# DRIFT: A RESIDUAL FLOW ADAPTER FOR DECODING CONTINUOUS OUTPUTS IN VISION-LANGUAGE MODELS

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Many modern vision-language models (VLMs) build on autoregressive decoding of discrete tokens. While text-based output interfaces enable scalable pretraining and strong zero-shot generalization across diverse tasks, they are poorly suited for problems that require precise continuous outputs, such as localizing temporal boundaries of events or generating robotic control actions. To address this challenge, we propose DRIFT, a general framework for adapting pretrained VLMs to continuous decoding tasks. DRIFT combines a base predictor, which provides a coarse estimate of the target output, with a generative refinement module based on flow matching that iteratively improves the prediction. This residual formulation transforms the generative modeling problem from learning a global output distribution to modeling a localized residual distribution around a strong prior, substantially simplifying optimization. We evaluate DRIFT on both perception and planning tasks, including visual grounding and robotic control. Across multiple tasks and architectures spanning MLLMs, VLAs, and WAMs, DRIFT consistently outperforms a strong set of regression- and generative-based solutions.

## 1 INTRODUCTION

Vision-language models (VLMs) (Liu et al., 2023b; Alayrac et al., 2022) pretrained on large-scale visual and textual corpora have achieved remarkable success across perception and planning tasks, leading to the emergence of multimodal large language models (MLLMs) (Bai et al., 2025a; Deitke et al., 2025), vision-language-action models (VLAs) (Kim et al., 2024; Black et al., 2025b;a), and world action models (WAMs) (Yuan et al., 2026; Ye et al., 2026). Modern VLMs typically adopt a generic autoregressive backbone (*e.g.*, Transformers (Vaswani et al., 2017)) and decode discrete text tokens as their primary output interface. This design enables scalable pretraining via next token prediction, and supports flexible zero-shot transfer to downstream tasks whose outputs can be naturally expressed as text, including classification, captioning, and question answering.

However, many important perception and planning tasks fundamentally require predicting continuous quantities. Examples include estimating event timestamps in temporal video grounding (TVG) (Sigurdsson et al., 2016; Krishna et al., 2017) and generating control signals for robotic manipulation in VLA tasks (Liu et al., 2023a; Li et al., 2025). The discrete token interface is poorly suited to such tasks: tokenization introduces quantization errors, discards ordinal structure, and ignores fine-grained variation within bins. In this work, we ask the following question: "*How can pretrained VLMs with discrete decoding interfaces be adapted for precise continuous prediction while preserving and leveraging the knowledge acquired during pretraining?*"

Existing approaches only partially address this challenge. In MLLMs, prior work often discretizes continuous targets into tokens (Kim et al., 2024; Bai et al., 2025a) or appends special tokens with lightweight regression heads (Zhang et al., 2026). In VLAs and WAMs, diffusion- and flow-based action decoders have emerged as dominant approaches for generating continuous outputs (Black et al., 2025b;a; Yuan et al., 2026). Tokenization and direct regression are deterministic and thus struggle to model uncertainty in continuous prediction, including multimodal output distributions (Chi et al., 2025) and high-variance modes (Moltisanti et al., 2017; Zhang et al., 2025) (Figure 1-2,3). Diffusion- and flow-based approaches alleviate this issue by iteratively refining pre-



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

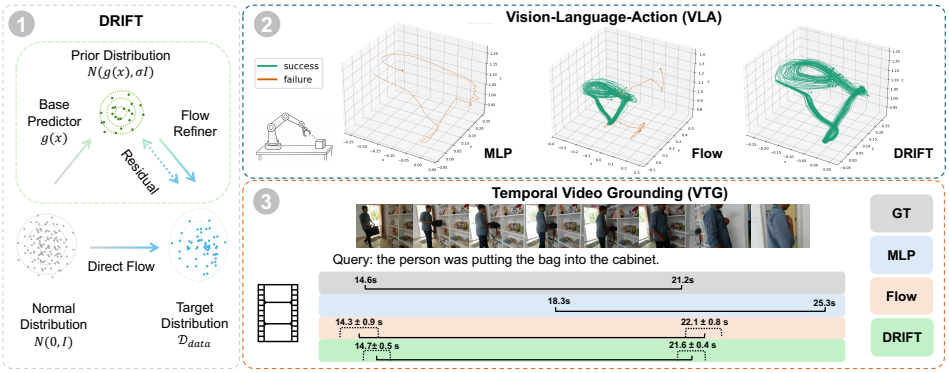


Figure 1: **(1)**: **DRIFT** learns a velocity field that transports a starting distribution centered at an initial coarse prediction toward the target distribution, simplifying optimization by modeling localized residuals. **(2, 3)**: **Visualization** of predicted action trajectories (on Libero-Long) and temporal event boundaries (on Charades-STA) under different decoding strategies. Compared to MLP or flow matching, DRIFT produces more accurate and stable predictions with lower variance.

dictions through generative modeling in continuous space. However, they typically require more training data and optimization effort, can exhibit high prediction variance in low-data regimes (Figure 1-2,3), and must devote modeling capacity to global structure that can be readily captured by deterministic predictors.

*Our key insight* is that deterministic prediction and iterative generative refinement can form a natural cascade. Deterministic predictors leverage pretrained representations to produce efficient coarse estimates, while generative refinement effectively captures uncertainty and recovers fine-grained structure around these estimates. Leveraging this insight, we propose **DRIFT** (Decoding via Residual Flow Adapter), a general framework for adapting pretrained VLMs to continuous decoding tasks. DRIFT introduces a lightweight residual flow adapter that augments a base predictor such as a token decoder or multilayer perceptron (MLP) regressor, with a flow-matching-based refinement module. Rather than generating predictions from scratch, DRIFT models the residual error distribution around an initial coarse prediction and progressively refines it using flow matching (Figure 1-1).

We evaluate DRIFT on perception (spatial and temporal grounding with MLLMs) and planning (robotic control with VLAs and WAMs). Across all settings, DRIFT consistently improves over both direct regression baselines and standalone flow-matching decoders, while achieving state-of-the-art performance on several TVG and VLA benchmarks. On TVG benchmarks, DRIFT achieves gains of +2.6% on Charades-STA and +2.1% on ActivityNet Captions compared to the latest methods (Pramanick et al., 2025). In VLA tasks, DRIFT built on Qwen3-VL-2B (Bai et al., 2025a) achieves 97.9% average accuracy on Libero benchmarks and 61.5% accuracy on Simpler WidowX.

**Our main contributions** are summarized as follows.

- We introduce **DRIFT**, a general framework for adapting pretrained vision-language models with discrete decoding interfaces to precise continuous prediction tasks.
- We develop a residual refinement adapter that combines efficient coarse prediction with flow-matching-based refinement, together with theoretical analysis of its optimization advantages and an effective training strategy.
- Through extensive experiments across visual grounding and robotic manipulation benchmarks, we demonstrate consistent gains over strong baselines across multiple VLM architectures.

## 2 Related Work

**Vision Language Models (VLMs).** VLMs learn from paired visual and textual data to address a broad range of vision-language tasks. In particular, multi-modal large language models (MLLMs) (Liu et al., 2023b; Bai et al., 2025a; Deitke et al., 2025; Li et al., 2023; Alayrac et al., 2022) have emerged as powerful frameworks for learning transferable cross-modal representations from massive-scale image-text and video-text datasets, enabling sophisticated multimodal under-

standing and reasoning (Fu et al., 2025; Yue et al., 2024; Li et al., 2024). Beyond passive perception, another line of VLM research focuses on action control, leading to vision-language-action (VLA) models (Kim et al., 2024; Black et al., 2025b;a). These models are typically pretrained on large-scale action trajectories to improve cross-embodiment generalization and can be adapted to diverse downstream manipulation tasks with limited fine-tuning. More recently, researchers have begun to connect scene generation with action control in world-action models (WAMs) (Yuan et al., 2026; Ye et al., 2026), which aim to model both environmental dynamics and action-conditioned interactions.

**Vision Tasks with Continuous Output.** Many perception and planning tasks require outputs that are inherently continuous rather than purely discrete. In this work, we focus on two representative tasks: (1) visual grounding (Xiao et al., 2025), which localizes the spatial or temporal extent of a visual concept in an image or video given a language query, and (2) VLA tasks (Liu et al., 2023a; Li et al., 2025), which predict continuous robot action trajectories conditioned on visual observations and language instructions. Prior work has explored several ways to adapt VLMs to such continuous-output tasks. For visual grounding, some methods (Bai et al., 2025a; Clark et al., 2026; Liu et al., 2024a; Sun et al., 2025) discretize box coordinates or timestamps into tokens, while others (Pramanick et al., 2025; Mu et al., 2024; Lin et al., 2023) attach an MLP head to directly regress continuous values. For VLA tasks, existing approaches either use a simple MLP decoder to predict action trajectories from VLM representations (Zhang et al., 2026), or adopt diffusion- or flow-matching-based heads (Black et al., 2025b; Chen et al., 2025; Zitkovich et al., 2023) to improve action precision.

Unlike these task-specific decoding strategies, DRIFT provides an efficient and general adapter for extending pretrained VLMs to continuous-output tasks. By refining continuous predictions around a learned prior, our framework supports diverse perception and planning problems while preserving the reusable Transformer backbone of existing VLMs.

**Diffusion and Flow Matching.** Diffusion models (Ho et al., 2020; Song et al., 2021) seek to reverse a forward "diffusion" process that progressively injects noise into data by learning a denoising function. Flow matching models (Lipman et al., 2023; Liu et al., 2023c) simplify the training of continuous-time normalizing flows as supervised regression over target vector fields, leading to an objective closely connected to diffusion models. While both frameworks were initially designed for generative modeling (Rombach et al., 2022; Polyak et al., 2024), recent research has expanded their utility to continuous value decoding in VLA (Chi et al., 2025; Black et al., 2025b; Yang et al., 2026).

# 3 DRIFT: Decoding Continuous Output via Residual Flow Adapter

## 3.1 Preliminaries: Decoding Continuous Outputs in Autoregressive VLMs

**Autoregressive VLMs**. An autoregressive VLM takes an image or video $\mathbf{X}^v$ and a text query $\mathbf{X}^q = \{x^q\}$ as its input, and generates an answer $\mathbf{X}^a = \{x^a\}$ in text form. Specifically, $\mathbf{X}^v$ is first encoded by a visual encoder $E_v(\cdot)$ (including a vision backbone and its projector) into a set of vision tokens $\{\mathbf{z}^v \in \mathbb{R}^d\}$. Similarly, $\mathbf{X}^q$ is processed by a text encoder $E_t(\cdot)$, which embeds the words $\{x^q\}$ into a set of text tokens $\{\mathbf{z}^q \in \mathbb{R}^d\}$. These tokens are combined into $\{\mathbf{z}^{v|q}\} = [\{\mathbf{z}^v\}, \{\mathbf{z}^q\}]$, and processed by $f(\cdot)$ (often an LLM) that decodes the answer $\mathbf{X}^a$ autoregressively

$$z_i^a = f\left(\left[\{\mathbf{z}^{v|q}\}, \{\mathbf{z}_{<i}^a\}\right]; \theta_f\right), \quad x_i^a = D_t(z_i^a), \tag{1}$$

where $D_t(\cdot)$ is the decoding function that converts the embedding $z_i^a$ to the answer token $x_i^a$ in text, and $\{\mathbf{z}_{<i}^a\}$ are text tokens from previously generated answer $x_{<i}^a$, *i.e.* $\mathbf{z}^a = E_t(x^a)$.

