

# ATA: Bridging Implicit Reasoning with Attention-Guided and Action-Guided Inference for Vision-Language Action Models

Cheng Yang<sup>1\*</sup>, Jianhao Jiao<sup>2\*</sup>, Lingyi Huang<sup>1</sup>, Jinqi Xiao<sup>1</sup>, Zhexiang Tang<sup>1</sup>, Yu Gong<sup>1</sup>,  
Yibiao Ying<sup>1</sup>, Yang Sui<sup>3</sup>, Jintian Lin<sup>4</sup>, Wen Huang<sup>4</sup>, Bo Yuan<sup>1†</sup>

**Abstract**—Vision-Language-Action (VLA) models rely on current observations, including images, language instructions, and robot states, to predict actions and complete tasks. While accurate visual perception is crucial for precise action prediction and execution, recent work has attempted to further improve performance by introducing explicit reasoning during inference. However, such approaches face significant limitations. They often depend on data-intensive resources such as Chain-of-Thought (CoT) style annotations to decompose tasks into step-by-step reasoning, and in many cases require additional visual grounding annotations (e.g., bounding boxes or masks) to highlight relevant image regions. Moreover, they involve time-consuming dataset construction, labeling, and retraining, which ultimately results in longer inference sequences and reduced efficiency. To address these challenges, we propose ATA, a novel training-free framework that introduces implicit reasoning into VLA inference through complementary attention-guided and action-guided strategies. Unlike CoT or explicit visual-grounding methods, ATA formulates reasoning implicitly by integrating attention maps with an action-based region of interest (RoI), thereby adaptively refining visual inputs without requiring extra training or annotations. ATA is a plug-and-play implicit reasoning approach for VLA models, lightweight yet effective. Extensive experiments show that it consistently improves task success and robustness while preserving, and even enhancing, inference efficiency.

## I. INTRODUCTION

### A. Motivation

Vision-Language Action (VLA) models [1] have made significant strides in advancing the interpretation and synthesis of multimodal information, bridging the gap between textual, visual, and action-based inputs. Rather than using modular designs, VLA models establish a new paradigm for connecting multimodal information. Building on the success of large language models, this paradigm directly leverages the characteristics of large-scale data training, aiming to achieve highly generalizable and versatile robot action models. These models combine visual inputs, textual instructions, and robot states, such as an image with accompanying instructional text, to predict the required actions and execute tasks in the physical world. By processing and integrating these diverse modalities, VLA models have demonstrated strong capabilities in a variety of embodied tasks, such as object manipulation, long-horizon planning, and real-world household task execution.

\*Equal Contribution.

†Corresponding Author.

<sup>1</sup>Rutgers University, <sup>2</sup>University College London, <sup>3</sup>Rice University,

<sup>4</sup>TCL High-Tech Development Co., Ltd.

### B. Challenges

Despite their impressive capabilities, VLA models face several challenges that hinder their scalability and practical applicability. The former is mostly hindered by the high data collection cost, while the latter impedes the model training.

From the *data collection* perspective, existing reasoning-based methods heavily rely on data-intensive resources. Tasks involving fine-grained manipulations require datasets with a high degree of precision, and explicit reasoning approaches further demand step-by-step annotations such as Chain-of-Thought (CoT) [2]. These requirements make large-scale data acquisition extremely costly and labor-intensive.

From the *annotation requirement* perspective, many methods often rely on extra supervisory annotations, such as region-level labels or pixel-wise masks, to highlight relevant parts of the input. While these annotations can improve model perception, they are typically obtained either through costly manual labeling or by leveraging external models (e.g., object detectors or segmenters), which makes the process time-consuming, expensive, and difficult to scale across diverse environments and object categories.

The *training* issues are also considerable. Pure VLA models, which directly map multimodal inputs to actions without additional reasoning, often remain fragile in complex tasks. To address this limitation, reasoning-based methods such as CoT have been proposed to enhance VLA performance. However, these methods themselves introduce significant drawbacks, requiring costly and time-consuming dataset construction, labeling, and retraining. Furthermore, they introduce longer inference sequences, which reduces inference efficiency. Training these large-scale VLA models, which can have billions of parameters, also demands extensive computational resources, often requiring hundreds of GPU-hours and substantial time investment [1].

### C. Contributions

To address these challenges, we propose ATA (**AT**tention-Guided and **Ac**tion-Guided inference), a training-free framework that introduces implicit reasoning during inference. ATA integrates attention maps with the action region of interest (RoI) to refine visual inputs adaptively in a single forward pass. Instead of slowing down inference, ATA improves efficiency by introducing implicit reasoning, leading to shorter inference time while simultaneously enhancing model performance.

Our contributions are summarized as follows:

- We propose ATA, a novel training-free framework introducing implicit reasoning into VLA models during inference.
- We introduce an attention-guided and action-guided approach that enhances visual input based on attention maps and action-based regions of interest (RoI).
- We conduct extensive experiments on multiple state-of-the-art models, including OpenVLA [1],  $\pi_0$ -fast [3] [4], HybridVLA [5], and GR00T-N1.5 [6], demonstrating the effectiveness and efficiency of ATA in improving VLA model performance across various tasks in simulation and real-world environments. Specifically, in the LIBERO [7] environment, ATA improves the performance of the OpenVLA model by 5.2% and the  $\pi_0$ -fast model by 2.0%. In the RLBench [8] environment, ATA enhances the HybridVLA model by 5.3%. In real-world block-stacking tasks using GR00T-N1.5, where the robot stacks three-layer towers with only 3 cm  $\times$  3 cm  $\times$  3 cm blocks, ATA achieves up to a 10% performance improvement in complex scenarios.

