

FD-VLA: Force-Distilled Vision-Language-Action Model for Contact-Rich Manipulation

Ruiteng Zhao¹, Wenshuo Wang¹, Yicheng Ma², Xiaocong Li³, Francis E.H. Tay⁴,
Marcelo H. Ang Jr.⁴ and Haiyue Zhu^{5†}

Abstract—Force sensing is a crucial modality for Vision-Language-Action (VLA) frameworks, as it enables fine-grained perception and dexterous manipulation in contact-rich tasks. We present Force-Distilled VLA (FD-VLA), a novel framework that integrates force awareness into contact-rich manipulation without relying on physical force sensors. The core of our approach is a Force Distillation Module (FDM), which distills force by mapping a learnable query token, conditioned on visual observations and robot states, into a predicted force token aligned with the latent representation of actual force signals. During inference, this distilled force token is injected into the pretrained VLM, enabling force-aware reasoning while preserving the integrity of its vision-language semantics. This design provides two key benefits: first, it allows practical deployment across a wide range of robots that lack expensive or fragile force-torque sensors, thereby reducing hardware cost and complexity; second, the FDM introduces an additional force-vision-state fusion prior to the VLM, which improves cross-modal alignment and enhances perception-action robustness in contact-rich scenarios. Surprisingly, our physical experiments show that the distilled force token outperforms direct sensor force measurements as well as other baselines, which highlights the effectiveness of this force-distilled VLA approach.

I. INTRODUCTION

Vision-Language-Action (VLA) architectures [1]–[5] have emerged as a dominant paradigm in robot imitation learning, enabling scalable manipulation skills for a wide spectrum of everyday and industrial tasks. Building on large pretrained Vision-Language Models (VLMs) [6], [7], modern VLAs integrate perception, reasoning, planning, and control into a unified end-to-end framework. This allows them to map RGB inputs and natural-language instructions directly to low-level robot commands, while benefiting from strong

This research is supported by National Robotics Programme (NRP) 2.0 funding initiative "Domain-specific Robotics Foundation Models for Manufacturing (DS-RFM)".

¹Ruiteng Zhao and Wenshuo Wang are with the Advanced Robotics Centre, National University of Singapore, and are also attached students of SIMTech, A*STAR. (email: {ruiteng, wenshuo_wang}@u.nus.edu)

²Yicheng Ma is with School of Electrical & Electronic Engineering, Nanyang Technological University (NTU) and is also an attached student of SIMTech, A*STAR. (email: yichatma@gmail.com)

³Xiaocong Li is with the College of Information Science and Technology, Eastern Institute of Technology, Ningbo, Ningbo 315200, China, and also with the John A. Paulson School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02134 USA. (email: xiaocongli@eitech.edu.cn, xiaocongli@seas.harvard.edu).

⁴Francis E.H. Tay and Marcelo H. Ang Jr are with Advanced Robotics Centre at National University of Singapore, Singapore 117608, (email: mpetayeh@nus.edu.sg, mpeangh@nus.edu.sg)

⁵Haiyue Zhu is with the Singapore Institute of Manufacturing Technology, Agency for Science, Technology and Research (A*STAR), Singapore 138634, (email: zhu_haiyue@simtech.a-star.edu.sg)

† Corresponding author: zhu_haiyue@simtech.a-star.edu.sg

semantic priors and broad generalization capabilities. Yet, real-world deployments reveal fundamental limits of vision-only policies: occlusions, lighting shifts, depth ambiguities, and subtle contact effects often remain poorly captured in images. This underscores the need for truly multimodal VLAs that combine additional sensory signals with visual and linguistic context.

Among the possible modalities, force/tactile sensing stands out as a particularly critical but still underexplored dimension for VLAs, as it provides direct information about contact dynamics, compliance, and physical interactions that are often invisible to vision [8], [9]. Recent studies [10]–[13] have begun to incorporate force/tactile signals into VLAs, which raises a key question of how to efficiently incorporate this additional modality within the VLA framework. Fig. 1 overviews the different architecture of those VLAs with force/tactile modality. Tactile-VLA encodes raw force signals into compact latent embeddings that are projected and fused with pretrained vision-language features via cross-modal alignment layers, which enriches visuomotor reasoning with contact awareness. However, it risks disturbing the pretrained VLM's carefully aligned vision-language semantics through modality misalignment and catastrophic interference. Instead, ForceVLA encodes force signals into a dedicated token that is injected post-VLM and fused through a force-aware Mixture-of-Experts (MoE) module, thereby enriching the policy with contact-sensitive cues while its late fusion preserves the integrity of the pretrained vision-language representations. Though effective, its MoE design potentially increases the risk of training instability as well as inference complexity. Moreover, the late-fusion mechanism limits fine-grained force-vision-state interactions, reducing tight perception-action coupling and generalization.

In this work, we propose a novel FD-VLA framework that incorporates a distilled force token, rather than raw sensor signals, into the VLA model to enhance contact-rich manipulation. The distillation is achieved via a Force Distillation Module (FDM) that maps a learnable query token, conditioned on visual observations and robot states through attention, into a latent predicted force token. During the training, the supervision of FDM is achieved by aligning the predicted force token with the latent representation of actual force signals. At inference time, this distilled force token can then be injected into the VLM, enabling force-aware reasoning without requiring direct sensor measurements. This enables practical deployment on a wide range of robot platforms that lack force sensors, reducing hardware cost and