**Decoding Continuous Outputs**. For many vision tasks, $\mathbf{X}^a$ involves a continuous output. For simplicity, we consider a single continuous output $y_i = x_i^a \in \mathbb{R}^c$ at slot $i$ with dimension $c$, though the formulation can be easily extended to multiple variables. Without loss of clarity, we drop the subscript $i$ and the superscript $a$ hereafter. To decode a continuous $y$, a separate decoding function $g(z) \to y$ is often needed. When *tokenization* is applied to $y$, $g$ can re-use the discrete decoder $D_t$. In *direct regression*, a special token $x^s$ representing the continuous output is introduced, and

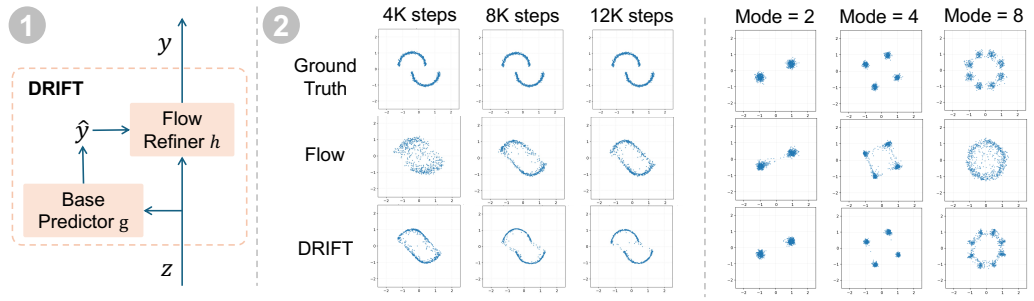


Figure 2: **(1)**: **Overview of DRIFT**, which combines a base predictor with a generative refinement module for continuous decoding conditioned on VLM embedding $z$. **(2)**: **Toy experiments** on a synthetic 2D dataset projected into a high-dimensional space with a fixed, random, column-orthogonal projection matrix, following (Li & He, 2025). A 3-layer MLP is trained to model the velocity field in the projected space. The projection matrix is unknown to the model and used only for visualization. **(2-Left)**: DRIFT achieves faster convergence than flow matching. **(2-Right)**: As the underlying distribution becomes more multimodal, DRIFT more accurately recovers the target distribution.

$g$ is realized as a simple MLP. The decoder $g$ is activated only when $x^s$ is generated. When using *diffusion or flow match* heads, a noise or time dependent version of $g$ will be learned to model the conditional probability $p(y|z)$, again, with the help of special tokens $x^s$.

This continuous decoding function $g$ can be provided as part of the modern VLMs, often in the form of tokenization or a simple MLP. For example, Qwen3-VL (Bai et al., 2025a) has incorporated visual grounding into pre-training by representing numerical quantities as discrete text tokens. When $g$ is not provided, we implement $g$ with an MLP that is jointly learned with DRIFT.

**Flow Matching**. Flow matching (Lipman et al., 2023) or diffusion (Chi et al., 2025) heads can be used for continuous decoding. Conditional flow matching learns a velocity field $v_\theta(y(t), t, z) \triangleq dy(t)/dt$, typically a neural network with parameters $\theta$, that transports samples from a known noise distribution to a target distribution $p(y|z)$. With target $y(t=1) \in \mathbb{R}^c$, noise $y(t=0) \in \mathbb{R}^c \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$, and time $t \sim \mathcal{U}(0,1)$, the linear interpolant $y(t) = (1-t)\, y(t=0) + t\, y(t=1)$ has ground-truth velocity $v^* = y(t=1) - y(t=0)$. The standard flow-matching loss is

$$\mathcal{L}_{\text{FM}} = \mathbb{E}\left[\|v_\theta(y(t), t, z) - v^*\|_2^2\right]. \tag{2}$$

Training minimizes Eq. (2) with respect to the parameters $\theta$ in the velocity field $v_\theta(\cdot)$. Inference predicts $y(t=1)$ by solving the ODE $dy/dt = v_\theta(y(t), t, z)$ given $y(t=0)$ and $z$.

### 3.2 Design of DRIFT

**Problem Formulation**. Our goal is to decode continuous output $y \in \mathbb{R}^c$ from VLM embedding $z$ without modifying existing VLM architecture. Our key intuition is that a simple base predictor $g$, built on the learned VLM representation, can provide a strong initial estimation and such estimation can be further refined by a generative model.

To this end, we propose learning a velocity field $v_\theta(y(t), t, z)$ that transports a Gaussian distribution $\mathcal{N}(\hat{y} \triangleq g(z), \sigma \boldsymbol{I})$ centered at the coarse prediction $\hat{y} \triangleq g(z)$, towards the target conditional distribution $p(y|z)$. Here $\sigma \in \mathbb{R}^c$ scales the variance of the Gaussian and is learned from data. In contrast to standard flow matching, which transports $\mathcal{N}(\mathbf{0}, \boldsymbol{I})$ to $p(y|z)$, our approach, illustrated in Figure 2-1, initializes starting distribution using the coarse prediction $\hat{y}$, offering a strong prior over $y$.

**Parameterization**. We parameterize the velocity field as

$$v_\theta(y(t), t, z) = \frac{y_\theta(y(t), t, z) - y(t)}{1-t}, \tag{3}$$

where $y_\theta(\cdot)$ aims to directly predict the target output $y(t=1)$. This parameterization was recently shown to be competitive with velocity prediction in generation tasks (Li & He, 2025).

Predicting the target $y$ allows us to add *a skip connection* to the base predictor $g$. This is given by

$$y_\theta(y(t), t, z) = \alpha\, g(z) + (1 - \alpha)\, h_\theta(y(t), t, z). \tag{4}$$

$h_\theta(\cdot)$ is time-dependent and can be considered as the *refinement module*. It iteratively refines the residual between $y$ and the coarse prediction $\hat{y}$. $\alpha$ is a learnable gating vector that selectively combines the outputs of the base predictor $g$ and the refinement module $h_\theta(\cdot)$.

**Learning Target**. We consider the following linear interpolant, also referred to as the *bridge state*,

$$y(t) = (1 - t)(\sigma\epsilon + g(z)) + ty, \tag{5}$$

where $\epsilon \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$. The target velocity is given by $v^* = dy(t)/dt = y - \sigma\epsilon - \hat{y}$. We follow the standard flow matching loss in Eq. 2. We sample $t \in \mathcal{U}(0, 1)$, construct $y(t)$ by mixing a noise $\epsilon$ with a target output $y$, compute its velocity at $t$ with the corresponding condition $z$ following Eq. 3, and then minimize the mean squared error (MSE) between the predicted velocity and the target flow $v^*$. Our learnable parameters include: (1) $\theta$ in $h_\theta(\cdot)$, which models the residual error; (2) the gating vector $\alpha$, which balances the coarse prediction and its refinement; and (3) the scaling vector $\sigma$, which controls the variance of the initial Gaussian distribution.

It is worth noting that an alternative linear interpolant is $y(t) = (1 - t)\sigma\epsilon + t(y - g(z))$. When the base predictor $g(\cdot)$ is fixed, these two interpolants are mathematically equivalent. However, the equivalence no longer holds when $g(\cdot)$ is jointly learned. In that case, they present two conceptually different ways of modeling. Eq. 5 first "shifts" the Gaussian distribution to $\sigma\epsilon + g(z)$ and then models the transport from the shifted distribution to the target distribution. In contrast, the alternative interpolant here directly models the gaps between the initial Gaussian and the distribution of the residual error. Practically, Eq. 5 offers greater flexibility in modeling. For example, it allows the skip connection in our refinement module (Eq. 4). Empirically, we validate that Eq. 5 provides minor performance benefits in our ablation studies (see Section 4.4).

**Joint Training with the Base Predictor**. When the base predictor $g$ is already provided as part of the VLM, we either keep its parameters frozen or update them, based on the setting. When such a predictor is unavailable, we instantiate $g$ as an MLP regressor, supervise it with an additional MSE loss $\mathcal{L}_p$ (with the overall loss as $\mathcal{L}_{\text{FM}} + \mathcal{L}_p$), and jointly optimize its parameters. A key challenge in this joint optimization setting is that the model can collapse to relying primarily on the flow matching model while ignoring the coarse predictor $g$. We find that stable learning requires applying gradient stopping to $g(\cdot)$ when computing the linear interpolant using Eq. 5.

**Inference**. At inference, we sample Gaussian noise $\epsilon \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$ and initialize $y(0) = g(z) + \sigma\epsilon$. We then solve the ODE $dy(t)/dt = v_\theta(y(t), t, z)$ to obtain the final prediction $y(1)$.

### 3.3 Toy Experiments and Theoretical Justification

We present results of DRIFT and a vanilla flow matching model on several toy examples in Figure 2-2. DRIFT more faithfully captures the underlying data distribution by disentangling its modes.

We provide a detailed theoretical analysis in Appendix A, assuming a fixed base predictor $g(z)$. Residual refinement strictly reduces prediction error compared with using the base predictor $g(z)$ alone, as long as the bridge state $y(t)$ carries information about the residual $r = y - \hat{y}$ beyond what $g(z)$ already captures. Intuitively, $y(t)$ is a noisy interpolation between the initial prediction $\hat{y}$ and the true target $y$, so it inherently contains a partial target signal that the refinement module $h(y(t), t, z)$ can exploit to better approximate $r$. Compared with applying flow matching directly to $y$, when the base predictor already captures the condition-dependent mean structure, both the residual signal $r$ and the residual velocity field have smaller second moments, and the bridge is confined to a narrower tube around the prior. The refinement module therefore only needs to learn a local correction rather than transport noise to the full target, which explains DRIFT's faster convergence and higher accuracy.

Table 1: **Results on VLA benchmarks.** We report action success rates, where * indicates reproduced results. DRIFT-MLP and DRIFT-Tokenizer use MLP and tokenization-based base predictors, respectively. Across Libero and Simpler benchmarks, DRIFT consistently improves upon OpenVLA and achieves state-of-the-art performance with the Qwen3-VL-2B backbone.

| Method | Base Model | Action Decoder | Libero (Liu et al., 2023a)↑ | | | | | Simpler (Li et al., 2025)↑ |
|---|---|---|---|---|---|---|---|---|
| | | | Spatial | Object | Goal | Long | AVG | |
| ***Methods with VLA pretraining*** | | | | | | | | |
| OpenVLA* (Kim et al., 2024) | Prismatic-7B (Karamcheti et al., 2024) | Tokenizer | 83.1 | 88.5 | 76.4 | 54.1 | 75.5 | 7.7 |
| $\pi_0$ (Black et al., 2025b) | PaliGemma-3B (Beyer et al., 2024) | Flow | 96.8 | 98.8 | 95.8 | 85.2 | 94.1 | 27.1 |
| $\pi_{0.5}$ (Black et al., 2025a) | PaliGemma-3B (Beyer et al., 2024) | Flow | 98.8 | 98.2 | **98.0** | 92.4 | 96.9 | 57.1 |
| ***Methods based on Adaptation*** | | | | | | | | |
| VLA-Adaptor (Wang et al., 2026) | Qwen-2.5-0.5B (Yang et al., 2024) | Diffusion | 97.8 | 99.2 | 97.2 | 95.0 | 97.3 | N/A |
| VLM4VLA (Zhang et al., 2026) | Qwen3-VL-2B (Bai et al., 2025a) | MLP | N/A | N/A | N/A | 55.8 | N/A | 49.0 |
| VLM4VLA (Zhang et al., 2026) | InternVL3.5-4B (Wang et al., 2025) | MLP | N/A | N/A | N/A | 62.8 | N/A | 57.3 |
| InternVLA (Chen et al., 2025) | Qwen2.5-VL-3B (Bai et al., 2025b) | Diffusion | 98.0 | 99.0 | 93.8 | 92.6 | 95.8 | 54.2 |
| DRIFT | OpenVLA-7B (Kim et al., 2024) | DRIFT-Tokenizer | 87.0 | 89.1 | 77.5 | 57.0 | 77.7 | 13.5 |
| DRIFT | Qwen3-VL-2B (Bai et al., 2025a) | DRIFT-MLP | **98.8** | **99.2** | 97.0 | **96.4** | **97.9** | **61.5** |

# 4 Experiments and Results

## 4.1 Results of Vision-Language-Action Tasks

VLA tasks aim to predict robot actions conditioned on visual inputs and language instructions. Given an observation of the environment and a task description, a VLA model generates continuous control signals for robotic manipulation. The predicted actions are executed by the robot, after which new observations are collected and fed back into the model, resulting in a closed-loop control process.