## II. RELATED WORK

### A. Vision-Language Action Models

Recent advances in VLA models have made significant strides in integrating vision, language, and action modalities, with models such as OpenVLA [1], OpenVLA-OFT [9],  $\pi_0$  [3], and  $\pi_0$ -fast [4]. OpenVLA utilizes an autoregressive architecture, processing visual inputs alongside language instructions to predict robot actions, demonstrating strong generalization in diverse manipulation tasks.  $\pi_0$  [3] enhances this paradigm by introducing flow matching to generate continuous action distributions, and  $\pi_0$ -fast [4] leverages autoregressive VLAs for highly dexterous and high-frequency tasks where standard discretization methods fail completely. Meanwhile, HybridVLA [5] further integrates diffusion and autoregressive methods within a unified architecture, aiming to exploit the generative strength of diffusion and the efficiency of autoregression to achieve robust performance across long-horizon tasks. These models contribute to better action prediction and offer promising solutions for general robot control. However, due to the use of Vision-Language Models (VLMs) [10] [11] [12], these VLA models require billion-level parameters, making them computationally expensive even when techniques like Low-Rank Adaptation (LoRA) [13] are used for fine-tuning. This makes them reliant on substantial computational resources, which poses challenges in scalability and efficiency, especially in resource-constrained environments. This motivates the need for lightweight, training-free approaches that can enhance VLA performance during inference without retraining or additional annotations.

### B. Reasoning for Vision-Language Action Models

To address the fragility of VLAs in complex tasks, recent works have incorporated explicit reasoning to improve robustness and interpretability. CoT-VLA [14] introduces visual CoT reasoning, where tasks are decomposed

into intermediate reasoning steps to guide decision-making. While this improves interpretability, it requires frame-level alignment across visual, language, and action modalities. Such alignment demands highly precise annotations for each step, which makes dataset collection extremely costly and time-consuming, and inference becomes inefficient due to longer reasoning sequences. Contrastive Region Guidance (CRG) [15] improves grounding by refining attention with region-based contrastive objectives. CRG relies on external multimodal models to explicitly mark foreground objects of interest, while suppressing or masking out other objects. This explicit region annotation enhances robustness but requires heavy dependence on external annotations and models, making data preparation labor-intensive and less scalable to broader domains. From Seeing to Doing bridges symbolic reasoning and visuomotor control by combining high-level planning with low-level execution policies. While this improves transparency, the approach requires computationally expensive training and leads to slower inference, thus limiting its applicability in large-scale or real-time VLA scenarios. API [16] exploits attention maps from the Vision Transformer (ViT) [17] during inference to enhance image inputs, thereby improving the performance of Vision-Language Models (VLMs). While API shares a similar spirit of leveraging implicit reasoning at inference time, it relies solely on ViT attention maps that are agnostic to the action modality, limiting its applicability to VLA models. In contrast, ATA further introduces an action-guided strategy that encodes end-effector motion intent into directional RoIs, along with a frequency-controlled mechanism to balance performance and efficiency. Furthermore, MOKA [18] proposes mark-based visual prompting, where affordance reasoning is transformed into visual question-answering by annotating marks on images. Although this method enables effective zero-shot generalization, it still relies on manual design of prompts and marks, which constrains scalability. Finally, OneTwoVLA [19] unifies reasoning and acting by adaptively switching between explicit reasoning and direct action generation, but still depends on large-scale reasoning-centric data and suffers from inference latency. Overall, existing reasoning-based methods improve interpretability and long-horizon planning, yet often face trade-offs between performance, data cost, and efficiency. In contrast, ATA is a training-free framework that injects implicit reasoning into VLA inference via attention- and action-guided strategies, improving both performance and efficiency without retraining or additional supervision.

## III. METHODOLOGY

In this section, we introduce the ATA inference framework for Visual-Language-Action (VLA) models, which integrates both Attention-guided and Action-guided strategies. First, we define the overall problem setup in Section III-A. We then present a detailed breakdown of the Attention-guided strategy in Section III-B and the Action-guided strategy in Section III-C. Finally, we describe how these two strategies

are integrated during the inference process in Section III-D to optimize model performance.

### A. Problem Formulation

VLA models aim to predict a sequence of executable actions  $\mathcal{A} = \{\mathbf{a}_{t+1:t+h}\}$  over a future prediction horizon of  $h$  timesteps, by taking multi-modal information as input at timestep  $t$ , including visual observations  $\mathbf{o}_t$ , language instructions  $\mathbf{z}$  (commonly represented as the text embedding vector), and robot states  $\mathbf{s}_t = \{\mathbf{x}_t, \boldsymbol{\theta}_t, \text{grip}_t\}$ . This process can be formalized as:

$$\mathbf{a}_{t+1:t+h} = \pi_\theta(\mathbf{z}, \mathbf{o}_t, \mathbf{s}_t), \quad \mathbf{a}_i = [\Delta\mathbf{x}, \Delta\boldsymbol{\theta}, \Delta\text{grip}]^\top, \quad (1)$$

where  $\pi_\theta$  represents the trained VLA model, and  $\mathbf{a}_{t+1:t+h}$  denotes an action chunk with horizon size  $h$ . In the single-arm setting, each action  $\mathbf{a}_i$  corresponds to a 7-DoF control vector, where  $\Delta\mathbf{x} \in \mathbb{R}^3$  denotes the incremental Cartesian displacement of the end-effector relative to the previous step,  $\Delta\boldsymbol{\theta} \in \mathbb{R}^3$  represents the incremental orientation offset in axis angles, and  $\Delta\text{grip} \in \mathbb{R}$  specifies the change in the gripper command. To achieve more realistic performance, we propose to enhance the visual input  $\mathbf{o}_t$  during the inference process by incorporating both attention-guided and action-guided strategies.