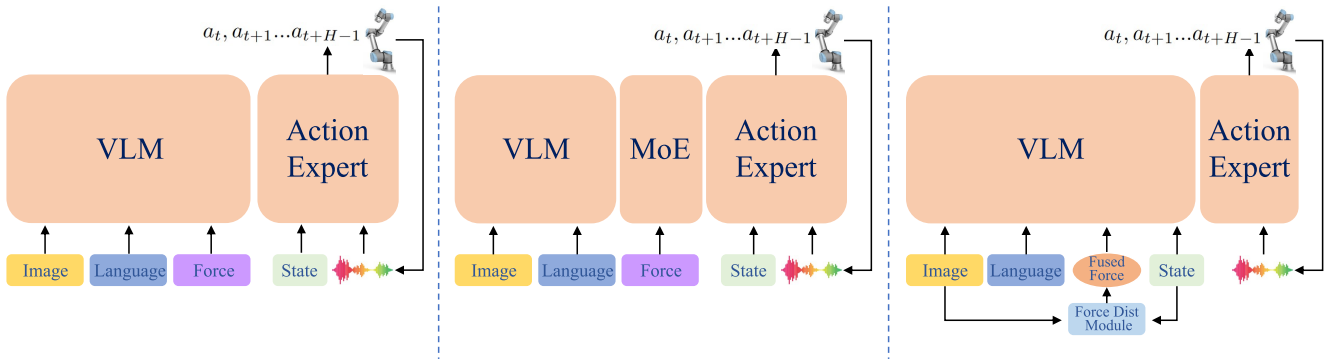


Fig. 1: Overview of differentiate architectures of force VLAs. **(Left)** Tactile-VLA with tactile encoder directly encode tactile information. **(Middle)** Force-VLA with MoE module between VLM and action expert. **(Right)** Our FD-VLA using predicted force instead of actual force to get better fused feature and sensor free at deployment.

complexity while still benefiting from force-aware reasoning. More important, FDM provides an additional force-vision-state fusion in the force distillation before VLM, which improves the consistency of cross-modal alignment and enhances the perception-action robustness in contact-rich manipulation.

In summary, the main contributions of this work are summarized as follows:

- We propose the FD-VLA framework that injects a distilled force token into the VLA model to improve contact-rich manipulation.
- The design of FDM is introduced for VLA to distill predicted force token by conditioning learnable query token on vision and state inputs.
- Our approach enables force-aware reasoning without requiring physical sensors, while provides additional force-vision-state cross-modal alignment.

II. RELATED WORKS

A. Vision-Language-Action Models

Compared to traditional robot grasp learning [14]–[16], The emergence of large-scale VLA models has redefined the landscape of robot learning by coupling pretrained vision-language representations with action generation. Pioneering systems [1]–[4], [17]–[21] built on large-scale datasets [22]–[24] demonstrated the potential of aligning visual-linguistic embeddings with robotic policies, which achieves strong generalization across diverse tasks and platforms. These models typically rely on large pretrained backbones, enabling them to parse natural language instructions and interpret complex visual scenes. Recent studies have moved beyond the conventional vision-language setting, introducing additional modalities such as audio information [25] and depth information [26], [27] to provide richer context for robotic perception and action generation. However, for contact-rich manipulation tasks, these additional modalities are less effective than force or tactile sensing, which provide direct measurements of physical interactions.

B. Force Integration in Contact-Rich Manipulation

Contact-rich manipulation tasks require precise modeling of physical interactions such as contact forces, slips, and deformations, which are difficult to infer from visual observations alone. Traditional vision-only approaches often struggle with unstable dynamics and lack the fine-grained feedback necessary for reliable control [28], [29]. To address these challenges, recent works have begun to integrate force and tactile sensing into robotic policies [30]–[34], demonstrating significant improvements in stability and precision. VLAs inherit the same fundamental limitation as traditional vision-centric policies: they lack direct access to physical interaction cues. In order to solve VLA’s limitation, recent research has begun to incorporate tactile and force sensing into the VLA paradigm. Examples include FuSe, ForceVLA, and Tactile-VLA [10], [11], [35]. These models explore auxiliary loss designs, modality-specific routing to capture embodied dynamics more faithfully or using prior knowledge and tactile sensors to achieve zero-shot generalization. Unlike current works that extend VLAs through heavy finetuning, architectural modifications or relying on special sensors, we explore enhancing model’s ability in contact-rich tasks by using images and robot state to generate force information and sensor-free at deployment.

C. Cross-Modal Knowledge Distillation

Transferring information across modalities is essential for bridging semantic perception and physical interaction in robotic systems. Cross-modal knowledge distillation provides a principled way to align heterogeneous sensory streams, allowing one modality to benefit from the complementary strengths of another. Earlier studies [36], [37] explored feature-level transfer through carefully designed objectives, laying the foundation for cross-modal knowledge distillation. Subsequent works [38] injected depth priors to enhance the visual representation for 3D perception tasks. And [39] introduced an adaptive transfer approach to achieve more robust and effective cross-modal knowledge distillation. Currently, little attention has been given to contact-rich robotic manipulation, where force or tactile sensing remains expensive and sparsely