**Experiment Design**. We instantiate DRIFT individually with Qwen3-VL-2B-Instruct (Bai et al., 2025a) and OpenVLA (Kim et al., 2024) as VLM backbones. Qwen3-VL is a recent MLLM pretrained on large-scale multimodal corpora, but it does not provide a dedicated base predictor for VLA tasks. We use this model to study how large-scale pretraining benefits continuous action decoding and to evaluate the joint optimization of the base predictor and flow-matching-based refinement. OpenVLA, in contrast, is a VLA model that predicts actions through tokenization. We include this model to demonstrate how DRIFT can refine predictions on top of an existing base predictor.

**Experiment Protocol**. We evaluate DRIFT on Libero-Spatial, Libero-Object, Libero-Goal, Libero-Long (Liu et al., 2023a), and Simpler (Li et al., 2025) with WidowX robot arm. Both Libero and Simpler are simulation based benchmarks, with a focus on tabletop object manipulation tasks. We train DRIFT on the LIBERO training set for LIBERO experiments, and on the Bridge and Fractal splits of OXE (O'Neill et al., 2024) when evaluating the policy on Simpler with WidowX. We report the action success rate as our evaluation metric.

**Implementation Details.** We attach LoRA (Hu et al., 2022) adapters to the backbone VLM, including the vision encoder and the LLM, and jointly train DRIFT with the base predictor. Appendix B.2 provides full implementation details and hyperparameters.

**Results**. The results are presented in Table 1. With OpenVLA as the base model, DRIFT achieves consistent improvement across different sub-tasks on Libero and Simpler. Particularly, DRIFT raises the Libero average from 75.5% to 77.7% and Simpler from 7.7% to 13.5%. With Qwen3-VL-2B, DRIFT reaches 97.9% average success rate on Libero and 61.5% on Simpler, achieving new state-of-the-art performance on both benchmarks. We outperform the VLM4VLA with Qwen3-VL-2B by a clear margin, demonstrating the effectiveness of DRIFT. Yet, we see a major gap between DRIFT initialized with OpenVLA and Qwen3-VL, which reveals that the pretraining of the VLM has a strong effect on the downstream task. The general improvement in the vision understanding of Qwen3-VL provides strong action representation helpful for DRIFT, and thus yields better results.

## 4.2 Results of Temporal Video Grounding

TVG aims to localize the temporal segment in a video that corresponds to a natural language query. Given a video and a text description, the model predicts continuous start and end timestamps indicating when the described event occurs.

**Experiment Design**. We adopt **ET-Chat** (Liu et al., 2024b) as the VLM backbone, following the ED-VTG (Pramanick et al., 2025) protocol for fair comparison with prior generalist baselines. ET-Chat natively localizes events by cosine-matching a generated $\langle \text{vid} \rangle$ token's hidden state to frame embeddings; we replace this matching head with an MLP that regresses $(t_{\text{start}}, t_{\text{end}})$ from the same

Table 2: **Results on temporal video grounding.** We evaluate the zero-shot generalization capability of DRIFT on the Charades-STA (Gao et al., 2017) and ActivityNet Captions (Krishna et al., 2017) test splits, and compare against existing baselines. Baseline results are taken from ED-VTG (Pramanick et al., 2025), except for ET-Chat, whose results are reproduced using the released checkpoint under the same evaluation protocol. "# Train" denotes the size of the temporal grounding adaptation corpus. DRIFT outperforms ED-VTG, the strongest generalist baseline, on both datasets, with the largest gains observed on R@1@0.3 and mIoU metrics.

| Method | # Train | Charades-STA (Gao et al., 2017)↑ | | | | ActivityNet-Captions (Krishna et al., 2017)↑ | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | R@0.3 | R@0.5 | R@0.7 | mIoU | R@0.3 | R@0.5 | R@0.7 | mIoU |
| ***Non-generalist (Specialist) Models*** | | | | | | | | | |
| UniVTG (Lin et al., 2023) | 4.2M | 44.1 | 25.2 | 10.0 | 27.1 | – | – | – | – |
| SeViLA (Yu et al., 2023) | 129M | – | – | – | – | 31.6 | 19.0 | 10.1 | 23.0 |
| PSVL (Nam et al., 2021) | – | 46.2 | 31.3 | 14.2 | 31.2 | 44.7 | 30.1 | 14.7 | 29.6 |
| LT-ZVG (Kim et al., 2023) | – | 52.9 | 37.2 | 19.3 | 36.0 | 47.6 | 32.6 | <u>15.4</u> | 31.8 |
| ***Generalist (LLM-based) Models*** | | | | | | | | | |
| Video-LLaMA (Zhang et al., 2023) | 2.7M | 25.2 | 10.6 | 3.4 | 16.8 | 21.9 | 10.8 | 4.9 | 16.5 |
| Video-ChatGPT (Maaz et al., 2024) | 100K | 27.2 | 6.2 | 1.9 | 19.7 | 19.5 | 10.6 | 4.8 | 14.2 |
| Valley (Luo et al., 2026) | 100K | 28.4 | 1.8 | 0.3 | 21.4 | 30.6 | 13.7 | 8.1 | 21.9 |
| VideoChat2 (Li et al., 2024) | 2M | 38.0 | 14.3 | 3.8 | 24.6 | 40.8 | 27.8 | 9.3 | 27.9 |
| Momentor (Qian et al., 2024) | 10M | 42.6 | 26.6 | 11.6 | 28.5 | 42.9 | 23.0 | 12.4 | 29.3 |
| VTimeLLM (Huang et al., 2024) | 170K | 51.0 | 27.5 | 11.4 | 31.2 | 44.0 | 27.8 | 14.3 | 30.4 |
| TimeChat (Ren et al., 2024) | 125K | – | 32.2 | 13.4 | – | – | – | – | – |
| HawkEye (Wang et al., 2024c) | 715K | 50.6 | 31.4 | 14.5 | 33.7 | 49.1 | 29.3 | 10.7 | 32.7 |
| ChatVTG (Qu et al., 2024) | 100K | 52.7 | 33.0 | 15.9 | 34.9 | 40.7 | 22.5 | 9.4 | 27.2 |
| ED-VTG (Pramanick et al., 2025) | 136K | 59.5 | 39.3 | 19.8 | 40.2 | <u>52.1</u> | <u>33.1</u> | **16.0** | <u>35.2</u> |
| ET-Chat (Liu et al., 2024b) | 164K | <u>65.7</u> | **45.9** | <u>20.0</u> | <u>42.3</u> | 24.1 | 12.7 | 6.2 | 18.9 |
| DRIFT (ET-Chat) | 100K | **67.2** | <u>44.8</u> | **20.6** | **43.8** | **57.9** | **35.8** | 14.1 | **37.3** |

Table 3: **Results on spatial grounding.** Integrating DRIFT into Qwen3-VL-2B improves performance on RefCOCO benchmarks.

| Method | refcoco | refcoco+ | refcocog | AVG |
|---|---|---|---|---|
| Qwen-VL (Bai et al., 2025a) | 89.6 | 81.8 | 85.5 | 85.6 |
| Qwen-VL (w/ DRIFT) | **91.7** | **85.8** | **87.9** | **88.5** |

Table 4: **Results with world action models.** Integrating DRIFT with FastWAM improves performance on the Libero benchmarks.

| Method | Spatial | Object | Goal | Long | AVG |
|---|---|---|---|---|---|
| FastWAM (Yuan et al., 2026) | 97.6 | 99.4 | 95.6 | 95.0 | 96.9 |
| FastWAM (w/ DRIFT) | **99.2** | 99.0 | **98.2** | **96.2** | **98.2** |

hidden state, providing the base prediction that DRIFT refines. The MLP, flow refiner, and LoRA adapters are jointly trained on TimeLens-100K (Zhang et al., 2025); Charades-STA and ActivityNet-Captions training splits are unseen, so all results are cross-dataset zero-shot.

**Experiment Protocol**. We evaluate on the Charades-STA (Gao et al., 2017) and ActivityNet-Captions (Krishna et al., 2017) test splits, and report Recall@1 at IoU thresholds $\tau \in \{0.3, 0.5, 0.7\}$ (R1@$\tau$) and mean IoU (mIoU). Implementation details are in Appendix B.3.

**Results.** Table 2 reports the VTG results. DRIFT (ET-Chat) outperforms ED-VTG, the strongest ED-VTG-reported generalist baseline, on 7 of 8 IoU/mIoU columns; the only exception is ActivityNet-Captions R@0.7, where DRIFT trails ED-VTG by 1.9%. On Charades-STA, where vanilla ET-Chat already exceeds ED-VTG by 6.2% at R@0.3, DRIFT adds $+1.5\%$ R@0.3 ($65.7\% \rightarrow 67.2\%$) and $+1.5\%$ mIoU ($42.3\% \rightarrow 43.8\%$), with R@0.5 within $\pm 1\%$ of ET-Chat. On ActivityNet-Captions, vanilla ET-Chat is weaker (R@0.3 24.1%, mIoU 18.9%); DRIFT lifts these to 57.9% and 37.3% respectively, surpassing ED-VTG on R@0.3, R@0.5, and mIoU.

### 4.3 Further Demonstration

**2D Spatial Grounding**. We further adapt DRIFT to 2D spatial grounding tasks. Specifically, we attach DRIFT to the Qwen3-VL-2B model. We then train and evaluate the model on the RefCOCO series datasets (Kazemzadeh et al., 2014; Mao et al., 2016; Yu et al., 2016). As shown in Table 3, DRIFT consistently improves the performance of Qwen3-VL by a clear margin across all benchmarks, demonstrating the versatility and generality of our framework. Details are in Appendix B.4.

**World Action Model**. We further explore adapting DRIFT to recent world action models. Specifically, we integrate DRIFT with FastWAM (Yuan et al., 2026). Unlike the other models considered in this paper, FastWAM already employs a flow matching based decoder for predicting continuous actions. In this setting, we introduce an additional base predictor using MLP and jointly opti-

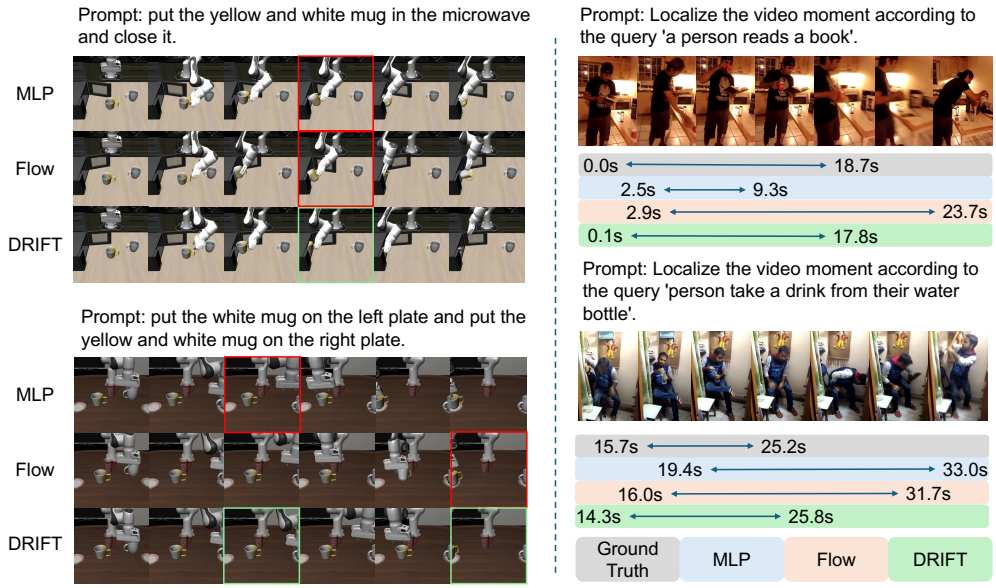


Figure 3: **Qualitative results. Left:** Integrating a base predictor and a flow-based refiner enables DRIFT to generate more accurate and stable action trajectories, improving action precision and overall task success rates. **Right:** Through iterative refinement of the coarse prediction, DRIFT produces sharper and more accurate event temporal boundaries.