We first define the sequential process to replace the original process at timestep  $t$ :

$$\mathbf{h}_t^{(\ell)} = \text{Layer}_\ell(\pi_\theta; \mathbf{z}, \mathbf{o}_t, \mathbf{s}_t), \quad \mathbf{a}_{t+1} = \pi_\theta(\mathbf{z}, \mathbf{o}_t, \mathbf{s}_t)[0], \quad (2)$$

$$\mathbf{M}_t^{\text{att}} = \Psi_{\text{att}}(\mathbf{h}_t^{(\ell)}), \quad \mathbf{M}_t^{\text{act}} = \Psi_{\text{act}}(\mathbf{a}_{t+1}), \quad (3)$$

$$\mathbf{o}'_t = \mathcal{E}(\mathbf{o}_t; \mathbf{M}_t), \quad \mathbf{M}_t \in \{\mathbf{M}_t^{\text{att}}, \mathbf{M}_t^{\text{act}}\}, \quad (4)$$

$$\mathbf{a}_{t+1:t+h} = \pi_\theta(\mathbf{o}'_t, \mathbf{z}, \mathbf{s}_t), \quad (5)$$

where  $\mathbf{h}_t^{(\ell)}$  denotes the hidden representation at the  $\ell$ -th attention layer, and  $\mathbf{a}_{t+1}$  denotes the first action step of the predicted action chunk  $\mathbf{a}_{t+1:t+h}$ .  $\Psi_{\text{att}}(\cdot)$  and  $\Psi_{\text{act}}(\cdot)$  denote the mapping functions that generate the attention-guided mask  $\mathbf{M}_t^{\text{att}}$  and the action-guided mask  $\mathbf{M}_t^{\text{act}}$ , respectively. The updated observation  $\mathbf{o}'_t$  is obtained by applying the inference-time guidance operator  $\mathcal{E}(\cdot)$  with either of the masks, *i.e.*,  $\mathbf{M}_t \in \{\mathbf{M}_t^{\text{att}}, \mathbf{M}_t^{\text{act}}\}$ .

As illustrated in Fig. 1, the original inference pipeline predicts actions solely from the raw observation  $\mathbf{o}_t$ . While this pipeline works reasonably well in many cases, it lacks the ability to leverage implicit reasoning available within the model. Consequently, an early misprediction at a single frame may propagate across the horizon, leading to cascading errors and eventual task failure. In contrast, our ATA framework introduces attention-guided and action-guided inference as implicit reasoning cues derived from attention maps and robot actions, which intervene directly on the observation stream at inference time. This implicit reasoning mechanism mitigates error propagation, improves robustness, and ultimately enhances the success rate of VLA models. To address this issue, our framework introduces periodic inference-time guidance on visual observations through attention-guided and action-guided strategies. Rather than relying solely on the

raw observation  $\mathbf{o}_t$  at every step, we selectively replace  $\mathbf{o}_t$  with the updated observation  $\mathbf{o}'_t$  defined in Eq. (5). Concretely, we propose two complementary strategies:

- **Attention-Guided Strategy:** An intermediate attention map from a designated layer of the VLA model is used as an implicit reasoning cue to highlight task-relevant objects and suppress irrelevant background regions, ensuring that the model's perception remains aligned with the instruction.
- **Action-Guided Strategy:** The robot's end-effector state is leveraged to construct a directional region of interest (RoI), implicitly reasoning about the intended action trajectory by emphasizing regions along the motion direction while de-emphasizing unrelated areas.

By injecting the guided observations into the pipeline at selected frames, the model receives refined visual inputs that help maintain a correct reasoning trajectory. As illustrated in Fig. 1, this stage-wise guidance enables the model to produce more reliable action predictions, thereby improving the overall task success rate. This naturally raises three key research questions: 1) *How to design the attention-guided strategy?* 2) *How to construct the action-guided strategy?* 3) *How to adaptively apply these strategies during inference?*

### B. Attention-Guided

To answer the first question, we begin by analyzing the structure of VLA blocks. Each block consists of Layer Normalization, an attention layer, another Layer Normalization, and an MLP, all connected through residual links to ensure stability and effective learning. The output token sequence of one block is then passed to the next. For the attention-guided strategy, we exploit the attention map generated inside the attention layer, since it reflects which regions of the input the model attends to during inference timestep  $t$ . The input sequence includes both image and text tokens. Within the attention layer, each token computes similarity scores between its query vector and the key vectors of all other tokens, followed by a softmax normalization to produce attention weights. To construct our guidance signal, we focus on the last query token, which typically aggregates high-level contextual information across the entire sequence. From this token, we extract the attention weights corresponding only to the image tokens, and average them across all attention heads to obtain the aggregated attention map,  $\Psi_{i,j}^{(L)} = \frac{1}{H} \sum_{h=1}^H \text{Attn}_h^{(L)}$ , where  $\text{Attn}_h^{(L)}$  denotes the attention weights at layer  $L$  between the last query token and the image tokens for head  $h$ . Here,  $(i, j)$  indexes the spatial location of the image token in the patch grid, so that  $\Psi_{i,j}^{(L)}$  represents the aggregated attention weight assigned to the patch at position  $(i, j)$ .

After extracting the attention map  $\Psi_{i,j}^{(L)}$ , we transform it into a mask that can serve as an inference-time guidance signal. To make the map effective, we first normalize it by subtracting the mean and dividing by the standard deviation

