

Cross-Modal Instructions for Robot Motion Generation

William Baron^{1,2}

Xiaoxiang Dong^{1,2}

Matthew Johnson-Roberson¹

Weiming Zhi^{3,1*}




Fig. 1: We enable robots to interpret *cross-modal instructions*, in the form of rough sketches and textual labels (shown on the left). Subsequent motion can be generated and generalized to novel setups and environments (shown on the right).

Abstract—Teaching robots novel behaviors typically requires motion demonstrations via teleoperation or kinaesthetic teaching, that is, physically guiding the robot. While recent work has explored using human sketches to specify desired behaviors, data collection remains cumbersome, and demonstration datasets are difficult to scale. In this paper, we introduce an alternative paradigm, Learning from Cross-Modal Instructions, where robots are shaped by demonstrations in the form of rough annotations, which can contain free-form text labels, and are used in lieu of physical motion. We introduce the *CrossInstruct* framework, which integrates cross-modal instructions as examples into the context input to a foundational vision-language model (VLM). The VLM then iteratively queries a smaller, fine-tuned model, and synthesizes the desired motion over multiple 2D views. These are then subsequently fused into a coherent distribution over 3D motion trajectories in the robot’s workspace. By incorporating the reasoning of the large VLM with a fine-grained pointing model, *CrossInstruct* produces executable robot behaviors that generalize beyond the environment of in the limited set of instruction examples. We then introduce a downstream reinforcement learning pipeline that leverages *CrossInstruct*

outputs to efficiently learn policies to complete fine-grained tasks. We rigorously evaluate *CrossInstruct* on benchmark simulation tasks and real hardware, demonstrating effectiveness without additional fine-tuning and providing a strong initialization for policies subsequently refined via reinforcement learning.

I. INTRODUCTION

Imitation learning, also known as Learning from Demonstration, is a dominant paradigm for teaching robots new skills [1]. Imitation learning requires the collection of a dataset of demonstrations, which are generally provided by a human user teleoperating the robot via a remote controller or by kinaesthetic teaching, i.e. physically handling the robot. In both of these cases, collecting a sufficient number of demonstrations for even marginal levels of generalization is challenging. There have also been efforts to reduce the human burden of providing demonstrations, including hardware adaptations [2], and using human sketches as a demonstration interface [3], [4], [5], [6], [7]. Humans have a remarkable ability to adapt the same skill to different environments and setups, even after observing one or two demonstrations. However, imitation learning is far less demonstration-efficient, and it can be challenging to collect a sufficient number of demonstrations to generalize

*email: Weiming.Zhi@sydney.edu.au.

¹ College of Connected Computing, Vanderbilt University, TN, USA

² Robotics Institute, Carnegie Mellon University, PA, USA

³ School of Computer Science, The University of Sydney, Australia.

⁴ Australian Centre for Robotics, The University of Sydney, Australia.

behavior when the scene changes. In this work, we aim to circumvent the cumbersome, large-scale human input required for imitation learning.

We introduce Learning from Cross-modal Instructions as an alternative paradigm to shape robot behavior. In the proposed method, human operators provide *cross-modal instructions*, which are free-form sketches and textual labels on an image of an operating environment. The robot infers the intended behavior from these inputs and generalizes execution across varying setups. To learn from cross-modal instructions, we propose the CrossInstruct framework. Figure 1 illustrates several examples in both simulation and the real world with cross-modal instructions and subsequent robot behavior generated by CrossInstruct. CrossInstruct takes cross-modal instructions as *in-context learning examples* [8] for a large VLM. The large model performs high-level task reasoning and delegates pixel-level keypoint localization to a smaller VLM fine-tuned for 2D pointing, which returns precise coordinates for task-critical features identified by the large model. The hierarchical coupling between the large reasoning VLM with the smaller, fine-tuned model enables the efficient and accurate identification of task-relevant keypoints. These are then used to produce motion trajectory sketches over a small set of multi-view images of the setup. The resulting 2D trajectories are lifted into the robot’s workspace to yield a distribution over 3D trajectories. Combined with end-effector orientations and gripper actions generated by the reasoning model, this produces executable motion sequences.