Table 5: **Ablation studies on Libero-Long.** (a) DRIFT outperforms MLP, diffusion, and flow-based decoders. (b) Joint training improves the action representations. (c) The best-performing configuration uses DRIFT with $x$-prediction and a skip connection.

**(a) Decoder**

| Decoder | LL |
|---|---|
| MLP | $88.2 \pm 0$ |
| Diffusion | $92.23 \pm 0.64$ |
| Flow | $93.67 \pm 0.35$ |
| DRIFT | $\mathbf{96.33 \pm 0.23}$ |

**(b) Co-training effect**

| Decoder | LL |
|---|---|
| MLP | 88.2 |
| MLP as base | 93.0 |
| DRIFT | **96.4** |

**(c) DRIFT design**

| | | Direct Flow | | DRIFT | |
|---|---|---|---|---|---|
| | | **v-pred** | **x-pred** | **v-pred** | **x-pred** |
| **w/ skip** | **w/** $\mathcal{L}_p$ | N/A | N/A | N/A | **93.8** |
| | **w/o** $\mathcal{L}_p$ | N/A | N/A | N/A | **93.6** |
| **w/o skip** | **w/** $\mathcal{L}_p$ | 92.4 | 91.2 | 91.0 | 92.2 |
| | **w/o** $\mathcal{L}_p$ | 91.0 | 91.2 | 93.2 | 93.4 |

mize it together with FastWAM's original flow decoder. We train the model following FastWAM's training recipe, and present the results in Table 4. DRIFT consistently improves FastWAM on the Libero benchmark suite, further demonstrating the generality of our framework. Details are in Appendix B.5.

**Qualitative Results**. Figure 3-left shows actions from different decoders on Libero-Long. Small action errors can cause task failures: a poor grasp angle blocks placing the cup into the microwave (top), and a slight release offset misses the plate (bottom). By refining the base prediction, DRIFT yields more accurate actions and higher success rates. Similarly, Figure 3-right compares predictions from different decoders on TVG tasks. Exact temporal boundaries are often ambiguous. Through iterative refinement, DRIFT produces sharper temporal boundaries, leading to improved localization accuracy and overall performance.

**Inference Latency**. Table 6(b) compares end-to-end latency under matched task settings. On VLA, DRIFT improves success by 8.13 percentage points over MLP while increasing latency from 28.4 to 31.4 ms (10.5%). Relative to direct flow, it improves success by 2.66 points for 0.3 ms additional latency. On TVG, DRIFT improves mIoU by 1.5 points while increasing latency from 603.3 to 605.8 ms (0.4%). The refiner uses 10 forward-Euler steps in these evaluations.

Table 6: **Prior quality and end-to-end inference latency.** (a) Libero-Long success rates; the random base predictor is untrained and has 0% success. (b) VLA reports Libero-Long success with Qwen3-VL-2B, and TVG reports Charades-STA mIoU with ET-Chat-3.8B. Latency on NVIDIA H200 GPUs is averaged over 1,000 VLA passes across 10 tasks and 510 TVG queries from 200 videos.

**(a) Prior quality**

| Action decoder | LL (%) |
|---|---|
| Flow | 94.2 |
| DRIFT (random base predictor) | 94.0 |
| DRIFT | **96.4** |

**(b) Inference latency**

| Task | Decoder | Perf. (%) | ms |
|---|---|---|---|
| VLA | MLP | $88.20 \pm 0.00$ | 28.4 |
| | Diffusion | $92.23 \pm 0.64$ | 33.5 |
| | Flow | $93.67 \pm 0.35$ | 31.1 |
| | DRIFT | $\mathbf{96.33 \pm 0.23}$ | 31.4 |
| TVG | ET-Chat | 42.3 | 603.3 |
| | ET-Chat + DRIFT | **43.8** | 605.8 |

### 4.4 Ablation Study

We ablate DRIFT on Libero-Long with Qwen3-VL-2B, using the training setup of Section 4.1.

**Decoder Structure and Prediction Variance.** We compare DRIFT against MLP regression, Flow Matching (Lipman et al., 2023), and Diffusion (Peebles & Xie, 2023) decoders. Table 5(a) summarizes the results, averaged across three runs. Both Flow Matching and Diffusion provide strong generative baselines for continuous decoding. Nevertheless, DRIFT consistently outperforms both.

**DRIFT Learns Improved Action Representations**. Table 5(b) shows the effects of jointly training the base predictor with the flow matching based refinement module. A standalone MLP baseline achieves 88.2% on Libero-Long, while the same MLP, when co-trained as the base predictor within DRIFT, improves to 93.0%. Incorporating the full DRIFT decoder further boosts performance to 96.4%. We hypothesize that joint training encourages the VLM to learn more informative action representations, thereby improving the performance of the base predictor itself.

**Effect of Prior Quality.** We test whether DRIFT still helps when the base predictor provides no useful prior. We omit its auxiliary loss $\mathcal{L}_p$ and stop gradients from the refinement module to the base predictor, leaving it untrained. Its success rate is 0% on Libero-Long. DRIFT then reaches 94.0% (Table 6(a)), close to direct flow matching (94.2%), whereas the trained-prior model reaches 96.4%.

**Model Design**. Finally, we compare design choices of DRIFT with a shorter 30K-step training schedule. We vary (1) the linear interpolant: $y_t = (1-t)\epsilon + t(y - \hat{y})$ (Direct Flow) vs. our design (DRIFT) $y_t = (1-t)(\epsilon + \hat{y}) + ty$ in Eq. 5, (2) velocity prediction ($v$-pred) vs. signal prediction ($x$-pred), (3) skip connection vs. no skip connection, and (4) with or without the base predictor loss $\mathcal{L}_p$. When without $\mathcal{L}_p$ loss, we allow the gradients from refinement module to back-propagate to the base predictor. Table 5(c) presents the design comparison. Among the valid configurations, the best result is our current design, $x$-prediction, skip-connection variant with $L_p$, which reaches 93.8%. Without the skip connection, the same interpolant with $x$-pred reaches 93.4%, still outperforming the Direct Flow alternatives at 91.2%. The effect of $\mathcal{L}_p$ is modest: with the skip connection, it changes 93.6% to 93.8%, while in the no-skip setting, the best score is obtained without it.

## 5 Conclusion

In this paper, we presented DRIFT, a general framework for adapting pretrained VLMs to continuous decoding tasks. DRIFT introduces a lightweight residual flow adapter that augments a base predictor with a flow-matching-based refinement module. It progressively refines an initial coarse prediction by modeling residual errors around a strong prior. This residual formulation simplifies generative modeling, improves optimization and learning efficiency, and enables effective continuous decoding across diverse tasks and architectures. Through extensive experiments on visual grounding and robotic control tasks spanning MLLMs, VLAs, and WAMs, we demonstrate consistent improvements over strong regression- and generative-based baselines. We hope our work sheds light on post-training strategies for extending VLMs beyond discrete token generation toward precise continuous outputs.

### AI USE STATEMENT

We did not use generative AI tools for any task that requires disclosure under the ICLR 2027 AI policy. We used them only to edit the paper for readability and length, adapt it to the ICLR 2027 format, and edit the presentation of the prior-quality and inference-latency results. We have reviewed all AI-assisted work, and no experimental result was generated by AI. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

## REFERENCES

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob L Menick, Sebastian Borgeaud, Andy Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikoł aj Bińkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karén Simonyan. Flamingo: a visual language model for few-shot learning. In *Advances in Neural Information Processing Systems*, volume 35, pp. 23716–23736. Curran Associates, Inc., 2022.

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, et al. Qwen3-vl technical report. *arXiv preprint arXiv:2511.21631*, 2025a.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025b.

Lucas Beyer, Andreas Steiner, André Susano Pinto, Alexander Kolesnikov, Xiao Wang, Daniel Salz, Maxim Neumann, Ibrahim Alabdulmohsin, Michael Tschannen, Emanuele Bugliarello, et al. Paligemma: A versatile 3b vlm for transfer. *arXiv preprint arXiv:2407.07726*, 2024.

Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Robert Equi, Chelsea Finn, Niccolo Fusai, Manuel Y. Galliker, Dibya Ghosh, Lachy Groom, Karol Hausman, brian ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Devin LeBlanc, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Allen Z. Ren, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz, James Tanner, Quan Vuong, Homer Walke, Anna Walling, Haohuan Wang, Lili Yu, and Ury Zhilinsky. $\pi_{0.5}$: a vision-language-action model with open-world generalization. In *9th Annual Conference on Robot Learning*, 2025a.

Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Lucy Xiaoyang Shi, James Tanner, Quan Vuong, Anna Walling, Haohuan Wang, and Ury Zhilinsky. $\pi_0$: A vision-language-action flow model for general robot control. In *RSS*, 2025b.

Xinyi Chen, Yilun Chen, Yanwei Fu, Ning Gao, Jiaya Jia, Weiyang Jin, Hao Li, Yao Mu, Jiangmiao Pang, Yu Qiao, et al. Internvla-m1: A spatially guided vision-language-action framework for generalist robot policy. *arXiv preprint arXiv:2510.13778*, 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*, 44(10-11):1684–1704, 2025.

Christopher Clark, Jieyu Zhang, Zixian Ma, Jae Sung Park, Mohammadreza Salehi, Rohun Tripathi, Sangho Lee, Zhongzheng Ren, Chris Dongjoo Kim, Yinuo Yang, et al. Molmo2: Open weights and data for vision-language models with video understanding and grounding. *arXiv preprint arXiv:2601.10611*, 2026.

Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, et al. Molmo and pixmo: Open weights and open data for state-of-the-art vision-language models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 91–104, 2025.

Chaoyou Fu, Yuhan Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, et al. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 24108–24118, 2025.

Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. TALL: Temporal activity localization via language query. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 5267–5275, 2017.

Lisa Anne Hendricks, Oliver Wang, Eli Shechtman, Josef Sivic, Trevor Darrell, and Bryan Russell. Localizing moments in video with natural language. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 33, pp. 6840–6851, 2020.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Liang Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *Iclr*, 1(2):3, 2022.

Bin Huang, Xin Wang, Hong Chen, Zihan Song, and Wenwu Zhu. VTimeLLM: Empower LLM to grasp video moments. In *CVPR*, 2024.

Siddharth Karamcheti, Suraj Nair, Ashwin Balakrishna, Percy Liang, Thomas Kollar, and Dorsa Sadigh. Prismatic vlms: Investigating the design space of visually-conditioned language models. In *International Conference on Machine Learning (ICML)*, 2024.

Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 787–798, 2014.

Dahye Kim, Jungin Park, Jiyoung Lee, Seongheon Park, and Kwanghoon Sohn. Language-free training for zero-shot video grounding. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 2539–2548, January 2023.

Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Foster, Grace Lam, Pannag Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. OpenVLA: An open-source vision-language-action model. In *CoRL*, 2024.

Ranjay Krishna, Kenji Hata, Frederic Ren, Li Fei-Fei, and Juan Carlos Niebles. Dense-captioning events in videos. In *Proceedings of the IEEE international conference on computer vision*, pp. 706–715, 2017.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 19730–19742. PMLR, 23–29 Jul 2023.

Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, et al. Mvbench: A comprehensive multi-modal video understanding benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22195–22206, 2024.

Tianhong Li and Kaiming He. Back to basics: Let denoising generative models denoise. *arXiv preprint arXiv:2511.13720*, 2025.

Xuanlin Li, Kyle Hsu, Jiayuan Gu, Oier Mees, Karl Pertsch, Homer Rich Walke, Chuyuan Fu, Ishikaa Lunawat, Isabel Sieh, Sean Kirmani, Sergey Levine, Jiajun Wu, Chelsea Finn, Hao Su, Quan Vuong, and Ted Xiao. Evaluating real-world robot manipulation policies in simulation. In *Proceedings of The 8th Conference on Robot Learning*, volume 270 of *Proceedings of Machine Learning Research*, pp. 3705–3728. PMLR, 06–09 Nov 2025.