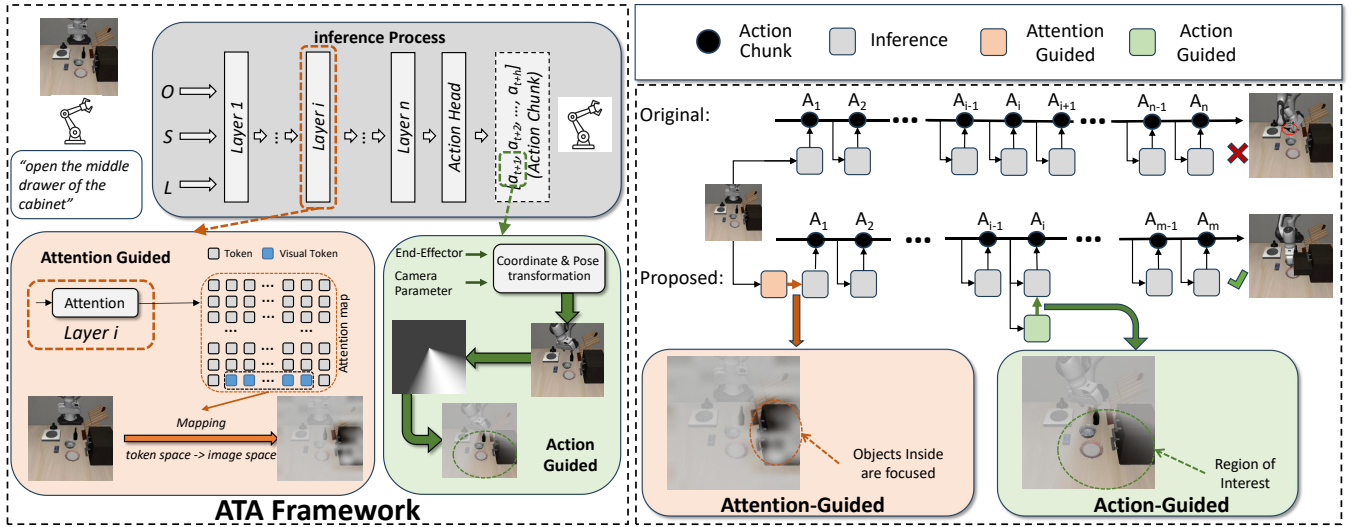


Fig. 1: Overview of the proposed ATA framework. **(Left)** Given a language instruction and multimodal observations, the VLA model processes tokens through stacked layers to produce an action chunk. Our framework injects implicit reasoning cues via two complementary strategies: (i) Attention-guided reasoning extracts attention maps from intermediate layers to focus on task-relevant visual regions; (ii) Action-guided reasoning leverages the end-effector pose and camera parameters to construct a directional region of interest (RoI). **(Right)** In the original pipeline (top), an error at an early step may propagate along the horizon (up to  $n$  steps), leading to task failure (marked with  $\times$ ) and risky behaviors such as knocking over a bottle (highlighted with a red circle). The proposed pipeline (bottom) applies guidance at selected steps, resulting in a shorter effective horizon ( $m < n$ ), successful task completion (marked with  $\checkmark$ ), and more robust execution.

across all pixels, followed by a sigmoid mapping:

$$\mathbf{M}_t^{\text{attn}}(i, j) = \text{Sigmoid} \left( \frac{\Psi_{i,j}^{(L)} - \mu(\Psi)}{\sigma(\Psi)} \right), \quad (6)$$

Where  $\mathbf{M}_t^{\text{attn}}(i, j)$  is the normalized attention mask,  $\mu(\Psi)$  and  $\sigma(\Psi)$  are the mean and standard deviation of the attention map  $\Psi_{i,j}^{(L)}$ , respectively. The sigmoid function maps the values to the range  $[0, 1]$ , ensuring that they are suitable for acting as a mask. Importantly, our attention maps are obtained by explicitly computing the similarity between queries and keys, independent of the value-update operation. This makes the procedure fully compatible with efficient implementations such as FlashAttention [20], without affecting its acceleration benefits. The final step is to apply the enhanced attention map as a mask to the original image. This mask will highlight the most relevant regions of the image while suppressing irrelevant ones. The updated visual observation  $\mathbf{o}'_t$  is obtained as follows:

$$\mathbf{o}'_t = \mathbf{o}_t \odot \mathbf{M}_t^{\text{attn}} + (1 - \mathbf{M}_t^{\text{attn}}) \odot \mathbf{bg}, \quad (7)$$

where  $\mathbf{M}_t^{\text{attn}}$  is broadcast to image channels, and  $\mathbf{bg}$  is typically set to a neutral value like gray. This process yields guided observations that emphasize task-relevant areas while suppressing distractors, thereby providing implicit reasoning cues for the model during inference.

### C. Action-Guided

To answer the second question, we propose an action-guided strategy that exploits the robot's end-effector (EEF)

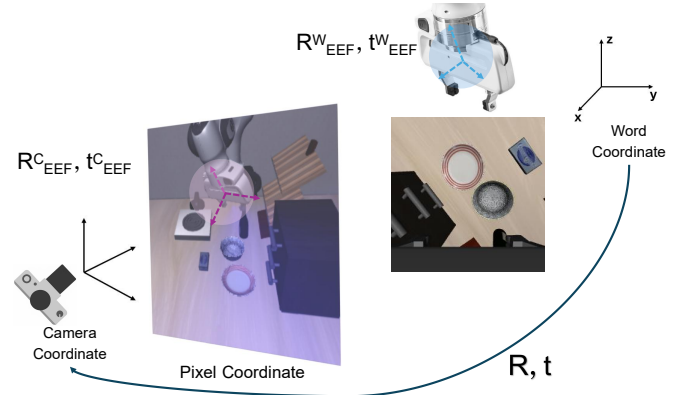


Fig. 2: Illustration of the action-guided strategy. The end-effector (EEF) pose ( $\mathbf{R}_{\text{EEF}}^W, \mathbf{t}_{\text{EEF}}^W$ ) is obtained in the world coordinate system and transformed into the camera coordinate ( $\mathbf{R}_{\text{EEF}}^C, \mathbf{t}_{\text{EEF}}^C$ ), allowing projection of the EEF direction onto the image plane. The resulting region of interest (RoI) is represented as a conic sector in the pixel coordinate system. For visualization, we overlay a semi-transparent red mask so that pixels inside the RoI appear in red, while the actual inference mask is computed from Eq. (9).

state as an implicit reasoning cue. Unlike the attention-guided strategy, which derives signals from the model's internal representations, the action-guided strategy directly encodes the robot's motion intent into a directional region of interest (RoI) on the image plane.

At timestep  $t$ , the EEF pose is given by  $(\mathbf{x}_t, \boldsymbol{\theta}_t)$ , where  $\mathbf{x}_t \in \mathbb{R}^3$  is the Cartesian position and  $\boldsymbol{\theta}_t$  is the axis angles

converted from the orientation quaternion. From  $\theta_t$  we obtain the rotation matrix  $\mathbf{R}_t$ , and select a local tool axis (e.g., the  $z$ -axis) to define the intended motion direction  $\mathbf{d}_t = \mathbf{R}_t \mathbf{e}_z$ . Together with the base point  $\mathbf{x}_t \in \mathbb{R}^3$ , this direction specifies a conic sector parameterized by an opening angle  $\alpha$  and a depth-adaptive length  $z_{\text{depth}}$ . To map this sector into the image, we use the camera parameters to project both the base point  $\mathbf{x}_t$  and the tip point  $\mathbf{x}_t + z_{\text{depth}} \mathbf{d}_t$  into the pixel plane, yielding the base pixel  $(u_0, v_0)$  and the projected direction vector  $(d_u, d_v)$ . For each pixel  $(u, v)$ , we measure the angular deviation from this direction:

$$\psi = \frac{(u - u_0, v - v_0) \cdot (d_u, d_v)}{\|(u - u_0, v - v_0)\| \cdot \|(d_u, d_v)\|}, \quad (8)$$

where pixels with  $\psi \geq \cos(\alpha/2)$  fall inside the sector. To provide smooth transitions at the boundary, we define a soft mask:

$$\mathbf{M}_t^{\text{act}}(u, v) = \max\left[\frac{\psi - \cos(\alpha/2)}{1 - \cos(\alpha/2)}, 0\right]. \quad (9)$$