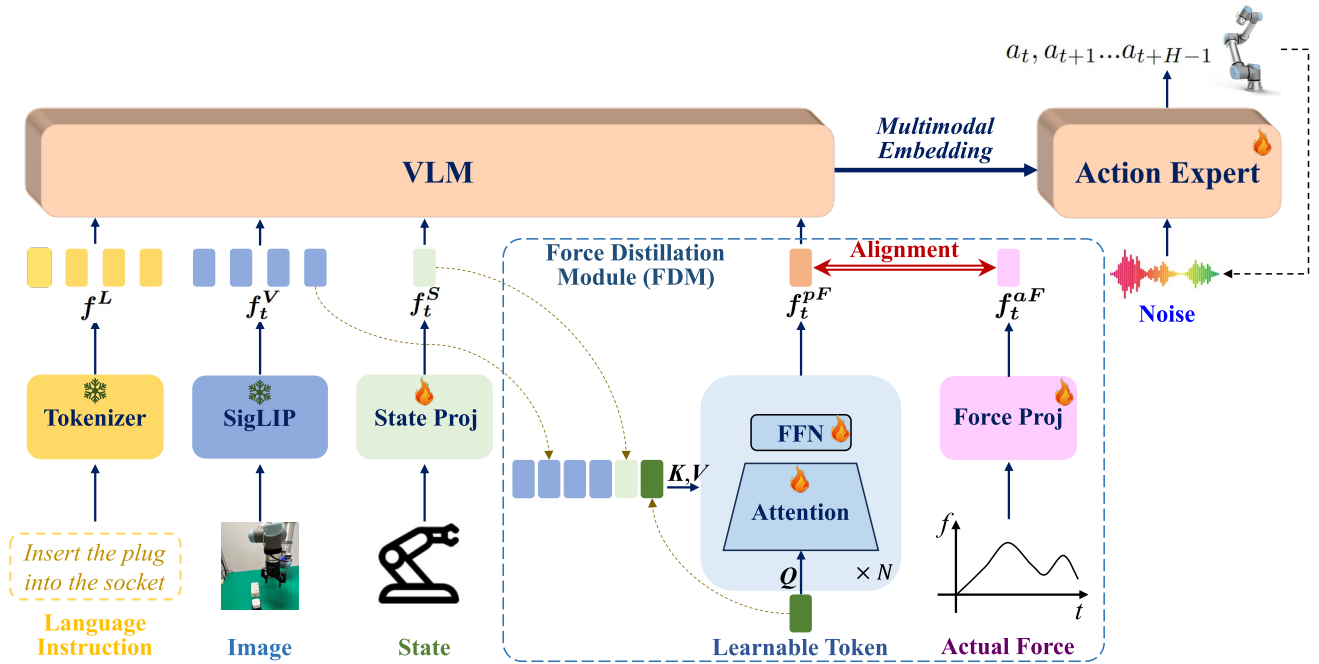


Fig. 2: Overview of our framework. During training, measured force signals are encoded into an actual force token via a lightweight projection. A learnable query attends to image and state tokens to predict a latent force token, which is supervised against the actual force token. At inference, the model synthesizes this latent force representations solely from vision and state inputs, eliminating the need for tactile hardware. The predicted force token is then fused with language, image, and state tokens inside the pretrained VLM, and an action expert consumes the fused representation to generate action sequences.

available. Our work tackles this challenge by distilling implicit force representations from vision and state inputs, eliminating the need for dedicated tactile hardware at deployment.

III. METHODOLOGY

This section presents the FD-VLA framework with detail methodologies.

A. Motivation

The objective of this work is to explore an effective and data-efficient approach to incorporate additional force modality into VLA model for contact-rich manipulation. Directly retraining large vision-language models to accommodate an additional modality is prohibitively expensive and risks degrading the well-aligned visual-linguistic representations. We seek a design that integrates force without modifying the pretrained VLM, thereby preserving its semantic alignment while avoiding costly retraining. However, the challenge is how to achieve good cross-modal fusion and alignment in the meantime.

Moreover, we also observed that raw force signals collected in real-world settings, as shown in Fig. 3, may not be optimal for policy learning, as they are often corrupted by high-frequency artifacts and low-frequency drift. Motivated by these challenges, we propose FD-VLA guided by three design principles: (i) leverage predicted force tokens obtained through distillation with actual force signals during training only; (ii) maximize feature-level cross-modal fusion and alignment before, within, and after the VLM so as to

preserve pretrained semantic priors while enriching action representations; and (iii) enable sensor-free inference at deployment, ensuring practical applicability across diverse robotic platforms.

B. Overall Framework

Fig. 2 illustrates the overall framework of the proposed FD-VLA. The multimodal inputs of VLA include language instruction L , visual observation V_t , robot state S_t , and force F_t , where t denotes the timestamp. Following SmolVLA, we employ a pretrained vision-language model SmolVLM-2 with SigLIP [40] as the perception backbone, and computations in




Fig. 3: Visualization of raw force in the plug insertion task.

the VLM are selectively skipped. First, the visual observation, language instruction, and robot proprioceptive state are encoded into their corresponding feature representations $\mathbf{f}_t^V \in \mathbb{R}^{N_v \times D}$, $\mathbf{f}_t^L \in \mathbb{R}^{N_l \times D}$, and $\mathbf{f}_t^S \in \mathbb{R}^{N_s \times D}$, respectively, by leveraging on the VLM’s pretrained tokenizer and vision encoder, and a lightweight projection layer consisting of a single-layer MLP. These embeddings are projected to match the hidden dimension D of the VLM, ensuring compatibility for subsequent fusion and processing.

Rather than directly incorporating the raw force measurements into the VLM, we introduce the Force Distillation Module (FDM) that can predict a latent force representation from the real-time visual observation and robot proprioceptive state,

$$\mathbf{f}_t^{pF} = \text{FDM}(\mathbf{f}_t^V, \mathbf{f}_t^S, \mathbf{p}), \quad (1)$$

where $\mathbf{p} \in \mathbb{R}^{1 \times D}$ is a learnable token for force prediction. Note that the FDM is trained to align with real force signals $\mathbf{f}^{aF} \in \mathbb{R}^{1 \times D}$ encoded by a projection layer only during the training stage, while at inference time it operates entirely without force measurements.