Kevin Qinghong Lin, Pengchuan Zhang, Joya Chen, Shraman Pramanick, Difei Gao, Alex Jinpeng Wang, Rui Yan, and Mike Zheng Shou. UniVTG: Towards unified video-language temporal grounding. In *ICCV*, 2023.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023.

Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. Libero: Benchmarking knowledge transfer for lifelong robot learning. *Advances in Neural Information Processing Systems*, 36:44776–44791, 2023a.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023b.

Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations*, 2023c.

Xu Liu, Junfeng Hu, Yuan Li, Shizhe Diao, Yuxuan Liang, Bryan Hooi, and Roger Zimmermann. Unitime: A language-empowered unified model for cross-domain time series forecasting. In *Proceedings of the ACM Web Conference 2024*, WWW '24, pp. 4095–4106, New York, NY, USA, 2024a. Association for Computing Machinery. ISBN 9798400701719. doi: 10.1145/3589334.3645434.

Ye Liu, Zongyang Ma, Zhongang Qi, Yang Wu, Ying Shan, and Chang Wen Chen. E.t. bench: Towards open-ended event-level video-language understanding. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024b.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.

Ruipu Luo, Ziwang Zhao, Min Yang, Zheming Yang, Minghui Qiu, Zhongyu Wei, Yanhao Wang, and Cen Chen. Valley: Video assistant with large language model enhanced ability. *ACM Trans. Multimedia Comput. Commun. Appl.*, February 2026. ISSN 1551-6857. doi: 10.1145/3796716.

Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-ChatGPT: Towards detailed video understanding via large vision and language models. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.

Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L. Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.

Davide Moltisanti, Michael Wray, Walterio Mayol-Cuevas, and Dima Damen. Trespassing the boundaries: Labeling temporal bounds for object interactions in egocentric video. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.

Fangzhou Mu, Sicheng Mo, and Yin Li. Snag: Scalable and accurate video grounding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 18930–18940, 2024.

Jinwoo Nam, Daechul Ahn, Dongyeop Kang, Seong Jong Ha, and Jonghyun Choi. Zero-shot natural language video localization. In *ICCV*, 2021.

Andreea-Maria Oncescu, João F. Henriques, Yang Liu, Andrew Zisserman, and Samuel Albanie. QuerYD: A video dataset with high-quality text and audio narrations. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021.

Abby O'Neill, Abdul Rehman, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, et al. Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 6892–6903. IEEE, 2024.

William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.

Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, et al. Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*, 2024.

Shraman Pramanick, Effrosyni Mavroudi, Yale Song, Rama Chellappa, Lorenzo Torresani, and Triantafyllos Afouras. Enrich and detect: Video temporal grounding with multimodal LLMs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025.

Long Qian, Juncheng Li, Yu Wu, Yaobo Ye, Hao Fei, Tat-Seng Chua, Yueting Zhuang, and Siliang Tang. Momentor: Advancing video large language model with fine-grained temporal reasoning. In *ICML*, 2024.

Mengxue Qu, Xiaodong Chen, Wu Liu, Alicia Li, and Yao Zhao. Chatvtg: Video temporal grounding via chat with video dialogue large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pp. 1847–1856, June 2024.

Shuhuai Ren, Linli Yao, Shicheng Li, Xu Sun, and Lu Hou. TimeChat: A time-sensitive multimodal large language model for long video understanding. In *CVPR*, 2024.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, June 2022.

Gunnar A Sigurdsson, Gül Varol, Xiaolong Wang, Ali Farhadi, Ivan Laptev, and Abhinav Gupta. Hollywood in homes: Crowdsourcing data collection for activity understanding. In *European Conference on Computer Vision*, pp. 510–526. Springer, 2016.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations (ICLR)*, 2021.

Xiaolong Sun et al. Moment quantization for video temporal grounding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2025.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.

Alex Jinpeng Wang, Linjie Li, Kevin Qinghong Lin, Jianfeng Wang, Kevin Lin, Zhengyuan Yang, Lijuan Wang, and Mike Zheng Shou. COSMO: COntrastive streamlined MultimOdal model with interleaved pre-training. *arXiv preprint arXiv:2401.00849*, 2024a.

Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, et al. Internvl3. 5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. *arXiv preprint arXiv:2508.18265*, 2025.

Yi Wang, Yinan He, Yizhuo Li, Kunchang Li, Jiashuo Yu, Xin Ma, Xinhao Li, Guo Chen, Xinyuan Chen, Yaohui Wang, Conghui He, Ping Luo, Ziwei Liu, Yali Wang, Limin Wang, and Yu Qiao. InternVid: A large-scale video-text dataset for multimodal understanding and generation. In *International Conference on Learning Representations (ICLR)*, 2024b.

Yihao Wang, Pengxiang Ding, Lingxiao Li, Can Cui, Zirui Ge, Xinyang Tong, Wenxuan Song, Han Zhao, Wei Zhao, Pengxu Hou, et al. Vla-adapter: An effective paradigm for tiny-scale vision-language-action model. In *Proceedings of the AAAI conference on artificial intelligence*, volume 40, pp. 18638–18646, 2026.

Yueqian Wang, Xiaojun Meng, Jianxin Liang, Yuxuan Wang, Qun Liu, and Dongyan Zhao. Hawk-Eye: Training video-text LLMs for grounding text in videos. *arXiv:2403.10228*, 2024c.

Linhui Xiao, Xiaoshan Yang, Xiangyuan Lan, Yaowei Wang, and Changsheng Xu. Towards visual grounding: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv e-prints*, pp. arXiv–2412, 2024.

Ganlin Yang, Tianyi Zhang, Haoran Hao, Weiyun Wang, Yibin Liu, Dehui Wang, Guanzhou Chen, Zijian Cai, Junting Chen, Weijie Su, Wengang Zhou, Yu Qiao, Jifeng Dai, Jiangmiao Pang, Gen Luo, Wenhai Wang, Yao Mu, and Zhi Hou. Vlaser: Vision-language-action model with synergistic embodied reasoning. In *The Fourteenth International Conference on Learning Representations*, 2026.

Seonghyeon Ye, Yunhao Ge, Kaiyuan Zheng, Shenyuan Gao, Sihyun Yu, George Kurian, Suneel Indupuru, You Liang Tan, Chuning Zhu, Jiannan Xiang, et al. World action models are zero-shot policies. *arXiv preprint arXiv:2602.15922*, 2026.

Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *European conference on computer vision*, pp. 69–85. Springer, 2016.

Shoubin Yu, Jaemin Cho, Prateek Yadav, and Mohit Bansal. Self-chained image-language model for video localization and question answering. In *NeurIPS*, 2023.

Tianyuan Yuan, Zibin Dong, Yicheng Liu, and Hang Zhao. Fast-wam: Do world action models need test-time future imagination? *arXiv preprint arXiv:2603.16666*, 2026.

Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9556–9567, 2024.

Abhay Zala, Jaemin Cho, Satwik Kottur, Xilun Chen, Barlas Oğuz, Yashar Mehdad, and Mohit Bansal. Hierarchical video-moment retrieval and step-captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.

Hang Zhang, Xin Li, and Lidong Bing. Video-LLaMA: An instruction-tuned audio-visual language model for video understanding. In *EMNLP (Demo)*, 2023.

Jianke Zhang, Xiaoyu Chen, Yanjiang Guo, Yucheng Hu, and Jianyu Chen. VLM4VLA: Revisiting vision-language-models in vision-language-action models. In *The Fourteenth International Conference on Learning Representations*, 2026.

Jun Zhang, Teng Wang, Yuying Ge, Yixiao Ge, Xinhao Li, Ying Shan, and Limin Wang. Timelens: Rethinking video temporal grounding with multimodal llms. *arXiv preprint arXiv:2512.14698*, 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 A Statistical Perspective on Residual Refinement

## A.1 Notation and Setup

Let $z \in \mathcal{Z}$ denote the conditional representation produced by the MLLM encoder and $y \in \mathbb{R}^c$ the continuous target. A base predictor $g(z)$ outputs the point estimate $\hat{y} \triangleq g(z)$, and the residual is $r \triangleq y - g(z)$. Unless otherwise stated, the theoretical discussion assumes a continuous-valued base predictor $g(z) \in \mathbb{R}^c$. The token-based base predictor is treated as a practical variant evaluated empirically.

In DRIFT, we sample noise $\epsilon \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$ and construct the prior-shifted bridge

$$y_t = (1-t)\big(g(z) + \sigma\epsilon\big) + ty, \qquad t \in [0,1], \tag{6}$$

where $\sigma \in \mathbb{R}^c_+$ are noise scales. Let $X \triangleq (y_t, t, z)$ denote the input to the refinement network $h$. The refinement network learns the mapping $h : X \mapsto r$, and the final prediction is $\bar{y} = g(z) + h(X)$.

We define the *unweighted risk* and *weighted risk* of any measurable predictor $h$ as

$$\mathcal{R}(h) \triangleq \mathbb{E}\big[\|h(X) - r\|_2^2\big] = \mathbb{E}\big[\|g(z) + h(X) - y\|_2^2\big] = \mathbb{E}\big[\|\bar{y} - y\|_2^2\big], \quad \mathcal{L}(h) \triangleq \mathbb{E}\left[\frac{\|h(X) - r\|_2^2}{(1-t)^2}\right]. \tag{7}$$

The weighted form $\mathcal{L}(h)$ arises from the equivalence between velocity matching and signal prediction in flow matching (cf. Eqs. (2) and (3) in the main text). We write $w(t) \triangleq 1/(1-t)^2$.

For direct flow matching (direct FM), the bridge is $\tilde{y}_t = (1-t)\sigma\epsilon + ty$, the input is $\tilde{X} \triangleq (\tilde{y}_t, t, z)$, and a predictor $H : \tilde{X} \mapsto y$ is trained with weighted risk $\mathcal{L}_{\mathrm{FM}}(H) \triangleq \mathbb{E}\big[w(t)\|H(\tilde{X}) - y\|_2^2\big]$.

The *direct MLP* baseline corresponds to setting $h \equiv 0$, so that the final output is simply $g(z)$.

## A.2 Assumptions

**Definition 1** (Assumptions A)**.** *The following conditions are assumed throughout.*

- ***A1* Truncated time sampling.** *There exists $\tau \in (0,1)$ such that $t \in [0, 1-\tau]$, which ensures $1 \le w(t) \le \tau^{-2}$.*
- ***A2* Bounded targets and outputs.** *There exist constants $B_r, B_y, B_h, B_H > 0$ such that almost surely*

$$\|r\|_2 \le B_r, \qquad \|y\|_2 \le B_y, \qquad \|h(X)\|_2 \le B_h, \qquad \|H(\tilde{X})\|_2 \le B_H.$$

- ***A3* Finite complexity.** *The refinement function class $\mathcal{H}$ and the direct FM class $\mathcal{F}$ have finite empirical Rademacher complexities $\mathfrak{R}_n(\mathcal{H})$ and $\mathfrak{R}_n(\mathcal{F})$, respectively.*

**A1** avoids the weight divergence as $t \to 1$; in practice this is handled by truncated time sampling or numerical clipping. **A2** states that training targets and network outputs are bounded, which can equivalently be stated as sub-Gaussian tail bounds. **A3** ensures that empirical risk minimization yields nontrivial generalization bounds.

## A.3 Advantage over Direct MLP

We first introduce the *Bayes residual predictor*

$$m(X) \triangleq \mathbb{E}[r \mid X], \tag{8}$$

which is the mean-squared-error-optimal prediction of the residual given the bridge state. If $m(X) \not\equiv 0$, then the bridge state carries information about the residual beyond what $g(z)$ already captures, and the refinement network can exploit this to reduce risk.