In practice, we set the opening angle to  $\alpha = 150^\circ$  (empirically,  $\alpha \geq 120^\circ$  also works). This wide opening angle ensures that the manipulated object remains within the RoI while still emphasizing the intended motion direction. Moreover, a larger  $\alpha$  produces smoother mask boundaries, which avoids jagged edges and maintains high visual quality during image blending. Finally, the enhanced observation is obtained by applying this mask to the original image:

$$\mathbf{o}'_t = \mathbf{o}_t \odot \mathbf{M}_t^{\text{act}} + (1 - \mathbf{M}_t^{\text{act}}) \odot \mathbf{bg}, \quad (10)$$

where  $\mathbf{M}_t^{\text{act}}$  is broadcast to image channels, and  $\mathbf{bg}$  is a neutral background (e.g., gray). As illustrated in Fig. 2, regions along the predicted motion direction are emphasized while irrelevant areas are suppressed, injecting action intent as an implicit reasoning signal into the inference pipeline.

#### D. Inference-Time Integration

To address the third question, we revisit the standard inference pipeline of VLA models. At each timestep  $t$ , the model predicts an action chunk from multimodal inputs. Since the initial robot state is usually fixed, variations in predictions depend largely on the quality of the visual observations. Early frames are especially critical: small errors at the beginning may propagate across the horizon and cause cascading failures.

Our key insight is that injecting implicit reasoning cues into the observation stream early in the task can stabilize the reasoning trajectory. Specifically, we apply **attention-guided reasoning** at the very first frame to direct the model’s focus toward task-relevant regions, thereby improving the prediction of the initial action chunk. This corresponds to the inference pipeline in Eq. (6) with  $\mathbf{M}_t^{\text{attn}}$ . While attention-guided reasoning improves semantic understanding of the scene, it does not fully leverage the robot’s physical interaction state. To complement it, we apply **action-guided reasoning** based on the EEF pose during the early stages of

---

#### Algorithm 1: ATA Inference Algorithm

---

**Input:**  $\mathbf{o}, \mathbf{z}, \mathbf{s}, \pi_\theta, M, Freq, i^{\text{act}}, \mathbf{bg}, f^{\text{attn}}, f^{\text{act}}$   
**Output:** Action chunk  $\mathcal{A}$

- 1  $i \leftarrow 0, \mathbf{s} \leftarrow \mathbf{s}_0;$
- 2 **repeat**
- 3     **if**  $i \bmod Freq = 0$  **then**
- 4          $\mathbf{M}_i^{\text{attn}} \leftarrow f^{\text{attn}}(\pi_\theta; \mathbf{z}, \mathbf{o}_i, \mathbf{s}_i) \quad \triangleright \text{Eq. 6}$
- 5          $\mathbf{o}_i \leftarrow \mathbf{o}_i \odot \mathbf{M}_i^{\text{attn}} + (1 - \mathbf{M}_i^{\text{attn}}) \odot \mathbf{bg} \quad \triangleright \text{Eq. 7}$
- 6     **if**  $i = i^{\text{act}}$  **then**
- 7          $\mathbf{M}_i^{\text{act}} \leftarrow f^{\text{act}}(\pi_\theta; \mathbf{z}, \mathbf{o}_i, \mathbf{s}_i) \quad \triangleright \text{Eq. 9}$
- 8          $\mathbf{o}_i \leftarrow \mathbf{o}_i \odot \mathbf{M}_i^{\text{act}} + (1 - \mathbf{M}_i^{\text{act}}) \odot \mathbf{bg} \quad \triangleright \text{Eq. 10}$
- 9      $\mathcal{A}_{i+1} \leftarrow \pi_\theta(\mathbf{z}, \mathbf{o}_i, \mathbf{s}_i);$
- 10     $i \leftarrow i + 1;$
- 11 **until success or**  $i \geq M;$

---

the task. This strategy highlights regions along the intended motion direction, ensuring that the model perceives both the semantic goal and the geometric action intent. Formally, this corresponds to Eq. (9) with  $\mathbf{M}_t^{\text{act}}$ .

In practice, we integrate the proposed two strategies by applying attention-guided reasoning at the first frame and action-guided reasoning in the early steps. This combined schedule improves action predictions with minimal computational overhead, as it only introduces an extra forward pass when guidance is applied. We find that periodic application of attention-guided reasoning cues further boost performance; however, overly frequent usage may inject noise into the visual stream and reduce efficiency, as discussed in Section IV.

#### E. Overall Algorithm

The sequential inference process of ATA is summarized in Alg. 1, where inference ends upon task success or when the maximum step  $M$  is reached. At each timestep  $t$ , the model receives the visual observation  $\mathbf{o}_t$ , instruction  $\mathbf{z}$ , and state  $\mathbf{s}_t$ , and outputs an action chunk. To inject implicit reasoning, we apply two guidance functions:  $f^{\text{attn}}$  (Sec. III-B) and  $f^{\text{act}}$  (Sec. III-C). These functions generate masks that update the observation before it enters the policy  $\pi_\theta$ . Here  $\mathbf{bg}$  is a neutral background,  $Freq$  is the trigger frequency of Attention-Guided reasoning, and  $i^{\text{act}}$  is the timestep for Action-Guided reasoning.

## IV. EXPERIMENTS

### A. Experimental Settings

1) *Model Settings:* To demonstrate the effectiveness of our proposed method, we apply it to several state-of-the-art VLA models, including OpenVLA,  $\pi_0$ -fast, HybridVLA, and GR00T-N1.5 models, and evaluate them across both simulation and real-world environments. Specifically, we first evaluate the performance of OpenVLA and  $\pi_0$ -fast models in the LIBERO simulation environment, and the HybridVLA model in the RLBench simulation environment. Additionally, we perform real-world experiments using the GR00T-N1.5 model to validate the practical effectiveness of ATA.

