage point gain over the naive AR-CoT baseline, with only a modest  $1.4\times$  latency overhead. These findings demonstrate that CoT supervision alone is insufficient; effective deployment of the "think before acting" paradigm requires an architecture that explicitly separates reasoning from action generation.

**RL for CoT–Action Alignment.** To further investigate how to align reasoning with effective actions, we conducted experiments applying RL on top of our SFT-trained DeepThinkVLA. The motivation is that while SFT provides CoT supervision, it relies on fixed demonstrations and may cause reasoning traces to be memorized rather than outcome-driven. RL, by contrast, can directly optimize task success, thereby encouraging the model to generate CoT that is causally linked to successful action execution.

As shown in Figure 4, applying RL increases the success rate on the challenging LIBERO-Long suite from 94.2% (SFT-only) to 96.2%, a +2 percentage point gain on top of an already strong model. This improvement demonstrates that RL effectively refines the policy by aligning reasoning and action through outcome-based optimization, overcoming the limitations of static SFT data. In a complementary finding, incorporating a wrist camera provides additional improvements by capturing near-field contact information that static scene views often miss. These findings highlight a complementary synergy: SFT equips the model with foundational reasoning ability, the wrist camera enriches perceptual grounding, and RL enforces outcome-driven alignment. Together, these factors yield superior performance on long-horizon tasks where precise, step-by-step reasoning and execution are most critical.

Case Study. To illustrate the practical benefits of our approach, we conduct a case study comparing the behavior of DeepThinkVLA with the  $\pi_0$ -FAST baseline, as depicted in Figure 5. The baseline policy successfully approaches the object but fails to secure a proper grasp. Following this error, it enters a repetitive failure loop, executing unproductive motions without making meaningful progress. This behavior highlights the inherent limitation of direct perception-to-action mapping: once an execution error occurs, the policy lacks an internal mechanism to recover, leading to task failure. In contrast, DeepThinkVLA generates a concise < think> . . . < / think> reasoning trace before each action chunk. When an accidental drop occurs, the model leverages its CoT to restate

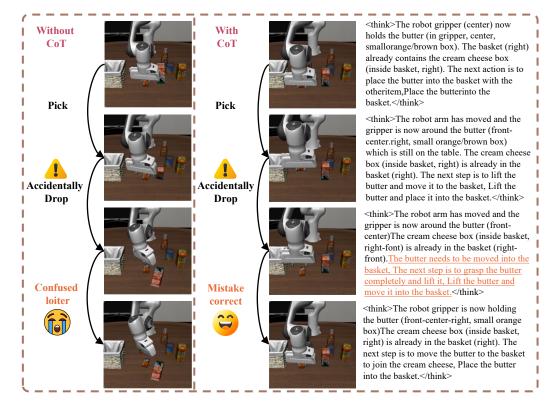


Figure 5: "Think before acting" enables error recovery. Comparison of rollouts on a LIBERO task. Left: the baseline misses the grasp and falls into a repetitive failure loop. Right Deep-ThinkVLA leverages a reasoning trace to restate the subgoal, correct mistakes and complete the task.

the subgoal (e.g., "the butter needs to be moved into the basket"), enabling it to reattempt the grasp and successfully complete the task. This demonstrates that our proposed method has effectively aligned reasoning with actions, rendering the CoT not just descriptive but actionable for guiding corrective behaviors. This case study highlights a key qualitative difference: while purely reactive policies are brittle to perturbations, DeepThinkVLA leverages CoT reasoning as an explicit planning and recovery scaffold, enabling more stable and resilient execution.

## 5 CONCLUSIONS AND FUTURE WORK

This work addresses critical challenges in the "think before acting" paradigm for VLA models. We identify two fundamental limitations in existing methods. First, they face an architectural conflict, using a single autoregressive decoder for both sequential reasoning and parallel motor commands. Second, they suffer from a training deficiency, where the generated reasoning is not causally linked to task success. To resolve these issues, we introduce DeepThinkVLA, an integrated approach that co-designs the model architecture and training strategy. Architecturally, we propose a hybridattention decoder that uses causal attention for sequential CoT generation and bidirectional attention for parallel action decoding. For training, we employ a two-stage pipeline: an initial SFT stage instills foundational reasoning, followed by an outcome-based RL stage that aligns this reasoning with task goals. Our empirical results demonstrate the effectiveness of this synergistic design. The hybrid architecture is essential for leveraging CoT and significantly outperforms a naive autoregressive baseline. Furthermore, the RL stage yields additional performance gains by optimizing reasoning for functional utility. Collectively, these innovations enable DeepThinkVLA to establish a new state-of-the-art on the LIBERO benchmark. Our work underscores that integrating explicit reasoning into robotics requires a holistic approach that tightly aligns model architecture with training objectives.