Next, all four modalities of inputs are concatenated along the token dimension, and fuse as $\mathbf{f}_t^{fs} = [\mathbf{f}_t^V, \mathbf{f}_t^L, \mathbf{f}_t^S, \mathbf{f}_t^{pF}]$ to pass into the VLM for multimodal reasoning. Within the VLM, the integrity of pretrained visual-linguistic knowledge are carefully preserved as the VLM parameters are frozen, and an attention masking strategy is employed to ensure that the state and predicted force tokens (\mathbf{f}_t^S and \mathbf{f}_t^{pF}) remain disentangled from the core visual and language streams. This separation prevents interference with the pretrained semantic representations while still enabling the action head to integrate all modalities effectively for downstream control. Finally, an action expert follows a transformer-based policy head with a conditional flow-matching decoder that maps the fused multimodal embeddings to action sequences,

$$[a_t, a_{t+1}, \dots, a_{t+H-1}] = \pi_\theta(\text{VLM}(\mathbf{f}_t^{fs})), \quad (2)$$

where $\pi_\theta(\cdot)$ represents the action expert policy head. This architecture allows our system to leverage the semantic richness of pretrained VLM while introducing stable, task-relevant physical reasoning through force distillation, achieving both robustness and efficiency without additional sensing requirements.

C. Force Distillation Module (FDM)

Our FDM generates a compact, state-aware force representation that can be seamlessly integrated into the VLA pipeline without requiring specialized tactile hardware in inference. The key insight is that force and tactile information are not isolated measurements but are inherently coupled with the robot’s proprioceptive state and its visual manipulation context: variations in joint torques, velocities, and positions, together with visual cues such as object deformation and spatial changes, jointly reflect the underlying contact forces. By exploiting this strong correlation, FDM distills force representations directly from multimodal inputs that are already available on most robotic platforms, thereby eliminating the need for additional physical sensors. Furthermore,

modeling this inherent mapping enhances cross-modal fusion and alignment and facilitate more accurate perception-action coupling in contact-rich tasks. Finally, FDM mitigates the noise and instability of raw sensor signals by learning a supervised latent embedding that serves as a denoised, task-relevant proxy for physical force, filtering out high-frequency artifacts while preserving the dynamics most critical for manipulation.

The realization of FDM consists of two parallel branches, i.e., the prediction branch and actual force branch. In the actual force branch, the raw force measurements collected from the robot’s force/torque sensor are encoded to \mathbf{f}_t^{aF} by a lightweight MLP to project them into the same embedding space as the predicted force representation. This projected embedding serves as the supervisory signal for the prediction branch only during the training stage.

In the prediction branch, the attention mechanism is utilized to predict the learnable token $\mathbf{p} \in \mathbb{R}^{1 \times D}$ into the latent force feature, given the current images embeddings \mathbf{f}_t^V and robot proprioceptive state embeddings \mathbf{f}_t^S . In detail, the design of prediction branch casts the force latent generation as a retrieval problem conditioned on the vision-state context by leveraging a single learned query \mathbf{p} . A context matrix is defined as $\mathbf{C}_t = [\mathbf{p}, \mathbf{f}_t^V, \mathbf{f}_t^S] \in \mathbb{R}^{N_c \times D}$, where $N_c = 1 + N_v + N_s$. Using a single-query multi-head attention, \mathbf{p} serves as the only query $\mathbf{Q} = \mathbf{p}\mathbf{W}_Q \in \mathbb{R}^{H \times d_k}$ while \mathbf{C}_t provides keys $\mathbf{K} = \mathbf{C}_t\mathbf{W}_K \in \mathbb{R}^{N \times H \times d_k}$ and values $\mathbf{V} = \mathbf{C}_t\mathbf{W}_V \in \mathbb{R}^{N \times H \times d_k}$. The output of the multi-head attention \mathbf{Z} can be calculated as

$$\alpha_t = \text{softmax}\left(\frac{(\mathbf{p}\mathbf{W}_Q)([\mathbf{p}, \mathbf{f}_t^V, \mathbf{f}_t^S]\mathbf{W}_K)^\top}{\sqrt{d_k}}\right), \quad (3)$$

$$\mathbf{Z} = \alpha_t \mathbf{V} \in \mathbb{R}^{H \times d_k},$$

where α_t represents a normalized weight vector indicating the relative importance of each token in the context. The final predicted force token can be evaluated as

$$\mathbf{f}_t^{pF} = \text{FFN}(\text{LN}([\mathbf{Z}_1, \dots, \mathbf{Z}_H]\mathbf{W}_O + \mathbf{p})). \quad (4)$$

Including the query token \mathbf{p} in both the key and value sets enables self-conditioning through the residual path, injecting the learned contact prior directly into the aggregation.

Finally, the force distillation is achieved by the alignment between the feature representations from these two branches, where an auxiliary distillation loss is used as

$$\mathcal{L}_{\text{dist}} = \|\mathbf{f}_t^{pF} - \mathbf{f}_t^{aF}\|_2^2. \quad (5)$$

D. Directional Attention Masking Mechanism

To incorporate state and force tokens into the pretrained VLM without corrupting its semantic priors, inspired by SmolVLA [20], we adopt a directional attention masking mechanism. The key idea is to treat vision and language tokens as the frozen perceptual stream and state/force tokens as the control stream.