TABLE I: **Best configurations of Attention-Guided and Action-Guided strategies.**  $L$  denotes the attention layer index,  $Freq$  denotes the trigger frequency. For  $\pi_0$ -fast and HybridVLA, the attention-guided strategy is applied only once ( $Freq = \infty$ ) at the first timestep ( $t = 0$ ).

Model	Dataset	L	Freq	Time	Step
OpenVLA	Spatial	20	100	periodic	50
	Goal	5	100	periodic	80
	Object	20	100	periodic	60
	10	20	50	periodic	170
$\pi_0$ -fast	LIBERO	16/12/12/12	$\infty$	$t = 0$	10/7/5/2
HybridVLA	RLBench	5	$\infty$	$t = 0$	3

2) *Evaluation and Datasets:* In the LIBERO simulation environment, we evaluate our method on four datasets: *LIBERO-Spatial*, *LIBERO-Goal*, *LIBERO-Object*, and *LIBERO-Long*, each containing 500 demonstrations across 10 tasks. Specifically, LIBERO-Spatial focuses on spatial generalization with identical objects placed in different spatial relations, LIBERO-Object evaluates generalization to novel objects with a unique object in each task, LIBERO-Goal tests goal generalization with fixed objects and layouts but varying task goals, and LIBERO-Long consists of long-horizon tasks to assess the model’s ability to handle extended action sequences. In the RLBench environment, we test the model on 8 tasks<sup>1</sup>.

In the real world, we collect a dataset of 500 stacking trajectories in a  $90cm \times 50cm$  workspace, where blue and red blocks are randomly placed at predefined positions. Each trajectory, built with  $3cm \times 3cm \times 3cm$  blocks, consists of robot actions, language instruction, and visual observations from three cameras (head, wrist, and external). The dataset includes 100 one-layer, 150 two-layer, and 250 three-layer towers, covering both simple and challenging tasks while balancing diversity and collection cost for fine-tuning GROOT-N1.5. For evaluation, we randomly sample 50 distinct block configurations and locations to compare performance with and without ATA. The real-world robotic system is shown in the left part of Tab. V, featuring a 7-DoF Discover ARM with wrist, head, and external cameras.

3) *Implementation Details:* For the OpenVLA and  $\pi_0$ -fast experiments in the LIBERO environment, as well as the HybridVLA experiments in the RLBench environment, all experiments are conducted on a single H20 GPU. For the GROOT-N1.5 experiments, we perform training on 8 NVIDIA A800 GPUs with mixed-precision training, resulting in a total training time of 35 hours over 90,000 steps.

In our experiments, we explore different configurations of the Attention-Guided and Action-Guided strategies. As shown in Tab. I, for OpenVLA, the best results are obtained with dataset-specific layer choices  $L \in \{20, 5, 20, 20\}$  and periodic triggers  $Freq \in \{100, 100, 100, 50\}$  across LIBERO-Spatial, Goal, Object, and Long, respectively. For  $\pi_0$ -fast, due to its long-horizon action chunking, the Attention-Guided strategy is more effective when applied

<sup>1</sup>Tasks include: close box, toilet seat down, sweep to dustpan, close fridge, close laptop, take umbrella out, wine at rack, water plants

only once at the first timestep ( $t = 0$ ). Similarly, for HybridVLA in RLBench, where tasks require only short action sequences, Attention-Guided reasoning also achieves the best performance when triggered once at  $t = 0$ . For all models, the Action-Guided strategy is applied only at the initial timestep, since early intervention provides the most benefit for guiding subsequent actions.

To ensure fairness, both LIBERO and RLBench experiments are executed under maximum step limits, where exceeding the predefined horizon is considered a task failure. For LIBERO, we follow the configurations in OpenVLA and  $\pi_0$ , with maximum step lengths of 220 for Spatial, 300 for Goal, 280 for Object, and 520 for LIBERO-10. For RLBench, consistent with HybridVLA, we set the maximum horizon to 10 steps for all tasks.

## B. Main Results

1) *Simulation Results:* The results presented in Tab. II show the performance of the models after applying our ATA framework, with the *Average Success Rates (S.R.)*, *Average Number of Inference Calls* for successfully completed tasks (S.I.C), and *Overall Average Number of Inference Calls (I.C.)*. For the API [16] method, the improvements are relatively limited since it only leverages the ViT component and does not incorporate deeper attention or action information. Specifically, on the LIBERO dataset, API brings only a 0.4% gain for OpenVLA and  $\pi_0$ -fast, while in RLBench, it improves HybridVLA by just 1.7% over the baseline. When only the Attention-Guided method is used, compared to the baseline, OpenVLA and  $\pi_0$ -fast models show improvements of 3.7% and 1.4% respectively on the LIBERO dataset. When combined with the Action-Guided method, the performance improves further to 5.2% for OpenVLA and 2% for  $\pi_0$ -fast, highlighting the complementary nature of the two strategies. Similarly, for the HybridVLA model tested in the RLBench, the Attention-Guided method alone provides a 2.3% improvement over the baseline. When the Action-Guided method is integrated, the improvement increases to 5.5%, demonstrating the additive benefits of combining both strategies.

Notably, ATA improves both performance and efficiency. For OpenVLA, the average inference calls drop from 235 to 225 (Avg I.C.) and from 165 to 164 (Avg S.I.C.). Similar reductions are observed for  $\pi_0$ -fast (41  $\rightarrow$  39, 30  $\rightarrow$  29) and HybridVLA (8  $\rightarrow$  7, 7  $\rightarrow$  6). Although ATA requires an extra step to introduce implicit reasoning, it ultimately lowers the overall number of inference calls. Since simulator runs involve costly environment resets, we report average inference calls instead of wall-clock time as a reliable measure of efficiency.