We demonstrate the capabilities of CrossInstruct to shape behavior to complete and generalize across a wide range of tasks, in both simulation and in the real world, without additional fine-tuning. Beyond directly executing the produced motions, we can enable the robots to achieve an even higher level of precision by joining CrossInstruct with a downstream reinforcement learning pipeline. Concretely, the technical contributions in this paper include:

- The Learning from Cross-Modal Instructions paradigm, circumventing the need for physical demonstrations;
- The CrossInstruct framework to leverage cross-modal instructions in an in-context fashion to synthesize robot behavior. This is achieved by a hierarchical coupling of two VLMs and using raycasting to fuse multi-view 2D sketches into executable robot motion;
- Evaluations on benchmark tasks, demonstrating robustness of CrossInstruct out of the box in both simulation and the real world, as well as leveraging the distribution of trajectories to enable efficient reinforcement learning.

II. RELATED WORK

Robot Motion Generation: Generating movements for robots is a critical area of study in robotics. Classical approaches formulate a motion planning problem [9], [10], [11], where desired goals and constraints are specified *a priori*. Another popular paradigm, imitation learning or learning from demonstrations [1], circumvents the specification of “how” the




Fig. 2: Example cross-modal instructions which a human can specify to complete the *Basketball in Hoop* problem in RL-Bench [27].

motion should be, by mimicking a set of shown robot trajectory demonstrations. Various approaches have been proposed to formulate mimicking as a deterministic or probabilistic supervised learning problem to learn actions [12], [13], predict high-level way-points [14], or as a *inverse* reinforcement learning problem [15].

Sketches for Robot Learning: Providing robot trajectories as demonstrations often requires careful tele-operation or cumbersome kinesthetic teaching. Sketches over images have emerged as a lightweight channel for specifying motion. These include techniques for directly interpreting sketches of precise desired robot motion [3], [16], [17]. Others translate sketches into 3D demonstrations that bootstrap data collection and reinforcement learning [18]. A related direction conditions imitation policies directly on sketches, where drawings of desired outcomes are provided as goals, guiding behavior [19]. Sketching has emerged as an impactful interface for specifying fine-grained details for robot motions.

Large Pre-trained Models in Robotics: Large language and vision–language models have been used as task planners, decomposing natural-language goals into perception queries and motion objectives [20]. Earlier work revolves around grounding language in action, these include Perceiver-Actor [21] and SayCan [22]. Many approaches construct 3D value maps optimized by motion planners [23], while others output coarse sketches refined by low-level controllers [24]. In these cases, models remain high-level reasoners, with trajectory generation handled by separate planning or control modules. Vision–language models have also been employed to structure perception through semantic cues. Examples include open-vocabulary detectors that produce voxelized affordance fields optimized by planners [23], and hierarchical structures for reasoning [25]. Instruction-tuned VLMs have also been applied for pixel-level pointing, producing 2D keypoints or affordances without explicit control code or volumetric fields [26]. These strategies highlight the utility of VLMs in providing precise and semantically rich cues for downstream planning.

III. LEARNING FROM CROSS-MODAL INSTRUCTIONS

We consider a robot operating in a workspace $\mathcal{W} \subseteq \mathbb{R}^3$, where the goal is to generate executable motion trajectories conditioned on *cross-modal instructions*. A cross-modal instruction is defined as a tuple