The input four-modality tokens are divided into two streams, $\mathbf{X} = \{\mathbf{X}^{\text{percept}}, \mathbf{X}^{\text{control}}\}$, where $\mathbf{X}^{\text{percept}} = \{\mathbf{f}^V, \mathbf{f}^L\}$ are the frozen perceptual stream, and $\mathbf{X}^{\text{control}} = \{\mathbf{f}^S, \mathbf{f}^{pF}\}$ are

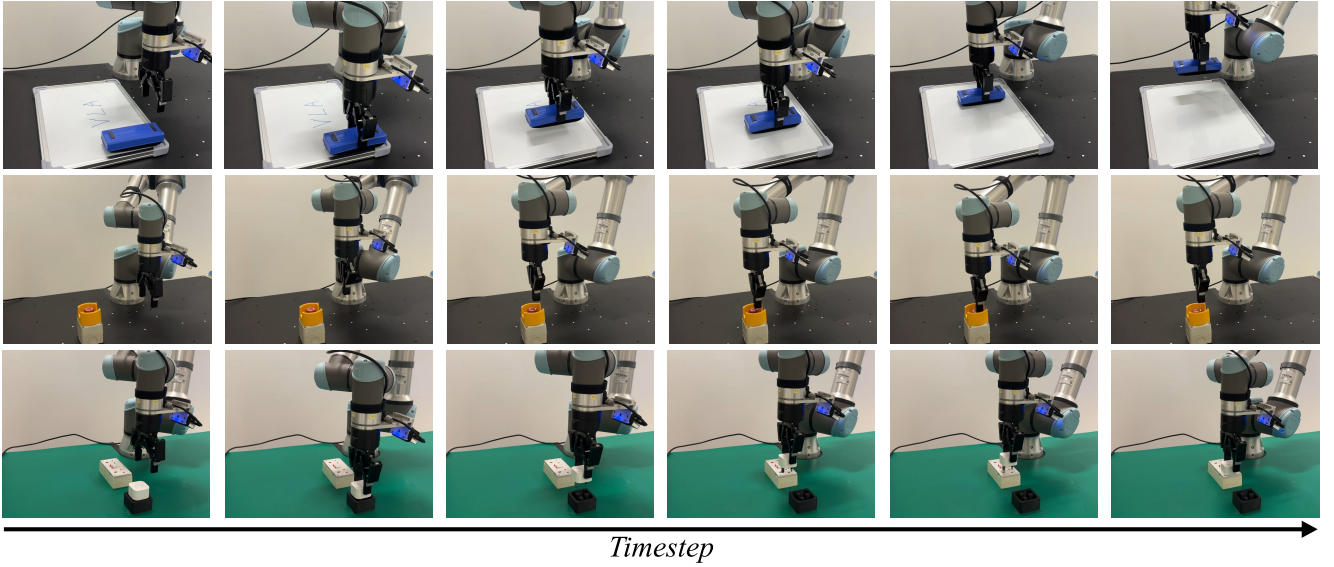


Fig. 4: Visualization of real-world experimental tasks: 1) Clean the whiteboard, 2) Press the emergency button, 3) Insert the plug into the socket.

the control stream. To preserve pretrained semantics while enabling fusion, we construct a Directional Attention Mask $M \in \{0, 1\}^{N_c \times N_c}$, where

$$M_{ij} = \begin{cases} 1, & \text{if } x_i \in \mathbf{X}^{\text{percept}} \text{ and } x_j \in \mathbf{X}^{\text{percept}}, \\ 1, & \text{if } x_i \in \mathbf{X}^{\text{control}} \text{ and } x_j \in \mathbf{X}^{\text{percept}}, \\ 1, & \text{if } x_i \in \mathbf{X}^{\text{control}}, x_j \in \mathbf{X}^{\text{control}}, \text{ and } i \geq j, \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

Applying the mask inside the VLM’s transformer yields masked self-attention, $\mathbf{X}' = \text{SelfAttention}(\mathbf{X}, M)$. Specifically,

$$\begin{aligned} \mathbf{X}'^{\text{percept}} &= \text{SelfAttention}(\mathbf{X}^{\text{percept}}), \\ \mathbf{X}'^{\text{control}} &= \text{SelfAttention}(\mathbf{X}^{\text{control}}, \mathbf{X}^{\text{percept}}). \end{aligned} \quad (7)$$

This design enforces one-way information flow, i.e., perceptual tokens can only attend to each other, preserving pretrained vision-language alignment, while control tokens can attend to perceptual tokens, thereby fusing proprioception and force with perceptual context for downstream action prediction.

E. Action Expert

The action expert π_θ is instantiated as a transformer that predicts an action chunk $\mathbf{A}_t = [a_t, \dots, a_{t+H-1}]$ conditioned on VLM features \mathbf{X}_t . We train the action expert with a conditional flow-matching objective,

$$\mathcal{L}^\tau(\theta) = \mathbb{E}_{p(\mathbf{A}_t|\mathbf{X}_t), q(\mathbf{A}_t^\tau|\mathbf{A}_t)} \left[\left\| \mathbf{v}_\theta(\mathbf{A}_t^\tau, \mathbf{X}_t) - \mathbf{u}(\mathbf{A}_t^\tau | \mathbf{A}_t) \right\|^2 \right], \quad (8)$$

where $\mathbf{A}_\tau = \tau \mathbf{A} + (1 - \tau) \epsilon$, with $\epsilon \sim \mathcal{N}(0, I)$. $\pi_\theta(\cdot)$ is trained to output the vector field $\mathbf{u}(\mathbf{A}_t^\tau | \mathbf{A}_t) = \epsilon - \mathbf{A}_t$ from the VLM features and the noisy actions \mathbf{A}_t^τ , τ is sampled from a Beta distribution. The policy learns a velocity field rather than stepwise residuals, which is well suited to chunked action prediction.




Fig. 5: Visualization of the real robotic platform. We use a UR5e robot arm as the main manipulation platform, the Kinect Azure camera as the main camera and the RealSense D405 as the gripper camera.

F. Overall Training Objective

The overall objective of our FD-VLA framework combines two complementary components, i.e., a standard policy learning loss and a force-distillation loss, which is defined as

$$\mathcal{L} = \mathcal{L}^\tau(\theta) + \lambda \mathcal{L}_{\text{dist}}, \quad (9)$$

where λ is the weight for controlling force supervision.

IV. EXPERIMENTS

This section presents an extensive evaluation of the FD-VLA model through real-world experiments and analytical studies.