2) *Real-World Results:* In the real world, we collected 500 stacking block trajectories, where blocks are stacked between one, two, and three towers with red and blue colors. We validated the performance using a single-arm 7-DoF Discover Arm. As shown in Tab. V, the ATA method improves the performance by 2%, 2%, and 6% for stacking one, two, and three towers of blocks, respectively, compared

TABLE II: **Comparison of ATA and Baselines on LIBERO and RLbench.** LIBERO includes four datasets, each with 10 tasks and 50 demonstrations per task. The RLbench environment evaluates 8 tasks, each with 100 demonstrations. *Avg S.R.* denotes the average success rate. *Attn* denotes the use of only the Attention-Guided strategy, and *ATA* integrates both Attention-Guided and Action-Guided strategies. *Avg S.I.C.* indicates the average number of inference calls required for successfully completed tasks, and *Avg I.C.* represents the overall average number of inference calls across all tasks.

Models	Method	LIBERO-Spatial	LIBERO-Goal	LIBERO-Object	LIBERO-Long	Avg S.R.	Avg S.I.C.↓	Avg I.C.↓
OpenVLA	Baseline	84.2	75.6	88.6	55.2	75.9	165	235
OpenVLA	API	86.2	78.2	88.6	52	76.3	167	232
OpenVLA	ATTN	87	82	92	57.4	79.6	165	228
OpenVLA	ATA	88.8	82.6	94.6	58.2	81.1	164	225
$\pi_0$ -fast	Baseline	96.2	89	98.2	60.2	85.9	30	41
$\pi_0$ -fast	API	97.6	88.8	98	60.8	86.3	31	41
$\pi_0$ -fast	ATTN	98.2	90.2	98.6	62.6	87.3	30	40
$\pi_0$ -fast	ATA	98.6	90.8	98.8	63.2	87.9	29	39

Models	Method	Close box	Close laptop lid	Toilet seat down	Sweep to dustpan	Close fridge	Umbrella out	Wine at rack	Water plants	Avg S.R.	Avg S.I.C.↓	Avg I.C.↓
HyBridVLA	Baseline	85	90	90	100	90	40	50	25	71.3	7	8
HyBridVLA	ATTN	87	92	92	97	95	43	53	30	73.6	7	7
HyBridVLA	ATA	88	95	100	98	97	46	55	35	76.8	6	7

TABLE III: **Comparison of Blur, RandomBlur, Attention-Guided, and Baselines on LIBERO and RLbench datasets.** *Blur* represents the Gaussian Blur applied to the first frame, *RandomBlur* denotes the Gaussian Blur applied randomly to frames after the first frame, and *ATTN\** indicates the attention-guided strategy applied only on the first frame.

Models	Method	LIBERO-Spatial	LIBERO-Goal	LIBERO-Object	LIBERO-Long	Avg S.R.
OpenVLA	Baseline	84.2	75.6	88.6	55.2	75.9
OpenVLA	Blur	83.4	74.6	84.2	53.2	73.9
OpenVLA	RandomBlur	84.6	77.2	87.2	53.8	75.7
OpenVLA	ATTN*	86.6	80.8	89.2	56.4	78.3

Models	Method	Close box	Close laptop lid	Toilet seat down	Sweep to dustpan	Close fridge	Umbrella out	Wine at rack	Water plants	Avg S.R.
HyBridVLA	Baseline	85	90	90	100	90	40	50	25	71.3
HyBridVLA	Blur	40	45	25	20	85	35	15	5	42.7
HyBridVLA	RandomBlur	55	85	95	20	75	50	20	25	53.1
HyBridVLA	ATTN*	87	92	92	97	95	43	53	30	73.6

TABLE IV: **Comparison of Different Frequencies of Attention-Guided Strategy on LIBERO Dataset Performance.**  $0$  denotes the baseline,  $\infty$  denotes applying Attention-Guided (*ATTN\**) only on the first timestep ( $t = 0$ ), and  $20, 50, 100, 200$  indicate applying Attention-Guided every 20, 50, 100, and 200 inference steps, respectively.

Models	Freq	Spatial	Goal	Object	Long
OpenVLA	0 (Baseline)	84.2	75.6	88.6	55.2
OpenVLA	$\infty$ ( <i>ATTN*</i> )	86.6	80.8	89.2	56.4
OpenVLA	20	83.9	78.8	87.8	55.6
OpenVLA	50	86	79.6	89.2	57.4
OpenVLA	100	87	82	92	56.4
OpenVLA	200	86.2	81.2	89.6	55.8

to the baseline. Additionally, during the three-tower stacking test, we introduced irrelevant and unseen objects such as blocks of different colors, pens, scissors, etc. We found that the ATA method further improves robustness, achieving a 10% increase in performance compared to the baseline.

The bottom part of Tab. V illustrates task execution sequences: the top sequence depicts stacking a three-layer tower in a simple scene containing only red blocks, corresponding to the language instruction “Pick up three red blocks and stack them on the red point to form a three-layer

tower.”. In contrast, the bottom sequence depicts stacking a three-layer tower in a more complex scene with unseen distractor objects such as scissors, pens, and additional colored blocks (orange and green), guided by the instruction “Pick up three blue blocks and stack them on the red point to form a three-layer tower.”.

### C. Ablation Studies.

1) *The Impact of the First Frame on Success Rates:* As discussed in Sec. III-D, the first frame plays a crucial role in setting the task context. To validate this, we compare three scenarios on OpenVLA (LIBERO) and HybridVLA (RLBench): 1) Gaussian Blur ( $9 \times 9$ ) on the first frame; 2) Gaussian Blur on random subsequent frames; 3) Attention-Guided strategy on the first frame.

As shown in Tab. III, blurring the first frame causes a 2% drop on LIBERO and a 28.6% drop on RLbench, whereas blurring random later frames only leads to 0.2% and 18.2% drops, respectively. In contrast, applying the Attention-Guided strategy to the first frame improves performance by 2.4% on LIBERO and 2.3% on RLbench. The larger impact on RLbench is due to its shorter task horizons, which amplify early frame effects.

TABLE V: **Real-world stacking results with and without ATA.** Performance is measured on stacking *one*, *two*, and *three*-layer towers of blocks. For the three-tower case, we also tested robustness by adding irrelevant and unseen objects (*Three\**). The bottom-left shows the real-world robotic setup, while the right part illustrates task execution sequences in both simple (*top*) and complex (*bottom*) scenes.