# REPRODUCIBILITY STATEMENT

All code, configuration files, and scripts to reproduce our results are available in the anonymous repository submitted with this paper. Code available at https://anonymous.4open.science/r/DeepThinkVLA

# **ETHICS STATEMENT**

This work uses publicly available, license-compliant datasets and simulated environments only; no personally identifiable information, human subjects, or animal experiments are involved. For embodied settings, all evaluations are conducted in simulation with safety constraints (termination conditions, joint/force limits, and reset policies) to avoid unsafe behaviors; any real-world deployment should include additional risk assessments and hardware interlocks.

We acknowledge that foundation models may carry societal biases. To mitigate this, we employ schema checks and consistency filters during data construction, report failure cases, and avoid tasks with sensitive attributes. Our method could be misused to automate unsafe manipulation; we strongly discourage use in safety-critical contexts without appropriate supervision and compliance.

We report configurations and code to facilitate reproducibility. There are no conflicts of interest or undisclosed sponsorships. The authors have read and will adhere to the Code of Ethics.

# REFERENCES

- 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*, 2024.
- Johan Bjorck, Fernando Castañeda, Nikita Cherniadev, Xingye Da, Runyu Ding, Linxi Fan, Yu Fang, Dieter Fox, Fengyuan Hu, Spencer Huang, et al. Gr00t n1: An open foundation model for generalist humanoid robots. *arXiv preprint arXiv:2503.14734*, 2025.
- Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. \pi\_0: A vision-language-action flow model for general robot control. In *Robotics: Science and Systems*, 2025.
- Qingwen Bu, Jisong Cai, Li Chen, Xiuqi Cui, Yan Ding, Siyuan Feng, Shenyuan Gao, Xindong He, Xuan Hu, Xu Huang, et al. Agibot world colosseo: A large-scale manipulation platform for scalable and intelligent embodied systems. *arXiv preprint arXiv:2503.06669*, 2025a.
- Qingwen Bu, Yanting Yang, Jisong Cai, Shenyuan Gao, Guanghui Ren, Maoqing Yao, Ping Luo, and Hongyang Li. Univla: Learning to act anywhere with task-centric latent actions. *arXiv* preprint arXiv:2505.06111, 2025b.
- William Chen, Suneel Belkhale, Suvir Mirchandani, Oier Mees, Danny Driess, Karl Pertsch, and Sergey Levine. Training strategies for efficient embodied reasoning. *arXiv* preprint arXiv:2505.08243, 2025.
- 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*, pp. 02783649241273668, 2023.
- Can Cui, Pengxiang Ding, Wenxuan Song, Shuanghao Bai, Xinyang Tong, Zirui Ge, Runze Suo, Wanqi Zhou, Yang Liu, Bofang Jia, et al. Openhelix: A short survey, empirical analysis, and open-source dual-system vla model for robotic manipulation. *arXiv preprint arXiv:2505.03912*, 2025.
- Hao-Shu Fang, Hongjie Fang, Zhenyu Tang, Jirong Liu, Chenxi Wang, Junbo Wang, Haoyi Zhu, and Cewu Lu. Rh20t: A comprehensive robotic dataset for learning diverse skills in one-shot. In *IEEE International Conference on Robotics and Automation*, 2024.

- Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, Pierre Sermanet, Tomas Jackson, Noah Brown, Linda Luu, Sergey Levine, Karol Hausman, and brian ichter. Inner monologue: Embodied reasoning through planning with language models. In *Proceedings of The Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research*, pp. 1769–1782. PMLR, 2023.
  - Chia-Yu Hung, Qi Sun, Pengfei Hong, Amir Zadeh, Chuan Li, U Tan, Navonil Majumder, Soujanya Poria, et al. Nora: A small open-sourced generalist vision language action model for embodied tasks. *arXiv preprint arXiv:2504.19854*, 2025.
  - Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, et al. Droid: A large-scale in-the-wild robot manipulation dataset. In *Robotics: Science and Systems*, 2024.
  - Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan P Foster, Pannag R Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. Open-VLA: An open-source vision-language-action model. In *Conference on Robot Learning*, 2024.
  - Moo Jin Kim, Chelsea Finn, and Percy Liang. Fine-tuning vision-language-action models: Optimizing speed and success. In *Robotics: Science and Systems*, 2025.
  - Fanqi Lin, Ruiqian Nai, Yingdong Hu, Jiacheng You, Junming Zhao, and Yang Gao. Onetwovla: A unified vision-language-action model with adaptive reasoning. *arXiv preprint arXiv:2505.11917*, 2025.
  - 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.
  - Songming Liu, Lingxuan Wu, Bangguo Li, Hengkai Tan, Huayu Chen, Zhengyi Wang, Ke Xu, Hang Su, and Jun Zhu. RDT-1b: a diffusion foundation model for bimanual manipulation. In *The Thirteenth International Conference on Learning Representations*, 2025a.
  - Yang Liu, Weixing Chen, Yongjie Bai, Xiaodan Liang, Guanbin Li, Wen Gao, and Liang Lin. Aligning cyber space with physical world: A comprehensive survey on embodied ai. *IEEE/ASME Transactions on Mechatronics*, 2025b.
  - Yuecheng Liu, Dafeng Chi, Shiguang Wu, Zhanguang Zhang, Yaochen Hu, Lingfeng Zhang, Yingxue Zhang, Shuang Wu, Tongtong Cao, Guowei Huang, et al. Spatialcot: Advancing spatial reasoning through coordinate alignment and chain-of-thought for embodied task planning. *arXiv* preprint arXiv:2501.10074, 2025c.
  - Yuejiang Liu, Jubayer Ibn Hamid, Annie Xie, Yoonho Lee, Max Du, and Chelsea Finn. Bidirectional decoding: Improving action chunking via closed-loop resampling. *International Conference on Learning Representations*, 2025d.
  - Guanxing Lu, Wenkai Guo, Chubin Zhang, Yuheng Zhou, Haonan Jiang, Zifeng Gao, Yansong Tang, and Ziwei Wang. Vla-rl: Towards masterful and general robotic manipulation with scalable reinforcement learning. *arXiv preprint arXiv:2505.18719*, 2025.
  - Yueen Ma, Zixing Song, Yuzheng Zhuang, Jianye Hao, and Irwin King. A survey on vision-language-action models for embodied ai. *arXiv preprint arXiv:2405.14093*, 2024.
- 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. In *IEEE International Conference on Robotics and Automation*, pp. 6892–6903. IEEE, 2024.
- 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. In *Robotics: Science and Systems*, 2025.