Fig. 6: Success rates for three contact-rich manipulation tasks: Plug in Socket, Clean Whiteboard, and Press Button. Results are averaged over 30 evaluation episodes per task. We compare FD-VLA (ours) with SmolVLA, π_0 and DP3, SmolVLA and π_0 are evaluated with and without force inputs. Our model achieves consistently higher performance, which highlights the benefit of force distillation for accurate and robust manipulation.

A. Experimental Setup

In the experiments, we use a UR5e robotic arm as our main manipulation platform, an Azure Kinect camera for main RGB image acquisition and a RealSense D405 for robot gripper RGB image acquisition. The real experimental platform is shown in Fig. 5. This platform provides a large workspace, high repeatability, and easy programmability. During operation, the robotic arm and the cameras work together through real-time data exchange to achieve precise manipulation.

B. Experimental Tasks

To evaluate the effectiveness of FD-VLA, we conducted experiments on three diverse contact-rich manipulation tasks: 1) Clean Whiteboard, demanding consistent motion planning with continuous contact across the surface. 2) Press Emergency Button, requiring precise vertical pressing and exceeding the resistance and spring position of the button. 3) Insert Plug, which requires precise positional alignment and regulated force application during engagement. The visualization of the tasks is shown in Fig. 4. During data collection, we use the force sensor inside the UR5e robot arm to obtain actual force information. To efficiently collect expert demonstrations, we use the 3Dconnexion SpaceMouse to teleoperate the robotic arm. We collect 50 demonstrations for each task. The demonstration procedure is standardized to ensure that all operators follow a predefined protocol consistently.

For evaluation, each task was trained using a set of 50 demonstrations and subsequently evaluated over 30 independent test trials to ensure statistical robustness. Success criteria were defined based on precise task-specific outcomes. For the Clean Whiteboard task, a trial was considered successful only if all visible markings on the whiteboard were completely

erased. In the Press Emergency Button task, success required the button to be fully depressed and remain engaged without rebound. For the Insert Plug task, success was achieved when the plug was fully seated within the socket, indicating proper alignment and insertion.

C. Main Results

To ensure a fair and comprehensive evaluation, we benchmark FD-VLA against three representative policies: DP3 [41], π_0 [18], and SmolVLA [20]. DP3 is selected as a strong diffusion-based control framework with a parameter scale comparable to ours, which provides a capacity-matched baseline that excels at generating precise motion trajectories. π_0 represents the state-of-the-art VLA paradigm with substantially larger capacity and powerful language-conditioned control. Including π_0 in baselines allows us to compare against a substantially larger model with scale and broad pretraining. SmolVLA is a recent compact VLA model that integrates multimodal reasoning and control through a compact transformer-based design, which also serves as our foundational model. For all baselines, we standardize training data, evaluation tasks, and optimization budgets to control for dataset and compute confounds, and we report all evaluation metrics following identical experimental protocols to guarantee a rigorous and fair comparison between methods.

Fig. 6 presents the success rates across three contact-rich manipulation tasks. For example, in Clean Whiteboard task our FD-VLA achieves a success rate of 73.3%, which is at least 23.3% higher than any other baselines. Across all the tasks, our FD-VLA achieves the highest overall performance with a mean success rate of 61.1%, substantially outperforming both SmolVLA without force encoding (23.3%), DP3 (11.1%) and even π_0 without force encoding (46.7%) which contains approximately ten times as many parameters as our

TABLE I: Ablation study of Force Distillation Module.

Model	Plug	Clean	Press	Mean
w/o FDM	8/30	15/30	12/30	38.9%
FDM w/ force token	12/30	17/30	17/30	51.1%
FDM w/ learnable token	12/30	22/30	21/30	61.1%

model.

To further evaluate the contribution of our force distillation approach, we compare FD-VLA with models trained using direct raw force inputs. By injecting raw force information, SmolVLA achieves an improvement of 15.6% but still falls short of our method by 22.2%. π_0 also shows an improvement of 20.0%, yet our approach continues to outperform it by 14.4% despite π_0 operating at a substantially larger parameter scale. These results suggest that simply providing raw force signals is insufficient for optimal performance, as the model must directly process noisy, high-dimensional sensory inputs. In contrast, our predicted fused force module integrates information from both vision and robot state, enabling the extraction of compact and task-relevant force representations. In addition, models trained with raw force require access to real-time force measurements during inference, which limits their applicability to platforms with specialized hardware.

Our approach overcomes this limitation by predicting force representations directly from visual and state inputs, eliminating the need for force sensors at test time while still achieving superior performance. This highlights both the effectiveness and practicality of our method for deploying contact-rich manipulation policies in real-world settings.

D. Ablation Study

To validate the proposed Force Distillation module, we conduct an ablation by replacing the learnable force tokens in the Force Distillation module with direct ground-truth force measurements obtained from sensors. We compare three variants on three contact-rich tasks, each with 30 independent rollouts under the same initialization, time budget, and success criteria as the main experiments: 1) Without FDM, directly using an MLP layer to encode force. 2) FDM with force token, replacing the learnable token with the actual force token encoded by an MLP layer. 3) FDM with learnable token, our proposed method.

Table I shows the clear improvement in adding FDM, which indicates that deep cross-modal interactions in the early fusion are a key driver for the improvements. From the table, we could also find that when replacing the learnable token with the latent force features encoded from ground-truth force measurements, the mean success rate drops. This result highlights the challenge of directly incorporating raw force: raw signals are often high-dimensional and noisy, which makes them difficult for the policy to integrate effectively. In contrast, the learnable token not only avoids the pitfalls of directly feeding fixed force features but also retains the benefits of real physical supervision, providing reliable contact information.