**Lemma 2** (Orthogonal decomposition of squared loss)**.** *Let $m(X) = \mathbb{E}[r \mid X]$. Then for any measurable $h$,*

$$\mathcal{R}(h) = \mathcal{R}(m) + \mathbb{E}\big[\|h(X) - m(X)\|_2^2\big], \tag{9}$$

$$\mathcal{L}(h) = \mathcal{L}(m) + \mathbb{E}\big[w(t)\|h(X) - m(X)\|_2^2\big]. \tag{10}$$

*Proof.* We prove the unweighted case; the weighted case follows identically since $w(t)$ is $X$-measurable. Write

$$h(X) - r = \big(h(X) - m(X)\big) + \big(m(X) - r\big).$$

Expanding the squared norm:

$$\|h(X) - r\|_2^2 = \|h(X) - m(X)\|_2^2 + \|m(X) - r\|_2^2 + 2\langle h(X) - m(X),\, m(X) - r\rangle.$$

Taking expectations, it suffices to show the cross term vanishes. Since $h(X) - m(X)$ is $X$-measurable,

$$\mathbb{E}\big[\langle h(X) - m(X),\, m(X) - r\rangle \mid X\big] = \big\langle h(X) - m(X),\, \mathbb{E}[m(X) - r \mid X]\big\rangle = 0,$$

because $\mathbb{E}[m(X) - r \mid X] = m(X) - \mathbb{E}[r \mid X] = 0$ by definition. By the tower property, the unconditional cross term is also zero, giving $\mathcal{R}(h) = \mathcal{R}(m) + \mathbb{E}\,\|h(X) - m(X)\|_2^2$. □

**Theorem 3** (Population-level improvement over direct MLP)**.** *Direct MLP corresponds to the zero residual predictor $h \equiv 0$. Then*

$$\mathcal{R}(0) - \inf_h \mathcal{R}(h) = \mathbb{E}\big[\|m(X)\|_2^2\big], \tag{11}$$

$$\mathcal{L}(0) - \inf_h \mathcal{L}(h) = \mathbb{E}\big[w(t)\|m(X)\|_2^2\big]. \tag{12}$$

*In particular, whenever $m(X) \not\equiv 0$, we have $\inf_h \mathcal{R}(h) < \mathcal{R}(0)$, i.e., DRIFT strictly improves over direct MLP at the population level.*

*Proof.* By Theorem 2 with $h \equiv 0$:

$$\mathcal{R}(0) = \mathcal{R}(m) + \mathbb{E}\big[\|m(X)\|_2^2\big].$$

Since $\inf_h \mathcal{R}(h) = \mathcal{R}(m)$ (the Bayes predictor attains the infimum), we obtain Eq. (11). The weighted case is analogous: $\mathcal{L}(0) = \mathcal{L}(m) + \mathbb{E}[w(t)\|m(X)\|_2^2]$ and $\inf_h \mathcal{L}(h) = \mathcal{L}(m)$. □

**Interpretation.** The improvement gap $\mathbb{E}[\|m(X)\|_2^2]$ measures how much residual structure the bridge state can explain. Direct MLP stops at the coarse prediction $g(z)$; DRIFT further leverages the bridge state to refine the residual. As long as the bridge state carries any predictive information about the residual, DRIFT achieves strictly lower risk.

**Theorem 4** (Finite-sample sufficient condition for improvement over direct MLP)**.** *Under Assumptions A, let $S = \{(X_i, r_i, t_i)\}_{i=1}^n$ be an i.i.d. training sample. Define the empirical risk minimizer $\hat{h} \in \arg\min_{h\in\mathcal{H}} \hat{\mathcal{L}}(h)$ where $\hat{\mathcal{L}}(h) \triangleq \frac{1}{n}\sum_{i=1}^n w(t_i)\|h(X_i) - r_i\|_2^2$, and the best-in-class predictor $h^\star_{\mathcal{H}} \in \arg\min_{h\in\mathcal{H}} \mathcal{L}(h)$ with approximation error $\mathrm{App}_{\mathcal{H}} \triangleq \mathcal{L}(h^\star_{\mathcal{H}}) - \mathcal{L}(m) \geq 0$.*

*Then there exists a constant $C > 0$ such that, with probability at least $1 - \delta$,*

$$\mathcal{L}(\hat{h}) - \mathcal{L}(h^\star_{\mathcal{H}}) \leq C\tau^{-2}\left[(B_r + B_h)\mathfrak{R}_n(\mathcal{H}) + (B_r + B_h)^2\sqrt{\frac{\log(2/\delta)}{n}}\right]. \tag{13}$$

*Furthermore, since $w(t) \geq 1$,*

$$\mathcal{R}(0) - \mathcal{R}(\hat{h}) \geq \mathbb{E}\big[\|m(X)\|_2^2\big] - \mathrm{App}_{\mathcal{H}} - C\tau^{-2}\left[(B_r + B_h)\mathfrak{R}_n(\mathcal{H}) + (B_r + B_h)^2\sqrt{\frac{\log(2/\delta)}{n}}\right]. \tag{14}$$

*Therefore, whenever*

$$\mathbb{E}\big[\|m(X)\|_2^2\big] > \mathrm{App}_{\mathcal{H}} + C\tau^{-2}\left[(B_r + B_h)\mathfrak{R}_n(\mathcal{H}) + (B_r + B_h)^2\sqrt{\frac{\log(2/\delta)}{n}}\right], \tag{15}$$

*we have $\mathcal{R}(\hat{h}) < \mathcal{R}(0)$, i.e., DRIFT with finite-sample ERM strictly outperforms direct MLP.*

*Proof.* The proof proceeds in three steps.

**Step 1: Uniform deviation bound.** Define the per-sample loss $\ell_h(X, r, t) = w(t)\|h(X) - r\|_2^2$. By **A1** and **A2**, $1 \le w(t) \le \tau^{-2}$ and $\|h(X) - r\|_2 \le B_h + B_r$, so $0 \le \ell_h \le \tau^{-2}(B_h + B_r)^2$.

For fixed $(r, t)$, the map $u \mapsto \phi_{r,t}(u) = w(t)\|u - r\|_2^2$ is Lipschitz in $u$ with constant at most $2\tau^{-2}(B_h + B_r)$. By the vector-valued Rademacher contraction inequality,

$$\mathfrak{R}_n\big(\{\ell_h : h \in \mathcal{H}\}\big) \le C_1\tau^{-2}(B_h + B_r)\mathfrak{R}_n(\mathcal{H}),$$

where $C_1$ is an absolute constant. The standard Rademacher generalization bound then gives, with probability at least $1 - \delta$,

$$\sup_{h\in\mathcal{H}}\big|\mathcal{L}(h) - \hat{\mathcal{L}}(h)\big| \le C_2\tau^{-2}\left[(B_r + B_h)\mathfrak{R}_n(\mathcal{H}) + (B_r + B_h)^2\sqrt{\frac{\log(2/\delta)}{n}}\right]. \tag{16}$$

**Step 2: Excess risk of ERM.** Since $\hat{h}$ minimizes empirical risk, $\hat{\mathcal{L}}(\hat{h}) \le \hat{\mathcal{L}}(h^\star_{\mathcal{H}})$. Applying Eq. (16) to both $\hat{h}$ and $h^\star_{\mathcal{H}}$:

$$\mathcal{L}(\hat{h}) \le \hat{\mathcal{L}}(\hat{h}) + \Delta_n \le \hat{\mathcal{L}}(h^\star_{\mathcal{H}}) + \Delta_n \le \mathcal{L}(h^\star_{\mathcal{H}}) + 2\Delta_n,$$

where $\Delta_n$ denotes the right-hand side of Eq. (16). Absorbing the factor of 2 into the constant $C$ yields Eq. (13).

**Step 3: Converting to improvement over direct MLP.** By Theorem 2, $\mathcal{R}(\hat{h}) - \mathcal{R}(m) = \mathbb{E}\,\|\hat{h}(X) - m(X)\|_2^2$. Since $w(t) \ge 1$,

$$\mathbb{E}\big[\|\hat{h}(X) - m(X)\|_2^2\big] \le \mathbb{E}\big[w(t)\|\hat{h}(X) - m(X)\|_2^2\big] = \mathcal{L}(\hat{h}) - \mathcal{L}(m).$$

Decomposing the right-hand side:

$$\mathcal{L}(\hat{h}) - \mathcal{L}(m) = \big(\mathcal{L}(\hat{h}) - \mathcal{L}(h^\star_{\mathcal{H}})\big) + \big(\mathcal{L}(h^\star_{\mathcal{H}}) - \mathcal{L}(m)\big) = \big(\mathcal{L}(\hat{h}) - \mathcal{L}(h^\star_{\mathcal{H}})\big) + \mathrm{App}_{\mathcal{H}}.$$

Hence $\mathcal{R}(\hat{h}) - \mathcal{R}(m) \le \mathrm{App}_{\mathcal{H}} + (\text{RHS of Eq. (13)})$.

From Theorem 3, $\mathcal{R}(0) - \mathcal{R}(m) = \mathbb{E}[\|m(X)\|_2^2]$. Therefore

$$\mathcal{R}(0) - \mathcal{R}(\hat{h}) = \big(\mathcal{R}(0) - \mathcal{R}(m)\big) - \big(\mathcal{R}(\hat{h}) - \mathcal{R}(m)\big) \ge \mathbb{E}\big[\|m(X)\|_2^2\big] - \mathrm{App}_{\mathcal{H}} - (\text{RHS of Eq. (13)}),$$

which gives Eq. (14). When Eq. (15) holds, the right-hand side is strictly positive. □

### A.4 Comparison with Direct Flow Matching

We next compare DRIFT with direct full-target flow matching. The purpose of this subsection is not to claim that residualization changes the unrestricted population optimum. In fact, the two objectives are Bayes-equivalent after a deterministic shift of the bridge state. The advantage suggested by the teaser visualization is instead geometric and statistical: after the base predictor removes the predictable component of the target, the flow refiner learns a smaller residual correction and its bridge remains localized around the coarse anchor. This distinction is important when the target $y$ has already been normalized to a bounded range such as $[-1, 1]^c$; worst-case range bounds alone do not explain the advantage, but residual second moments and bridge radii still can.

**Proposition 5** (Bayes equivalence of DRIFT and direct FM)**.** *With unrestricted measurable function classes,*

$$\inf_h \mathcal{L}(h) = \inf_H \mathcal{L}_{\mathrm{FM}}(H). \tag{17}$$

*Proof.* By definition, the DRIFT bridge and the direct FM bridge satisfy

$$y_t = \tilde{y}_t + (1 - t)g(z),$$

so given $(t, z)$, the states $y_t$ and $\tilde{y}_t$ differ by a deterministic shift.

For any residual predictor $h$, define the full-target predictor $H_h(\tilde{X}) \triangleq h(X) + g(z)$, where $X = (\tilde{y}_t + (1-t)g(z), t, z)$. Then

$$H_h(\tilde{X}) - y = h(X) + g(z) - y = h(X) - r,$$

so $\|H_h(\tilde{X}) - y\|_2^2 = \|h(X) - r\|_2^2$. Multiplying by $w(t)$ and taking expectations gives $\mathcal{L}_{\mathrm{FM}}(H_h) = \mathcal{L}(h)$.

Conversely, for any full-target predictor $H$, define $h_H(X) \triangleq H(\tilde{X}) - g(z)$, where $\tilde{X} = (y_t - (1-t)g(z), t, z)$. Then

$$h_H(X) - r = H(\tilde{X}) - g(z) - (y - g(z)) = H(\tilde{X}) - y,$$

so $\mathcal{L}(h_H) = \mathcal{L}_{\mathrm{FM}}(H)$. This loss-preserving bijection implies Eq. (17). □

**Interpretation.** Since the unrestricted population optima coincide, DRIFT should not be presented as having a lower Bayes risk than direct FM. The relevant comparison is instead about the target that the learned generative component has to model. Direct FM learns the full target transport, whereas DRIFT assigns the predictable component to the base predictor and asks the flow network to learn the remaining residual transport.