- Delin Qu, Haoming Song, Qizhi Chen, Yuanqi Yao, Xinyi Ye, Yan Ding, Zhigang Wang, JiaYuan Gu, Bin Zhao, Dong Wang, et al. Spatialvla: Exploring spatial representations for visual-language-action model. *arXiv preprint arXiv:2501.15830*, 2025.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- Wenxuan Song, Jiayi Chen, Pengxiang Ding, Han Zhao, Wei Zhao, Zhide Zhong, Zongyuan Ge, Jun Ma, and Haoang Li. Accelerating vision-language-action model integrated with action chunking via parallel decoding. *arXiv preprint arXiv:2503.02310*, 2025.
- Huajie Tan, Yuheng Ji, Xiaoshuai Hao, Minglan Lin, Pengwei Wang, Zhongyuan Wang, and Shanghang Zhang. Reason-rft: Reinforcement fine-tuning for visual reasoning. *arXiv* preprint *arXiv*:2503.20752, 2025.
- BAAI RoboBrain Team, Mingyu Cao, Huajie Tan, Yuheng Ji, Minglan Lin, Zhiyu Li, Zhou Cao, Pengwei Wang, Enshen Zhou, Yi Han, et al. Robobrain 2.0 technical report. *arXiv preprint arXiv:2507.02029*, 2025.
- Gemini Team. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, et al. Octo: An open-source generalist robot policy. *arXiv preprint arXiv:2405.12213*, 2024.
- Homer Rich Walke, Kevin Black, Tony Z Zhao, Quan Vuong, Chongyi Zheng, Philippe Hansen-Estruch, Andre Wang He, Vivek Myers, Moo Jin Kim, Max Du, et al. Bridgedata v2: A dataset for robot learning at scale. In *Conference on Robot Learning*, pp. 1723–1736. PMLR, 2023.
- Zhiqiang Wang, Hao Zheng, Yunshuang Nie, Wenjun Xu, Qingwei Wang, Hua Ye, Zhe Li, Kaidong Zhang, Xuewen Cheng, Wanxi Dong, et al. All robots in one: A new standard and unified dataset for versatile, general-purpose embodied agents. *arXiv preprint arXiv:2408.10899*, 2024.
- Kun Wu, Chengkai Hou, Jiaming Liu, Zhengping Che, Xiaozhu Ju, Zhuqin Yang, Meng Li, Yinuo Zhao, Zhiyuan Xu, Guang Yang, et al. Robomind: Benchmark on multi-embodiment intelligence normative data for robot manipulation. In *Robotics: Science and Systems*, 2025.
- Zhiyuan Xu, Kun Wu, Junjie Wen, Jinming Li, Ning Liu, Zhengping Che, and Jian Tang. A survey on robotics with foundation models: toward embodied ai. *arXiv preprint arXiv:2402.02385*, 2024.
- Jingyun Yang, Max Sobol Mark, Brandon Vu, Archit Sharma, Jeannette Bohg, and Chelsea Finn. Robot fine-tuning made easy: Pre-training rewards and policies for autonomous real-world reinforcement learning. In *IEEE International Conference on Robotics and Automation*, pp. 4804–4811. IEEE, 2024.
- Michał Zawalski, William Chen, Karl Pertsch, Oier Mees, Chelsea Finn, and Sergey Levine. Robotic control via embodied chain-of-thought reasoning. In *Conference on Robot Learning*, 2024.
- Zijian Zhang, Kaiyuan Zheng, Zhaorun Chen, Joel Jang, Yi Li, Siwei Han, Chaoqi Wang, Mingyu Ding, Dieter Fox, and Huaxiu Yao. Grape: Generalizing robot policy via preference alignment. arXiv preprint arXiv:2411.19309, 2024.
- Baining Zhao, Ziyou Wang, Jianjie Fang, Chen Gao, Fanhang Man, Jinqiang Cui, Xin Wang, Xinlei Chen, Yong Li, and Wenwu Zhu. Embodied-r: Collaborative framework for activating embodied spatial reasoning in foundation models via reinforcement learning. *arXiv preprint arXiv:2504.12680*, 2025a.
- Qingqing Zhao, Yao Lu, Moo Jin Kim, Zipeng Fu, Zhuoyang Zhang, Yecheng Wu, Zhaoshuo Li, Qianli Ma, Song Han, Chelsea Finn, et al. Cot-vla: Visual chain-of-thought reasoning for vision-language-action models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 1702–1713, 2025b.

Ruijie Zheng, Yongyuan Liang, Shuaiyi Huang, Jianfeng Gao, Hal Daumé III, Andrey Kolobov, Furong Huang, and Jianwei Yang. Tracevla: Visual trace prompting enhances spatial-temporal awareness for generalist robotic policies. *arXiv preprint arXiv:2412.10345*, 2024.

Yifan Zhong, Fengshuo Bai, Shaofei Cai, Xuchuan Huang, Zhang Chen, Xiaowei Zhang, Yuanfei Wang, Shaoyang Guo, Tianrui Guan, Ka Nam Lui, et al. A survey on vision-language-action models: An action tokenization perspective. *arXiv preprint arXiv:2507.01925*, 2025.

Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, et al. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Conference on Robot Learning*, pp. 2165–2183. PMLR, 2023.

# A APPENDIX

1	A.1 PROMPT OF AUTOMATED PIPELINE FOR DATA CONSTRUCTION
	Role Assignment
	You are an advanced robotic intelligence agent.
	Inputs
	• Global Task Instruction - overall goal.
	Keyframes - an ordered list of N external-camera images capturi:
	critical moments.
	- Frame 1 is before any subtask; Frame i (i > 1) is after subtask
	-1 and before subtask i. Required Output
	For each keyframe (in order), output exactly:
	Produce exactly N consecutive ( <think> </think> <subtask> &lt;</subtask>
	subtask>) pairsone pair per image, same order, plain text, no
	extra lines. ( $\leq$ 50 words, single paragraph, no line breaks)
	<think> rules</think>
	1. Frame 1: declare initial frame.
	2. Frame i (i $>$ 1): internally compare the previous image with the
	current image to infer the effect of the last subtask, but **d
	not state Success or Failure explicitly**.
	3. Tag every instruction-relevant object with its location in the
	current image, e.g., bowl (right-front). 4. Describe spatial layout, affordances, obstacles, and reasoning
	that motivates an immediate next subgoal advancing the Global
	Task.
	5. Do not mention other frames, prior success, failure, or progres
	metrics. Just analyze the current frame.
	6. Do not mention frame indices like Frame 1; use initial frame.
	<subtask> rules</subtask>
	• State the immediate next subgoal in clear natural language, **
	without numerals or explicit quantities**.
	• Each subtask must correspond to the change that should occur
	between the current image and the next one (or to completion at the final image).
	• If the Global Task is fulfilled in the current image, output
	exactly <subtask>finish</subtask> and stop; omit remaining pai:
	·
	Global Constraints
	• Output pairs equal exactly the number of keyframes provided.
	<ul> <li>No extra text outside the mandated tags.</li> </ul>
	• No bullet points or lists.
	No numbers inside any <subtask> tag.</subtask>
	Reasoning and subtasks must align with the images, temporal flo
	and the global instruction.