$$\mathcal{I} = \{I, S, T\}, \quad (1)$$

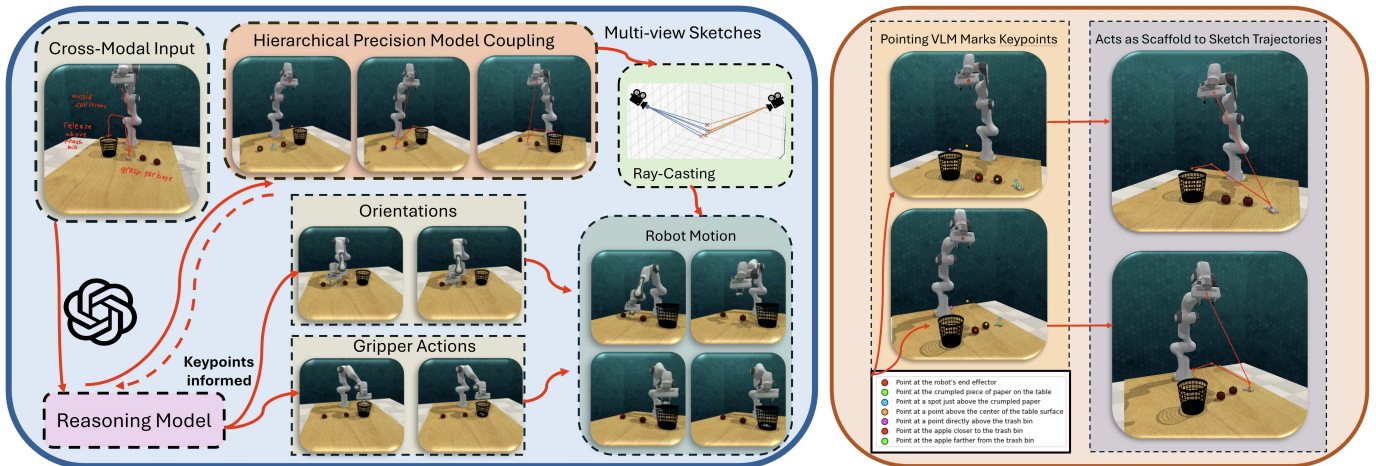


Fig. 3: An overview of CrossInstruct (left) with an example of precise pointing of keypoints (right). The hierarchical precision model coupling module enables the reasoning model to leverage a smaller fine-tuned VLM to precisely identify relevant keypoints, which then guide robot end-effector motion.

where I is an image of a relevant scene, S is a set of free-form sketches provided by a human operator (e.g., lines, arrows, or annotated paths over I), and T is an optional set of textual information describing constraints and the intended action, which can be directly handwritten over the image. An example of such cross-modal instructions, including arrow notations and text, are given in fig. 2.

The goal is for the robot to generalize the behavior captured in the instructions to a new environment or setup. Given a tuple of multi-view images of the new scene, \mathcal{V} , along with their poses, \mathcal{P} , the robot must map \mathcal{I} to a *distribution over motion trajectories* in the workspace:

$$p(\tau \mid \mathcal{I}, \mathcal{V}, \mathcal{P}), \quad \tau = \{(x_t, R_t, g_t)\}_{t=1}^H, \quad (2)$$

where $x_t \in \mathbb{R}^3$ is the position of the end-effector at time t , $R_t \in SO(3)$ is the orientation, $g_t \in \{0, 1\}$ encodes the open or closed grip state and H is the trajectory horizon. During execution, the mean of the trajectory distribution $\mathcal{E}(\tau)$ is tracked, while the entire distribution can be used to generate varying trajectories in order to adequately train downstream reinforcement learning policies.

The learning problem is thus, given cross-modal instructions \mathcal{I} , to infer a trajectory distribution $p(\tau \mid \mathcal{I}, \mathcal{V}, \mathcal{P})$ that generalizes across environments. Unlike imitation learning, where supervision comes from expert trajectories $\mathcal{D}_{\text{IL}} = \{\tau^{(i)}\}_{i=1}^N$, our supervision consists only of weak, symbolic, or high-level annotations $\{\mathcal{I}^{(i)}\}_{i=1}^M$. The challenge is to *bridge modalities*, from free-form sketches to executable motion trajectories.

IV. CROSSINSTRUCT FRAMEWORK

The goal of CrossInstruct is to transform human-provided *cross-modal instructions*, sketches S and textual labels T defined over scene images I , into executable robot motion trajectories in \mathbb{R}^3 . This problem requires both high-level semantic interpretation of human intent and precise spatial grounding in the robot’s workspace. Our framework addresses this through a **hierarchical precision coupling** between a large reasoning vision–language model (VLM) and a lightweight, fine-tuned pointing model. The reasoning VLM primarily operates at




Fig. 4: CrossInstruct generates 2D trajectories over multi-view images (in red), which are subsequently fused into a coherent 3D trajectory (waypoints in blue-red color gradient shown). This can then be rolled out to slide the block to the target.

the symbolic level (task identification, task decomposition, intent parsing, trajectory refinement), while the fine-tuned model operates at the perceptual level (pixel-localization of keypoints). An overview of our CrossInstruct framework is given in Figure 3.