**Theorem 6** (Lower-variance residual transport)**.** *Let*

$$\mu(z) \triangleq \mathbb{E}[y \mid z], \qquad r \triangleq y - g(z),$$

*and assume that $\epsilon$ is independent of $(y, z)$ with $\mathbb{E}[\epsilon] = 0$. Define the direct-FM target velocity and the residual-FM target velocity as*

$$v^\star_{\mathrm{dir}} \triangleq y - \sigma\epsilon, \qquad v^\star_{\mathrm{res}} \triangleq y - g(z) - \sigma\epsilon = r - \sigma\epsilon.$$

*Then the signal targets satisfy*

$$\mathbb{E}\,\|y\|_2^2 = \mathbb{E}[\mathrm{Tr}\ \mathrm{Cov}(y \mid z)] + \mathbb{E}\,\|\mu(z)\|_2^2, \tag{18}$$

$$\mathbb{E}\,\|r\|_2^2 = \mathbb{E}[\mathrm{Tr}\ \mathrm{Cov}(y \mid z)] + \mathbb{E}\,\|\mu(z) - g(z)\|_2^2. \tag{19}$$

*The corresponding velocity targets satisfy*

$$\mathbb{E}\,\|v^\star_{\mathrm{dir}}\|_2^2 = \mathbb{E}\,\|y\|_2^2 + \mathbb{E}\,\|\sigma\epsilon\|_2^2, \tag{20}$$

$$\mathbb{E}\,\|v^\star_{\mathrm{res}}\|_2^2 = \mathbb{E}\,\|r\|_2^2 + \mathbb{E}\,\|\sigma\epsilon\|_2^2. \tag{21}$$

*Consequently,*

$$\mathbb{E}\,\|v^\star_{\mathrm{dir}}\|_2^2 - \mathbb{E}\,\|v^\star_{\mathrm{res}}\|_2^2 = \mathbb{E}\,\|y\|_2^2 - \mathbb{E}\,\|r\|_2^2 = \mathbb{E}\,\|\mu(z)\|_2^2 - \mathbb{E}\,\|\mu(z) - g(z)\|_2^2. \tag{22}$$

*Thus, whenever the base predictor explains more conditional-mean energy than the zero predictor,*

$$\mathbb{E}\,\|\mu(z) - g(z)\|_2^2 < \mathbb{E}\,\|\mu(z)\|_2^2, \tag{23}$$

*the residual correction target and the residual velocity target have smaller second moment than their direct-FM counterparts. Moreover,*

$$\mathrm{Cov}(r \mid z) = \mathrm{Cov}(y \mid z), \tag{24}$$

*so residualization does not remove the irreducible conditional uncertainty of the target distribution; it removes only the condition-dependent mean structure that the base predictor already captures.*

*Proof.* We first prove the two signal-target identities. By the conditional second-moment decomposition,

$$\mathbb{E}\big[\|y\|_2^2 \mid z\big] = \mathrm{Tr}\ \mathrm{Cov}(y \mid z) + \|\,\mathbb{E}[y \mid z]\|_2^2 = \mathrm{Tr}\ \mathrm{Cov}(y \mid z) + \|\mu(z)\|_2^2.$$

Taking expectation over $z$ gives Eq. (18). For the residual target, since $g(z)$ is deterministic given $z$,

$$\mathbb{E}[r \mid z] = \mathbb{E}[y - g(z) \mid z] = \mu(z) - g(z), \qquad \mathrm{Cov}(r \mid z) = \mathrm{Cov}(y \mid z).$$

Applying the same conditional second-moment decomposition to $r$ gives

$$\mathbb{E}\big[\|r\|_2^2 \mid z\big] = \mathrm{Tr}\ \mathrm{Cov}(y \mid z) + \|\mu(z) - g(z)\|_2^2,$$

and taking expectation over $z$ gives Eq. (19).

For the velocity targets, expand the squared norms:

$$\mathbb{E}\,\|y-\sigma\epsilon\|_2^2 = \mathbb{E}\,\|y\|_2^2 + \mathbb{E}\,\|\sigma\epsilon\|_2^2 - 2\,\mathbb{E}\langle y, \sigma\epsilon\rangle,$$

and

$$\mathbb{E}\,\|r-\sigma\epsilon\|_2^2 = \mathbb{E}\,\|r\|_2^2 + \mathbb{E}\,\|\sigma\epsilon\|_2^2 - 2\,\mathbb{E}\langle r, \sigma\epsilon\rangle.$$

Because $\epsilon$ is independent of $(y, z)$ and has zero mean, both cross terms vanish. This proves Eqs. (20) and (21). Subtracting the residual identity from the direct identity yields Eq. (22). The last claim follows from the fact that subtracting a deterministic function of $z$ does not change the conditional covariance given $z$. □

**Normalized-output implication.** The comparison above does not require $\|r\|_2$ to be uniformly smaller than $\|y\|_2$. Even when $y$ is normalized to $[-1,1]^c$, the direct decoder still has to learn the condition-dependent mean transport $\mu(z)$, while the residual decoder only has to learn $\mu(z)-g(z)$. If $g(z)$ is a useful base predictor, then $\mathbb{E}\,\|\mu(z)-g(z)\|_2^2$ can be much smaller than $\mathbb{E}\,\|\mu(z)\|_2^2$ despite both targets living in bounded ranges.

**Theorem 7** (Localization of the residual bridge)**.** *For a fixed $t \in [0,1]$, define the residual-bridge radius around the coarse anchor and the direct-FM bridge radius around the noise origin as*

$$\rho_{\mathrm{res}}(t) \triangleq \mathbb{E}\,\|y_t - g(z)\|_2^2, \qquad \rho_{\mathrm{dir}}(t) \triangleq \mathbb{E}\,\|\tilde{y}_t\|_2^2.$$

*Then*

$$\rho_{\mathrm{res}}(t) = (1-t)^2\,\mathbb{E}\,\|\sigma\epsilon\|_2^2 + t^2\,\mathbb{E}\,\|r\|_2^2, \tag{25}$$

$$\rho_{\mathrm{dir}}(t) = (1-t)^2\,\mathbb{E}\,\|\sigma\epsilon\|_2^2 + t^2\,\mathbb{E}\,\|y\|_2^2. \tag{26}$$

*Therefore, if $\mathbb{E}\,\|r\|_2^2 \le \eta\,\mathbb{E}\,\|y\|_2^2$ for some $\eta < 1$, then*

$$\rho_{\mathrm{res}}(t) \le \rho_{\mathrm{dir}}(t) - t^2(1-\eta)\,\mathbb{E}\,\|y\|_2^2. \tag{27}$$

*Moreover, for any radius $a > 0$,*

$$\mathbb{P}\big(\|y_t - g(z)\|_2 > a\big) \le \frac{(1-t)^2\,\mathbb{E}\,\|\sigma\epsilon\|_2^2 + t^2\,\mathbb{E}\,\|r\|_2^2}{a^2}. \tag{28}$$

*Proof.* Using the DRIFT bridge,

$$y_t - g(z) = (1-t)\sigma\epsilon + t(y-g(z)) = (1-t)\sigma\epsilon + tr.$$

Expanding the squared norm and using the independence and zero mean of $\epsilon$ gives

$$\mathbb{E}\,\|y_t - g(z)\|_2^2 = (1-t)^2\,\mathbb{E}\,\|\sigma\epsilon\|_2^2 + t^2\,\mathbb{E}\,\|r\|_2^2,$$

which proves Eq. (25). Similarly, the direct FM bridge satisfies

$$\tilde{y}_t = (1-t)\sigma\epsilon + ty,$$

so the same expansion gives Eq. (26). Substituting $\mathbb{E}\,\|r\|_2^2 \le \eta\,\mathbb{E}\,\|y\|_2^2$ into Eq. (25) and comparing with Eq. (26) gives Eq. (27). Finally, Eq. (28) follows from Markov's inequality applied to the nonnegative random variable $\|y_t - g(z)\|_2^2$. □

**Teaser interpretation.** Theorem 7 formalizes the qualitative picture in which residual refinement stays in a smaller tube around the coarse anchor, while direct generative decoding must transport noise toward the full target region. The statement is intentionally about the *correction* and the *bridge geometry*, not about artificially reducing the final conditional output diversity. At the Bayes optimum, the final conditional distribution should still match the true target distribution; DRIFT changes which part of the mapping is learned by the deterministic coarse prediction and which part is left to the generative refiner.

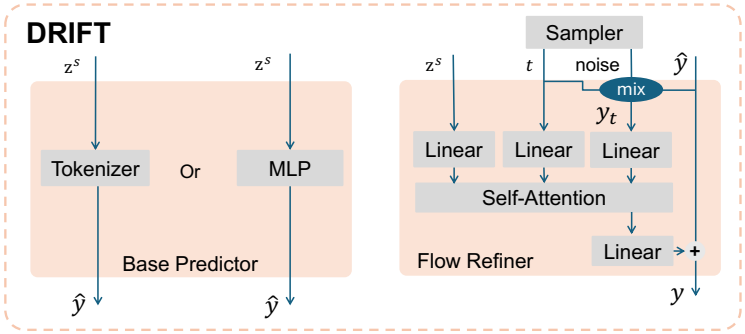


Figure 4: The overview of the DRIFT's structure.

# B Implementation Details

## B.1 DRIFT Architecture

Figure 4 illustrates the architecture of DRIFT. DRIFT consists of two main components: a Base Predictor and a Flow Refiner. The Base Predictor first produces a coarse continuous estimate $\hat{y}$ from the MLLM representation. The Flow Refiner then iteratively predicts the remaining residual from $\hat{y}$ to the ground-truth target.

The Flow Refiner contains three components: embedders, self-attention layers, and an output head. The embedders map the condition representation, the flow time $t$, and the bridge state $y_t$ into a shared hidden space. The self-attention layers aggregate information across these embedded tokens. Finally, the output head maps the updated representation to the flow prediction, which is combined with the $\hat{y}$ through the skip connection to produce the final prediction.

## B.2 VLA Implementation Details

**Training Data.** We follow the InternVLA-M1 (Chen et al., 2025) to prepare the training data. For the experiments on Libero (Liu et al., 2023a), we train on the Spatial, Object, Goal, and Long suites. Spatial, Object, Goal, and Long contain 432, 454, 428, and 379 training trajectories, respectively, and yield 53K, 67K, 52K, and 101K training samples, respectively. We normalize the action trajectories with the max and min values.

For the experiments on Simpler (Li et al., 2025) WidowX, we follow the setting used in (Chen et al., 2025) and train the model on the Bridge and Fractal splits of Open X-Embodiment (OXE) (O'Neill et al., 2024). Bridge and Fractal contain 60K and 87K training trajectories, respectively, and yield 1.28M and 3.4M training samples, respectively. We normalize the action trajectories with the 99th percentile value for the trajectories.

**Base VLMs.** We use Qwen3-VL-2B-Instruct (Bai et al., 2025a) and OpenVLA (Kim et al., 2024) as the base models. For OpenVLA, when training and evaluating on Simpler, we use the checkpoint pretrained on OXE. When training and evaluating on each Libero suite, we use the corresponding Libero-finetuned OpenVLA checkpoint. During DRIFT adaptation, we add LoRA (Hu et al., 2022) adapters to both the vision encoder and the LLM backbone. We use LoRA rank 32, $\alpha$= 64 for Libero experiments and rank 128, $\alpha$= 256 for Simpler experiments.

**Special Tokens.** We add one additional special token. For Qwen3-VL-2B-Instruct, we use the special token `<think>`, which already exists in the tokenizer but is not used during pre-training. For OpenVLA, we add a special token by extending the tokenizer, and we use both the added special token and the action tokens in the OpenVLA output as special tokens. We extract the hidden state of the special token from the last LLM layer as the condition $z^s$.

**Base Predictor.** When applying DRIFT to Qwen3-VL, we train a learnable MLP as the Base Predictor, which directly predicts a continuous action $\hat{y}$. We optimize this Base Predictor with MSE

loss. When applying DRIFT to OpenVLA, we use OpenVLA's tokenizer as the Base Predictor. The tokenizer discretizes each continuous value into 256 bins represented by 256 action tokens. During training, we sample the action from the multinomial distribution predicted by the model and use it as $\hat{y}$. We optimize the tokenizer-based Base Predictor with negative log-likelihood loss.

**Flow Refiner.** The Flow Refiner uses 4 self-attention layers with 8 heads and a hidden dimension of 768. We optimize the Flow Refiner with MSE loss between the target velocity $v^*$ and the predicted velocity $v_\theta(y(t), t, z)$.