Fig. 7: Illustration of our visual generalization settings: **(Left)** Novel Background, where the scene’s background differs from training; and **(Right)** Visual Perturbation, which involves changes in colorful lights.

E. Model Generalization

Visual generalization denotes a robot’s ability to sustain task performance under distribution shifts in scene appearance. We study this challenge in two representative scenarios: (1) **Novel Background**, where layouts and textures differ from training, and (2) **Visual Perturbations**, where lighting, contrast, and color distributions are altered. As illustrated in Fig. 7, these settings reflect common real-world challenges, requiring robots to operate beyond the narrow visual domain of training data. Qualitative results show that our method remains robust across both scenarios. On novel backgrounds, the robot leverages structural cues rather than overfitting to textures. Under photometric perturbations, it adapts smoothly to illumination and appearance changes without significant policy degradation. These findings highlight the model’s capacity to extract task-relevant features that remain invariant to superficial visual shifts.

V. CONCLUSIONS

In this work, we introduce FD-VLA, a novel VLA architecture that integrates force awareness and is tailored for contact-rich manipulation tasks. Instead of directly relying on noisy and hardware-dependent raw force signals, our method employs a Force Distillation Module (FDM), which distills task-relevant and denoised force representations from vision and state inputs. We leverage attention masks to keep state and predicted force tokens separate from the pretrained VLM stream. This design prevents interference with semantic representations while achieving efficient multimodal integration with minimal additional parameters. The fused features are then consumed by the action expert to generate precise, contact-aware control. Experiments across diverse real-world tasks show that FD-VLA consistently outperforms baselines without force reasoning or with raw force inputs. We hope this work paves the way for advancing multimodal learning in VLAs toward complex manipulation tasks. Looking ahead, we plan to extend this framework to multi-robot coordination and move toward scalable, general-purpose manipulation systems.

REFERENCES

- [1] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu *et al.*, “Rt-1: Robotics transformer for real-world control at scale,” *arXiv preprint arXiv:2212.06817*, 2022.