Models	Method	One	Two	Three	Three*
Gr00t-N1.5	Baseline	92	85	68	46
Gr00t-N1.5	ATA	94	87	74	56

2) *The Impact of Different Frequencies of Attention-Guided Strategy:* As shown in Tab. IV, applying the Attention-Guided strategy only at the first timestep ( $t = 0$ , *ATTN\**) improves the *average* success rate by 2.4% over the baseline. We then vary the trigger frequency on LIBERO. When the frequency is set between 50 and 100 inference steps, the model attains the best *average* performance, yielding a 3.7% gain over the baseline and a 1.3% gain over *ATTN\**. In contrast, excessively sparse (e.g., every 200 steps) or overly frequent (e.g., every 20 steps) triggers degrade performance, indicating diminishing returns.

## V. CONCLUSION

In this work, we presented ATA, a training-free framework that introduces implicit reasoning for VLA models through attention-guided and action-guided strategies. ATA adaptively refines visual inputs, remains compatible with efficient attention implementations, and requires no additional annotations or retraining. Experiments show that ATA improves task success and robustness while preserving, and even enhancing, inference efficiency, offering a lightweight yet effective path toward scalable VLA deployment. Looking forward, ATA provides a general paradigm for incorporating reasoning into VLA models without costly data collection, and can be extended to broader multimodal reasoning and control tasks. As future work, we aim to develop adaptive mechanisms for automatically selecting the optimal attention layer and trigger frequency in the Attention-Guided strategy, as well as determining the timing of Action-Guided intervention, thereby reducing dataset-specific tuning and further improving the generality of ATA.

## REFERENCES

[1] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. P. Foster, P. R. Sanketi, Q. Vuong *et al.*, “Openvla: An open-source vision-language-action model,” in *Conference on Robot Learning*. PMLR, 2025, pp. 2679–2713.

[2] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou *et al.*, “Chain-of-thought prompting elicits reasoning in large language models,” *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.

[3] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter *et al.*, “ $\pi_0$ : A vision-language-action flow model for general robot control,” *arXiv preprint arXiv:2410.24164*, 2024.

[4] K. Pertsch, K. Stachowicz, B. Ichter, D. Driess, S. Nair, Q. Vuong, O. Mees, C. Finn, and S. Levine, “Fast: Efficient action tokenization for vision-language-action models,” *arXiv preprint arXiv:2501.09747*, 2025.

[5] J. Liu, H. Chen, P. An, Z. Liu, R. Zhang, C. Gu, X. Li, Z. Guo, S. Chen, M. Liu *et al.*, “Hybridvla: Collaborative diffusion and autoregression in a unified vision-language-action model,” *arXiv preprint arXiv:2503.10631*, 2025.

[6] J. Bjorck, F. Castañeda, N. Cherniadev, X. Da, R. Ding, L. Fan, Y. Fang, D. Fox, F. Hu, S. Huang *et al.*, “Gr00t n1: An open foundation model for generalist humanoid robots,” *arXiv preprint arXiv:2503.14734*, 2025.

[7] B. Liu, Y. Zhu, C. Gao, Y. Feng, Q. Liu, Y. Zhu, and P. Stone, “Libero: Benchmarking knowledge transfer for lifelong robot learning,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 44 776–44 791, 2023.

[8] S. James, Z. Ma, D. R. Arrojo, and A. J. Davison, “Rlbench: The robot learning benchmark & learning environment,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3019–3026, 2020.

[9] M. J. Kim, C. Finn, and P. Liang, “Fine-tuning vision-language-action models: Optimizing speed and success,” *arXiv preprint arXiv:2502.19645*, 2025.

[10] S. Karamcheti, S. Nair, A. Balakrishna, P. Liang, T. Kollar, and D. Sadigh, “Prismatic vlms: Investigating the design space of visually-conditioned language models,” in *Forty-first International Conference on Machine Learning*, 2024.

[11] P. Wang, S. Bai, S. Tan, S. Wang, Z. Fan, J. Bai, K. Chen, X. Liu, J. Wang, W. Ge *et al.*, “Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution,” *arXiv preprint arXiv:2409.12191*, 2024.

[12] G. Team, T. Mesnard, C. Hardin, R. Dadashi, S. Bhupatiraju, S. Pathak, L. Sifre, M. Rivière, M. S. Kale, J. Love *et al.*, “Gemma: Open models based on gemini research and technology,” *arXiv preprint arXiv:2403.08295*, 2024.

[13] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen *et al.*, “Lora: Low-rank adaptation of large language models,” *Iclr*, vol. 1, no. 2, p. 3, 2022.

[14] Q. Zhao, Y. Lu, M. J. Kim, Z. Fu, Z. Zhang, Y. Wu, Z. Li, Q. Ma, S. Han, C. 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*, 2025, pp. 1702–1713.

[15] D. Wan, J. Cho, E. Stengel-Eskin, and M. Bansal, “Contrastive region guidance: Improving grounding in vision-language models without training,” in *European Conference on Computer Vision*. Springer, 2024, pp. 198–215.

[16] R. Yu, W. Yu, and X. Wang, “Attention prompting on image for large vision-language models,” in *European Conference on Computer Vision*. Springer, 2024, pp. 251–268.

[17] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly *et al.*, “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv preprint arXiv:2010.11929*, 2020.

[18] F. Liu, K. Fang, P. Abbeel, and S. Levine, “Moka: Open-world robotic manipulation through mark-based visual prompting,” *arXiv preprint arXiv:2403.03174*, 2024.

[19] F. Lin, R. Nai, Y. Hu, J. You, J. Zhao, and Y. Gao, “Onetwovla: A unified vision-language-action model with adaptive reasoning,” *arXiv preprint arXiv:2505.11917*, 2025.

[20] T. Dao, D. Fu, S. Ermon, A. Rudra, and C. Ré, “Flashattention: Fast and memory-efficient exact attention with io-awareness,” *Advances in neural information processing systems*, vol. 35, pp. 16 344–16 359, 2022.