A. Hierarchical Precision Coupling Module

The transformation from symbolic, cross-modal instructions $\mathcal{I} = \{I, S, T\}$ to grounded motion trajectories requires reasoning across two levels of abstraction: high-level semantic interpretation of task intent and low-level pixel precision in image space. To achieve this, we introduce a hierarchical precision coupling module that tightly integrates a large reasoning VLM \mathcal{R} with a smaller, fine-tuned pointing model \mathcal{G} . The hierarchical structure allows \mathcal{R} to maintain a global task context, while delegating precise spatial localization, or *pointing*, to \mathcal{G} , which is specifically trained to resolve ambiguous references in concrete pixel coordinates. In this work, our large reasoning VLM is OpenAI o3 [28], and the fine-tuned pointing VLM is Molmo [29].

Given an instruction \mathcal{I} and a tuple of multi-view images $\mathcal{V} = \{I_1, I_2\}$, the reasoning model \mathcal{R} first generates a set of

N semantic keypoint descriptors,

$$\mathcal{K} = \{k_i\}_{i=1}^N, \quad k_i = (\ell_i, \alpha_i), \quad (3)$$

where each ℓ_i is a natural language label and α_i denotes auxiliary text-based metadata and additional details, such as proximity descriptions or qualitative constraints. These descriptors are context-dependent and reflect \mathcal{R} 's global reasoning over both the free-form sketches contained in the instructions, and scene structure contained in the two multi-view images. However, the keypoints are not yet linked to the pixel locations.

The keypoint descriptors are then passed to a fine-tuned pointing model \mathcal{G} , which receives each descriptor ℓ_i together with the images I_1 and I_2 , and predicts a concrete pixel coordinate per keypoint and per multi-view image,

$$\{u_{i,m}^{(t)}, v_{i,m}^{(t)}\} = \mathcal{G}(\ell_i | I_m), \quad m = \{1, 2\}, i = 1, \dots, N. \quad (4)$$

yielding a set of 2D keypoints across views. As \mathcal{G} is trained for pixel-wise pointing within images, we obtain a set of precise keypoints indicating the most salient features relevant to task completion.

The set of keypoints is then added to the context maintained by the larger reasoning model \mathcal{R} . The context now contains the original keypoint descriptors \mathcal{K} , with precise locations $\{\{u_{i,m}, v_{i,m}\}_{i=1}^N\}_{m=1}^2$. Based on this updated context, \mathcal{R} produces two 2D trajectories over each multi-view image. That is,

$$\xi_1, \xi_2 = \mathcal{R}(\mathcal{I}, \mathcal{V}, \mathcal{K}, \{(u_{i,m}^{(t)}, v_{i,m}^{(t)})\}), \quad (5)$$

where ξ_1 and ξ_2 are equal-length pixels over the two views.

This design highlights two key principles. First, the hierarchy enables each model to operate at its natural level of abstraction: \mathcal{R} interprets high-level symbolic cues and manages task context, whereas \mathcal{G} delivers precise geometric observations. Example end-effector trajectories are shown in fig. 4, in the left subfigure. These are then ray-casted into the 3D workspace to estimate a 3D trajectory of end-effector motion. The right subfigure in fig. 4 shows the motion trajectory in 3D visualised as a sequence of waypoints. The reasoning model then takes the 3D trajectory as an in-context example and specifies end-effector orientation and gripper actions to arrive at a desired trajectory to follow, τ . Figure 6 shows the robot behavior generated when the 3D end-effector trajectory, orientation, and gripper actions are combined. In the next section, we shall elaborate on how to obtain estimates of 3D trajectories from ξ_1, ξ_2 .