- [2] B. Zitkovich, T. Yu, S. Xu, P. Xu, T. Xiao, F. Xia, J. Wu, P. Wohlhart, S. Welker, A. Wahid *et al.*, “Rt-2: Vision-language-action models transfer web knowledge to robotic control,” in *Conference on Robot Learning*. PMLR, 2023, pp. 2165–2183.
- [3] O. M. Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, T. Kreiman, C. Xu *et al.*, “Octo: An open-source generalist robot policy,” *arXiv preprint arXiv:2405.12213*, 2024.
- [4] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi *et al.*, “Openvla: An open-source vision-language-action model,” *arXiv preprint arXiv:2406.09246*, 2024.
- [5] Z. Zhang, H. Li, Y. Dai, Z. Zhu, L. Zhou, C. Liu, D. Wang, F. E. Tay, S. Chen, Z. Liu *et al.*, “From spatial to actions: Grounding vision-language-action model in spatial foundation priors,” *arXiv preprint arXiv:2510.17439*, 2025.
- [6] A. Steiner, A. S. Pinto, M. Tschannen, D. Keysers, X. Wang, Y. Bitton, A. Gritsenko, M. Minderer, A. Sherbondy, S. Long *et al.*, “Paligemma 2: A family of versatile vlms for transfer,” *arXiv preprint arXiv:2412.03555*, 2024.
- [7] 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.
- [8] W. Xie and N. Correll, “Towards forceful robotic foundation models: a literature survey,” *arXiv preprint arXiv:2504.11827*, 2025.
- [9] N. Funk, C. Chen, T. Schneider, G. Chahvatzaki, R. Calandra, and J. Peters, “On the importance of tactile sensing for imitation learning: A case study on robotic match lighting,” *arXiv preprint arXiv:2504.13618*, 2025.
- [10] J. Yu, H. Liu, Q. Yu, J. Ren, C. Hao, H. Ding, G. Huang, G. Huang, Y. Song, P. Cai *et al.*, “Forcevla: Enhancing vla models with a force-aware moe for contact-rich manipulation,” *arXiv preprint arXiv:2505.22159*, 2025.
- [11] J. Huang, S. Wang, F. Lin, Y. Hu, C. Wen, and Y. Gao, “Tactile-vla: Unlocking vision-language-action model’s physical knowledge for tactile generalization,” *arXiv preprint arXiv:2507.09160*, 2025.
- [12] Z. Zhang, H. Xu, Z. Yang, C. Yue, Z. Lin, H.-a. Gao, Z. Wang, and H. Zhao, “Ta-vla: Elucidating the design space of torque-aware vision-language-action models,” *arXiv preprint arXiv:2509.07962*, 2025.
- [13] J. Bi, K. Y. Ma, C. Hao, M. Z. Shou, and H. Soh, “Vla-touch: Enhancing vision-language-action models with dual-level tactile feedback,” *arXiv preprint arXiv:2507.17294*, 2025.
- [14] W. Wang, H. Zhu, and M. H. Ang Jr, “Sgsin: Simultaneous grasp and suction inference network via attention-based affordance learning,” *IEEE Transactions on Industrial Electronics*, 2024.
- [15] U. Asif, J. Tang, and S. Herrer, “Graspnet: An efficient convolutional neural network for real-time grasp detection for low-powered devices,” in *IJCAI*, vol. 7, 2018, pp. 4875–4882.
- [16] S. Liu, T. J. Teo, Z. Lin, and H. Zhu, “Relationgrasp: Object-oriented prompt learning for simultaneously grasp detection and manipulation relationship in open vocabulary,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 10 890–10 896.
- [17] C.-L. Cheang, G. Chen, Y. Jing, T. Kong, H. Li, Y. Li, Y. Liu, H. Wu, J. Xu, Y. Yang *et al.*, “Gr-2: A generative video-language-action model with web-scale knowledge for robot manipulation,” *arXiv preprint arXiv:2410.06158*, 2024.
- [18] K. Black, N. Brown, D. Driess, A. Esmail, M. Equi, C. Finn, N. Fusai, L. Groom, K. Hausman, B. Ichter *et al.*, “ π_0 : A vision-language-action flow model for general robot control,” *arXiv preprint arXiv:2410.24164*, 2024.
- [19] J. Wen, Y. Zhu, J. Li, M. Zhu, Z. Tang, K. Wu, Z. Xu, N. Liu, R. Cheng, C. Shen *et al.*, “Tinyvla: Towards fast, data-efficient vision-language-action models for robotic manipulation,” *IEEE Robotics and Automation Letters*, 2025.
- [20] M. Shukor, D. Aubakirova, F. Capuano, P. Kooijmans, S. Palma, A. Zouitine, M. Aractingi, C. Pascal, M. Russi, A. Marafioti *et al.*, “Smolvla: A vision-language-action model for affordable and efficient robotics,” *arXiv preprint arXiv:2506.01844*, 2025.
- [21] 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.
- [22] A. O’Neill, A. Rehman, A. Maddukuri, A. Gupta, A. Padalkar, A. Lee, A. Pooley, A. Gupta, A. Mandlekar, A. 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)*. IEEE, 2024, pp. 6892–6903.
- [23] A. Khazatsky, K. Pertsch, S. Nair, A. Balakrishna, S. Dasari, S. Karamcheti, S. Nasiriany, M. K. Srirama, L. Y. Chen, K. Ellis *et al.*, “Droid: A large-scale in-the-wild robot manipulation dataset,” *arXiv preprint arXiv:2403.12945*, 2024.
- [24] H. R. Walke, K. Black, T. Z. Zhao, Q. Vuong, C. Zheng, P. Hansen-Estruch, A. W. He, V. Myers, M. J. Kim, M. Du *et al.*, “Bridgedata v2: A dataset for robot learning at scale,” in *Conference on Robot Learning*. PMLR, 2023, pp. 1723–1736.
- [25] W. Zhao, P. Ding, M. Zhang, Z. Gong, S. Bai, H. Zhao, and D. Wang, “Vlas: Vision-language-action model with speech instructions for customized robot manipulation,” *arXiv preprint arXiv:2502.13508*, 2025.
- [26] C. Li, J. Wen, Y. Peng, Y. Peng, F. Feng, and Y. Zhu, “Pointvla: Injecting the 3d world into vision-language-action models,” *arXiv preprint arXiv:2503.07511*, 2025.
- [27] H. Zhen, X. Qiu, P. Chen, J. Yang, X. Yan, Y. Du, Y. Hong, and C. Gan, “3d-vla: A 3d vision-language-action generative world model,” *arXiv preprint arXiv:2403.09631*, 2024.
- [28] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *The International Journal of Robotics Research*, p. 02783649241273668, 2023.
- [29] W. Wang, H. Zhu, and M. H. Ang, “Grascontrast: Self-supervised contrastive learning with false negative elimination for 6-dof grasp detection,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 7294–7300.
- [30] Y. Wu, Z. Chen, F. Wu, L. Chen, L. Zhang, Z. Bing, A. Swikir, S. Haddadin, and A. Knoll, “Tacdifffusion: Force-domain diffusion policy for precise tactile manipulation,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 11 831–11 837.
- [31] W. Liu, J. Wang, Y. Wang, W. Wang, and C. Lu, “Foremimic: Force-centric imitation learning with force-motion capture system for contact-rich manipulation,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 1105–1112.
- [32] Z. He, H. Fang, J. Chen, H.-S. Fang, and C. Lu, “Foar: Force-aware reactive policy for contact-rich robotic manipulation,” *IEEE Robotics and Automation Letters*, 2025.
- [33] Y. Hou, Z. Liu, C. Chi, E. Cousineau, N. Kuppusswamy, S. Feng, B. Burchfiel, and S. Song, “Adaptive compliance policy: Learning approximate compliance for diffusion guided control,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 4829–4836.
- [34] H. Xue, J. Ren, W. Chen, G. Zhang, Y. Fang, G. Gu, H. Xu, and C. Lu, “Reactive diffusion policy: Slow-fast visual-tactile policy learning for contact-rich manipulation,” *arXiv preprint arXiv:2503.02881*, 2025.
- [35] J. Jones, O. Mees, C. Sferrazza, K. Stachowicz, P. Abbeel, and S. Levine, “Beyond sight: Finetuning generalist robot policies with heterogeneous sensors via language grounding,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Atlanta, USA, 2025.
- [36] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *arXiv preprint arXiv:1503.02531*, 2015.
- [37] S. Gupta, J. Hoffman, and J. Malik, “Cross modal distillation for supervision transfer,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2827–2836.
- [38] C. Wang, J. M. Buenaposada, R. Zhu, and S. Lucey, “Learning depth from monocular videos using direct methods,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2022–2030.
- [39] W. Xia, X. Li, A. Deng, H. Xiong, D. Dou, and D. Hu, “Robust cross-modal knowledge distillation for unconstrained videos,” *arXiv preprint arXiv:2304.07775*, 2023.
- [40] X. Zhai, B. Mustafa, A. Kolesnikov, and L. Beyer, “Sigmoid loss for language image pre-training,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 11 975–11 986.
- [41] Y. Ze, G. Zhang, K. Zhang, C. Hu, M. Wang, and H. Xu, “3d diffusion policy: Generalizable visuomotor policy learning via simple 3d representations,” in *Proceedings of Robotics: Science and Systems (RSS)*, 2024.