Figure 6: Prompt for Constructing CoT Data at Keyframes Using a Cloud-based LVLM: The prompt is designed to guide the LVLM in generating CoT data at keyframes, primarily focusing on two aspects: (i) describing the spatial relationships present in the visual scene, and (ii) capturing the temporal relationships of subtasks between consecutive keyframes.

# A.2 IMPLEMENTATION DETAILS

DeepThinkVLA is initialized from the public  $\pi_0$ -FAST weights (Pertsch et al., 2025). We refactor the baseline policy with our hybrid-attention decoder (Sec. 3.1), yielding a 2.9B parameter model.

**Dataset Construction.** Before SFT, we construct an embodied CoT dataset using the two-stage pipeline described in Sec. 3.3, based on the public LIBERO demonstrations. This process yields 273,465 annotated frames, which serve as the supervision source for the cold-start stage.

**Supervised Fine-Tuning (SFT).** For the SFT stage, we train with a batch size of 128 and a learning rate of  $2.5 \times 10^{-5}$  for 150k steps. A hybrid attention mask is employed so that CoT tokens are supervised autoregressively while action tokens are supervised bidirectionally within the same forward pass. Optimization is performed using a standard token-level cross-entropy loss.

**Reinforcement Learning (RL).** For the RL stage, we adopt Group Relative Policy Optimization (GRPO) (Shao et al., 2024). The action chunk size is set to 10. Each trajectory receives a sparse task-success reward plus a small format-regularization reward to maintain CoT quality. Policy updates use a mini-batch size of 128, a low clip ratio  $\epsilon = 0.2$ , a high clip ratio  $\epsilon = 0.28$ , and a KL penalty to the SFT reference model to avoid catastrophic forgetting.

**Infrastructure and Inference.** Training is conducted on 8×NVIDIA A800 GPUs. At inference time, we use greedy decoding for both reasoning and action tokens. For CoT ablation experiments, we evaluate three inference modes: Full CoT, Mask CoT, and Random CoT.

# A.3 BASELINES DETAILS

 We benchmark DeepThinkVLA against representative VLA systems spanning training paradigms and decoding families used in Table 1.

**Autoregressive SFT models.** TraceVLA (Zheng et al., 2024), Octo (Team et al., 2024), Open-VLA (Kim et al., 2024), SpatialVLA (Qu et al., 2025), NORA (Hung et al., 2025), and UniVLA (Bu et al., 2025b); these learn a direct perception-to-action mapping with a single AR decoder.

**RL-augmented AR variants.** GRAPE (Zhang et al., 2024), VLA-RL (Lu et al., 2025)) fine-tune AR VLAs with outcome-driven objectives but retain reactive decoding.

**Diffusion-based policies.** Diffusion Policy (Chi et al., 2023),  $\pi_0$  (Black et al., 2025)) generate actions via iterative denoising to better capture multimodal action distributions.

**Parallel-decoding policies.** CoT-VLA-7B (Zhao et al., 2025b), OpenVLA-OFT (Kim et al., 2025)) use bidirectional/block decoding to predict action tokens in chunks for higher throughput. We also include strong industrial AR references:  $\pi_0$ -FAST (Pertsch et al., 2025) and  $\pi_0$  (Black et al., 2025). This suite covers supervised, RL-augmented, diffusion, and parallel-decoding approaches from both academia and industry under the unified single-model, multi-suite LIBERO protocol.

# A.4 LLM USAGE STATEMENT

During the preparation of this manuscript, the authors utilized a large language model to assist with aspects of writing, including grammar correction, rephrasing for clarity, and improving the overall prose. In accordance with ICLR 2026 policy, we confirm that all authors have carefully reviewed and edited the final text and take full responsibility for its scientific accuracy, claims, and conclusions. Beyond language polishing, we also used a cloud-hosted LVLM to generate keyframe-level chain-of-thought annotations during dataset construction; all LVLM outputs were automatically schemachecked and consistency-filtered, with manual spot audits before use.