**Training Hyper-parameters.** We leverage both the main view and the wrist view to train DRIFT with Qwen3-VL-2B (Bai et al., 2025a), and DRIFT learns to predict the next 8 and 4 consecutive actions in the future for the Libero and Simpler settings, respectively. We only use the main view when training DRIFT with OpenVLA (Kim et al., 2024), and DRIFT predicts one future action to match the original setting of OpenVLA.

We train DRIFT with batch size 16 and learning rate $4 \times 10^{-5}$ for 30K steps on Spatial, Object, and Goal, and for 60K steps on Long. For Simpler WidowX, we train with batch size 128, learning rate $4 \times 10^{-5}$, and 50K steps. We use cosine learning-rate decay with 10% warmup steps. All the experiments are trained with AdamW (Loshchilov & Hutter, 2019) optimizer on 8*H100 in 8-16 hours.

**Inference and Evaluation.** We use $K = 10$ forward-Euler steps and one stochastic noise sample per query. We evaluate DRIFT on Libero-Spatial, Libero-Object, Libero-Goal, Libero-Long, and Simpler WidowX. On Libero, we run 50 trials for each of the 10 subtasks in every suite. On Simpler WidowX, we run 24 trials for each of 4 subtasks. We report the success rate as the metric.

### B.3 TVG Implementation Details

We provide the implementation details for DRIFT on Temporal Video Grounding (TVG). We evaluate DRIFT under the ED-VTG protocol with ET-Chat as the MLLM backbone.

**Training Data.** We train the TVG variant of DRIFT on TimeLens-100K (Zhang et al., 2025). TimeLens-100K contains approximately 20K source videos and 100K query–segment temporal grounding annotations. Its source videos are collected from a mixture of video-language datasets, including CosMo-Cap (Wang et al., 2024a), InternVid-VTime (Wang et al., 2024b), DiDeMo (Hendricks et al., 2017), QuerYD (Oncescu et al., 2021), and HiREST (Zala et al., 2023). The temporal-grounding annotations are MLLM-generated end-to-end: an MLLM first identifies temporally distributed events in each video and then produces grounding queries with timestamp pairs. We use TimeLens-100K only as the training corpus and evaluate on the ED-VTG benchmarks, Charades-STA (Gao et al., 2017) and ActivityNet-Captions (Krishna et al., 2017). The TimeLens-100K training annotations are distinct from the evaluation annotations used in these benchmarks.

**Base VLMs.** We use ET-Chat (Liu et al., 2024b) as the backbone. We freeze the ET-Chat visual encoder and frame compressor, apply LoRA to the LLM attention projections, and train the LoRA adapters jointly with the Flow Refiner. We add three temporal special tokens, `<CT>`, `<VALUE>`, and `</CT>`, and train the LLM with its standard next-token objective to emit these markers. We use LoRA rank 32, $\alpha = 64$, and dropout $0.05$.

**Input and Condition Representation.** We follow the ET-Chat video pipeline: videos are sampled at 1 FPS, encoded with EVA-CLIP ViT-G/14, compressed by ET-Chat's frame compressor, and then fed into Phi-3-Mini-4K-Instruct (3.8B) as the LLM backbone. The frame compressor consists of a Q-Former with learnable queries, a context aggregator, and a projector. For temporal grounding, the target text marks each continuous boundary with `<CT><VALUE></CT>`. We use the LLM hidden states at the `<VALUE>` positions as condition $z^s$.

**Base Predictor.** The coarse prediction is produced by a small MLP head over the `<VALUE>` hidden states. The prior head predicts the temporal segment in center-width coordinates, applies a $1.2\times$ width expansion factor, and converts the result to start-end coordinates.

**Flow Refiner.** The Flow Refiner uses 2 self-attention layers with GELU activations, nhead $= 16$, and hidden dimension equal to the LLM's hidden_size$/2$. It represents each temporal segment in center-width coordinates and predicts the velocity in residual center-width coordinates.

**Training Hyper-parameters.** We train DRIFT on TimeLens-100K with per-device batch size 1, learning rate $5 \times 10^{-4}$ for 4 epochs with AdamW optimization on 8 A100 GPUs. The training run takes approximately 10 hours. The training objective uses flow MSE on the velocity.

**Inference and Evaluation.** The start and end timestamps are normalized to $[-1, 1]$ by the video duration during training. At inference time, we unnormalize the prediction back to seconds and clip the final $(t_s, t_e)$ to the valid range $[0, \text{duration}]$. We use $K = 10$ forward-Euler steps. Integration is performed in center-width coordinates; the result is converted back to start-end coordinates and then clipped to the valid video-time range.

ET-Chat (Liu et al., 2024b) is instruction-tuned on E.T. Instruct 164K. According to the ET-Chat paper, its TVG samples are drawn from DiDeMo, TACoS, QueryD, and NaQ. ET-Chat does not report ActivityNet-Captions performance in the original paper, so we obtain the ET-Chat baseline in Table 2 by running the released checkpoint under the ED-VTG evaluation protocol.

We report Recall@1 at different IoU thresholds and mean IoU. Recall@1 at IoU threshold $\tau$ measures the fraction of queries whose top-1 predicted interval has IoU greater than $\tau$ with the ground-truth interval. mIoU is the mean IoU over all evaluated queries.

### B.4 Spatial Grounding Implementation Details

**Training Data.** We use the training set of the RefCOCO (Kazemzadeh et al., 2014), RefCOCO+ (Mao et al., 2016), and RefCOCOg (Yu et al., 2016), which contains 120K, 120K, and 80K training samples, respectively. All the bounding boxes' coordinates are normalized between 0 and 1, and with the bounding box format of [cx,cy,w,h].

**Base VLM.** We use Qwen3-VL-2B-Instruct as the base model. We add the LoRA (Hu et al., 2022) adaptor with rank 32, $\alpha$= 64, on both the vision encoder and the LLM backbone.

**Special Tokens.** We add one additional special token ¡think¿, which already exists in the tokenizer but was never used during training in the Instruct version of the model. We extract the hidden state of the special token from the last layer of the LLM as condition $c$.

**Base Predictor.** We train a learnable MLP as the base predictor, which directly predicts $\hat{y}$. We apply MSE loss to update the base predictor.

**Flow Refiner.** We use 8-layer self-attention, with 8 heads and hidden dimension of 768 in the flow refiner. We optimize the Flow Refiner with MSE loss between the target velocity $v^*$ and the predicted velocity $v_\theta(y(t), t, z)$.

**Training Hyper-parameters.** We train DRIFT with batch size 64 and learning rate $4 \times 10^{-5}$ for 1 epoch (about 5K steps). We use cosine-decay learning rate scheduling and with 3% warmup steps. All the experiments are trained with AdamW optimizer on 8*A100 in 2 hours.

**Inference and evaluation** $K = 10$ forward-Euler steps; one stochastic noise sample per query (deterministic). Following the setting used in the Qwen3-VL, we evaluate the validation set of DRIFT on RefCOCO, RefCOCO+, and RefCOCOg, and report the Accuracy@1 as the metric.

### B.5 FastWAM VLA Implementation Details

**Training Data.** Following the training setting of the FastWAM (Yuan et al., 2026), we use all the splits (Spatial, Object, Goal, and Long) in Libero to train the model. In total, it has 272K training samples.

**Base VLM.** We use the same architecture for the video generator (VideoDiT) and the action generator (ActionDiT) as the FastWAM. We follow FastWAM's implementation to initialize the model before the training.

**Base Predictor.** We use a 2-layer Qformer, with 16 learnable queries, 8 heads, hidden dimension 1024, to aggregate the information from the embedded text embedding and the vision token encoded by the VAE (Wan et al., 2025), and then we apply a mean operation on the query tokens from the Qformer. We send the averaged query token to MLP to predict the $\hat{y}$.

**Flow Refiner.** We use the ActionDiT in the FastWAM structure as the Flow Refiner, and we optimize the Flow Refiner with MSE loss between the target velocity $v^*$ and the predicted velocity $v_\theta(y(t), t, z)$.

**Training, Inference and Evaluation.** We use the same training and evaluation setting as FastWAM. We evaluate DRIFT on Libero-Spatial, Libero-Object, Libero-Goal, and Libero-Long. On Libero, we run 50 trials for each of the 10 subtasks in every suite. We report the success rate as the metric. All the experiments are trained on 8*H200 in 24 hours.

## C Additional Visualization

Figure 5 presents additional visualizations on TVG, comparing the start and end timestamps predicted by different modules. Samples 1–4 show that DRIFT achieves more accurate temporal boundaries by iteratively refining the coarse prediction using the flow refiner. Samples 5–6 illustrate error patterns shared by all decoders. These errors are mainly caused by ambiguous or inaccurate annotations: in Sample 5, the action *takes a bite of the sandwich* lasts only about two seconds, but the ground-truth boundary annotates it as a five-second interval; in Sample 6, the query *sits on a bed* is ambiguous, as it can refer either to the action of sitting down or to the state of already being seated. These cases suggest that further improving TVG performance may require both more carefully curated training and validation annotations, as well as models that are better able to handle ambiguous natural-language queries.

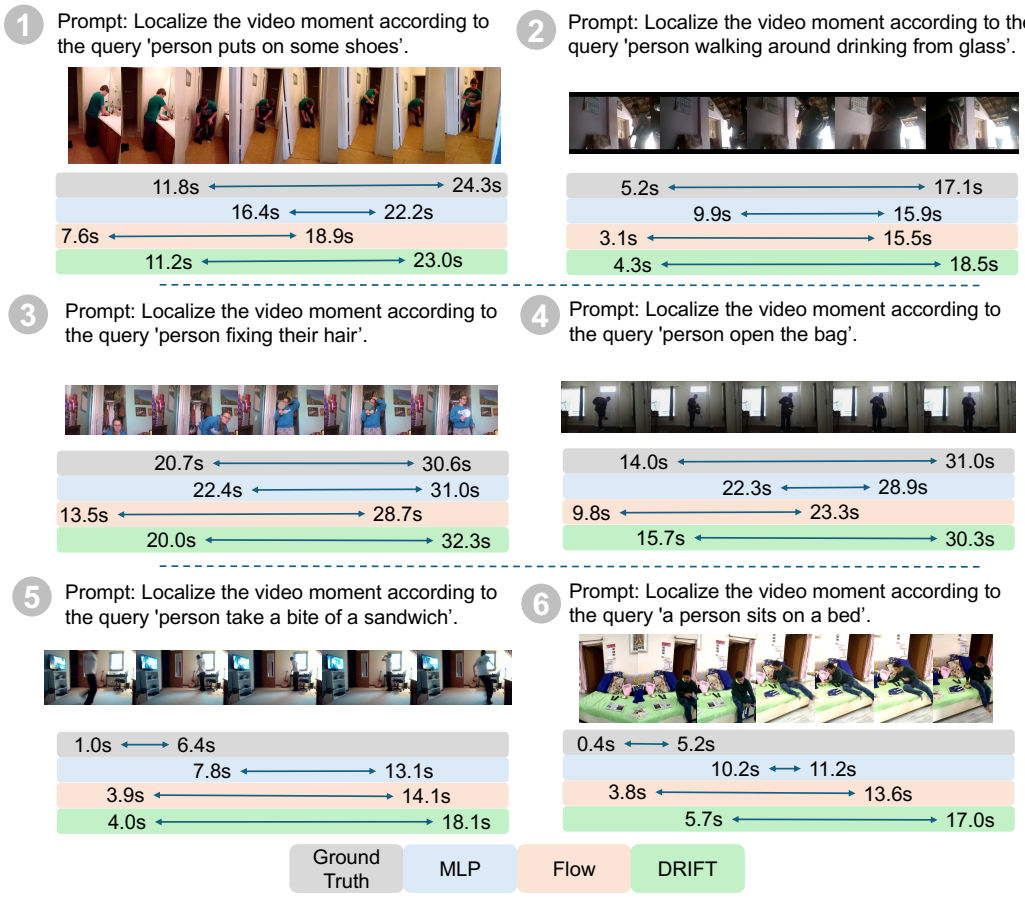


Figure 5: Additional visualization on TVG.