B. Multi-View Lifting via Ray Casting

The trajectories ξ_1, ξ_2 produced by the reasoning model define continuous curves in pixel space. To embed them into the 3D workspace $\mathcal{W} \subseteq \mathbb{R}^3$, we lift them through calibrated multi-view geometry. The camera poses can either be determined via external tags or prior calibration [30], [31], [32]. We interpret each 2D trajectory ξ_m as defining a time-indexed probability distribution over pixels, $p(u_m, v_m | t)$, for view $m \in \{1, 2\}$. Specifically, for a given time t , the pixel coordinate along the curve, $(u_{t,m}, v_{t,m}) = \xi_m(t)$, defines the

mean of a Gaussian density in image space:

$$p(u_m, v_m | t) = \mathcal{N}((u_m, v_m) | \xi_m(t), \Sigma_m), \quad (6)$$

where Σ_m is an assumed covariance that captures uncertainty, which can be predefined. Uncertainty in the robot pipeline provides additional robustness [33]. Thus, rather than a single pixel, each waypoint induces a distributional neighborhood in the view space. We adopt a standard pinhole camera model to project these pixel densities into the 3D workspace, by casting rays through each camera view and sampling to find where the rays intersect [34].

As illustrated in Figure 5, intersecting viewing rays from different calibrated poses localize a 3D waypoint that is consistent across both per-view selections. For view m , a pixel (u_m, v_m) corresponds to a ray $f_r(u_m, v_m, d) = \mathbf{o}^m + \omega(u_m, v_m) d$,




Fig. 5: Intersecting rays from different poses.

where \mathbf{o}^m is the camera origin, $\omega(u_m, v_m)$ is the ray direction derived from the intrinsic parameters of the camera, and d is the depth parameter. For a fixed time t , the 3D region consistent with the 2D Gaussian density is

$$R_t^m = \{\mathbf{x} \in \mathbb{R}^3 | \mathbf{x} = f_r(u_m, v_m, d), p(u_m, v_m | t) \geq \epsilon\}. \quad (8)$$

The 3D trajectory support at time t is then approximated by the intersection of these ray-cast regions from both views, $\mathcal{R}_t = R_t^1 \cap R_t^2$. In practice, we sample points along d and (u_m, v_m) in both views, and retain pairs whose cross-view projections are within a tolerance δ . This produces a set of spatial samples $\mathcal{S}_t = \{\mathbf{x}_i\}_{i=1}^{n_t}$, which approximate the feasible 3D positions of the end-effector at time t . Finally, we fit a concise Gaussian distribution to these samples to model the time-conditioned waypoint distribution:

$$p(x_t | t) = \mathcal{N}(x_t | \mu_t, \Sigma_t), \quad (9)$$

where μ_t is the mean of \mathcal{S}_t and Σ_t its empirical covariance. Repeating this across all $t = 1, \dots, H$ yields a trajectory distribution

$$p(\tau | \xi_1, \xi_2, \mathcal{P}) = \prod_{t=1}^H p(x_t | t). \quad (10)$$

The direction of the ray $\omega(u_m, v_m)$ is obtained by normalizing the coordinates of the pixels with the intrinsics of the camera of each view, K_m , and rotating into the world frame with R_m^\top . The mean trajectory $\mathbb{E}[\tau] = \{\mu_t\}_{t=1}^H$, combined with the end-effector orientation, can be executed directly. Otherwise, trajectories can be generated for further fine-tuning, where the covariance of the trajectory distribution is considered when generating samples.

C. Policy Execution and Reinforcement Learning Refinement

The synthesized trajectory $\mathbb{E}[\tau]$ can be executed directly as an open-loop plan π_τ . However, in general, τ provides a source of demonstration data that is used both to initialize the policy and to regularize learning through loss of behavior cloning (BC). This allows the system to bootstrap reinforcement learning with informative samples that remain aligned with




Fig. 6: CrossInstruct uses the accurately pointed keypoints to enable accurate reasoning of orientation. Here we see an example of the robot pulling a Jenga block, which requires an accurate movement direction.




Fig. 7: Representative tasks from the RL Bench [27] benchmark used in our evaluations.

the cross-modal instructions.

To this end, we adopt Twin Delayed DDPG (TD3) [35] augmented with Behavior Cloning. We first construct a data set of trajectories, which can be achieved by executing τ in simulation, yielding a set of state–action pairs. The robot manipulator’s state is described by the end-effector pose and a binary gripper state, and the action is given by the difference in end-effector pose and gripper state. By sampling $p(\tau|\xi_1, \xi_2, \mathcal{P})$ over a set of problem setups, we arrive at the dataset,

$$\mathcal{D} = \{(s_t, a_t)\}_{t=1}^H. \quad (11)$$

This dataset serves two functions. First, \mathcal{D} is used to pre-train the policy π by cloning behavior, providing a strong initialization that accelerates learning. Second, \mathcal{D} is retained throughout reinforcement learning, where it contributes to the BC regularization term in the actor loss. The TD3+BC actor loss is

$$\mathcal{L}_{\text{actor}} = \lambda \mathbb{E}_{(s,a) \sim \mathcal{D}} [\|\pi(s) - a\|^2] - (1 - \lambda) \mathbb{E}_s [Q(s, \pi(s))], \quad (12)$$

where π is the policy, Q is the critic, and $\lambda \in [0, 1]$ controls the balance between imitation and value-based optimization.

By construction, this loss ensures that the learned policy remains close to the synthesized trajectory while still improving beyond it through critic-guided updates. During training, the agent converges to a closed-loop policy π^* that is robust to disturbances, perception noise, and distributional shifts, while remaining grounded in the weak supervision provided by the original instruction. The resulting policy is capable of tackling problems where fine-grained motion is required.




Fig. 8: CrossInstruct, through hierarchical precision coupling, can generate spatially-accurate motions, enabling the robot to precisely pick up and place the basketball into the hoop.

V. EMPIRICAL EVALUATIONS

We rigorously evaluate CrossInstruct in both simulated and real-world robot learning tasks. Our experiments are designed to address the following research questions:

- How does CrossInstruct, which integrates a hierarchical precision-coupling module, compare against conventional VLM reasoning baselines and pure reinforcement learning (RL) baselines?
- Can CrossInstruct generalize to real world setups where the demonstrations are collected in environments that differ substantially, in both appearance and embodiment, from the robot’s execution environment?
- Can CrossInstruct be used to generate robot motion trajectories to warm-start reinforcement learning (RL), and what are the benefits of the downstream fine-tuning in terms of convergence for the RL policy?

A. Experimental Setup

a) *Simulation*: We conduct controlled experiments in simulation using the Robot Learning Benchmark (RLBench) [27]. RLBench provides a diverse suite of goal-directed manipulation tasks, each equipped with programmatic success checks and automatic scene randomization. These properties allow us to perform rigorous seed-based evaluations of both perception and control under varying conditions. The simulator is configured with a 7-degree-of-freedom (DoF) Franka Emika Panda arm mounted on a tabletop workspace. All policies are evaluated over 20 held-out random seeds per task, which are strictly disjoint from those used for demonstration collection.

b) *Hardware*: For real-world experiments, we deploy CrossInstruct on a 6-DoF AgileX manipulator with a parallel-jaw gripper. RGB images are captured from a fixed overhead camera. Importantly, the manipulator differs in morphology, kinematics, and appearance from the Franka Panda used in simulation.

c) *Tasks*: The selected RLBench tasks span both contact-rich and precision-sensitive skills: *Basketball-in-Hoop*, *Close Drawer*, *Insert Square Block on Peg*, *Lift Numbered Block*, *Play Jenga*, *Push Button*, *Put Rubbish in Bin*, and *Slide Block to Target*. Images of the representative RLBench tasks evaluated are shown in Figure 7. These tasks require reasoning over spatial layouts, geometric constraints, and trajectory shaping, stressing the coordination of end-effector (EE) position, orientation, and gripper actuation. Performance is reported as the task success rate, averaged over held-out seeds.

TABLE I: Comparison of CrossInstruct with a VLM reasoning baseline [24] and pure RL approaches (SAC, TD3) on RL Bench tasks. Numbers denote success rates (higher is better).

Method	basketball	peg	close drawer	slide block
CrossInstruct	0.90	0.25	0.90	0.90
VLM-Reasoning	0.00	0.20	0.45	0.20
Pure RL – SAC	0.00	0.00	0.95	0.10
Pure RL – TD3	0.00	0.00	0.40	0.00

	jenga	lift block	rubbish	push button
CrossInstruct	0.55	0.95	1.00	0.95
VLM-Reasoning	0.00	0.00	0.00	0.30
Pure RL – SAC	0.00	0.00	0.00	0.05
Pure RL – TD3	0.00	0.00	0.00	0.00

B. CrossInstruct Setup Details

For each RL Bench task, we evaluate the performance of CrossInstruct to execute on cross-modal instructions, without additional fine-tuning. We collect a single sketched demonstration on one random seed. A human annotator overlays freehand geometry (curves, arrows, boundaries) and scribbled text directly on task images. The CrossInstruct reasoning model consumes this mixed-modality input, infers the intent of the task, and identifies precision-critical waypoints. For each evaluation seed, we follow the implementation details in section IV to produce an ordered 3D waypoint sequence, along with end-effector orientations and gripper actions generated by the reasoning model. The trajectory is tracked to the closest joint-space trajectory by leveraging a lower-level controller, or alternatively with reactive motion primitives [36], [37], preserving sketch-implied trajectory shaping and collision clearance.

C. Baselines and Comparisons

To analyze the performance of CrossInstruct, we compare it against the following sets of baselines on a range of tasks from RL Bench.

VLM-Reasoning (no precision coupling): This baseline follows the same pipeline and uses the same sketched demonstration and evaluation seeds as CrossInstruct. However, the reasoning model directly draws trajectories over the task images and outputs EE poses and gripper actions without invoking the pointing VLM. This removes pixel-level keypoint supervision, allowing us to isolate the contribution of the hierarchical precision-coupling module.

Reinforcement Learning (RL): We train two widely used off-policy RL algorithms: Twin Delayed Deep Deterministic Policy Gradient (TD3) [35] and Soft Actor-Critic (SAC) [38]. Both agents are trained with sparse binary rewards and a budget of one million environment steps per task. TD3 is an actor-critic method that stabilizes deterministic policy gradients with clipped double Q-learning and delayed policy updates. SAC is an entropy-regularized method that learns stochastic policies to maximize exploration and robustness. We use their standard implementations without behavior cloning (BC) priors.

All baselines are evaluated on the held-out seeds. This suite of experiments allows us to (i) quantify the gains from precision coupling (CrossInstruct vs. VLM-Reasoning), and



Fig. 9: Without hierarchical precision coupling, directly seeking the VLM reasoning model to provide trajectories often leads to the robot not adequately reaching the object that we need to interact with, such as the button, basketball and Jenga block.

(ii) contrast instruction-driven learning with exploration-driven RL, i.e., CrossInstruct and VLM compared to TD3 and SAC.

D. Real-World Setup

To assess robustness under domain and embodiment shift, we qualitatively evaluate CrossInstruct on two real-world tabletop tasks: **Place Cups:** Four colored cups and four colored plates are scattered on a tabletop. The robot must match each cup with the plate of the corresponding color, while avoiding collisions and mismatches. Success requires inferring the latent color matching rule from the demonstration, selecting grasp-friendly waypoints, and executing trajectories robust to object displacement. **Saw Block:** The robot must grasp a hand saw and perform three full back-and-forth strokes along a marked line. This task stresses tool use, orientation control along a constrained path, and correct interpretation of repetition cycles (called '3x' in the demonstration).

For both tasks, demonstrations are collected on images captured in setups with different backgrounds, lighting, and end-effector appearances from the deployment environment. This directly tests CrossInstruct's ability to bind abstract, sketched intent to new scenes without additional training or domain adaptation. Images of these tasks can be seen in fig. 11.

E. Results and Analysis

a) Simulation: CrossInstruct consistently outperforms both TD3 and SAC across nearly all RL Bench tasks. The gains are particularly pronounced in precision-sensitive tasks such as *Basketball-in-Hoop* and *Push Button*, where pixel-level keypoint grounding prevents small misalignments that would otherwise cause failure. A qualitative rollout for *basketball-in-hoop* is shown in Figure 8, which illustrates spatially accurate pick-and-place into the hoop enabled by hierarchical precision coupling. A notable exception is *Close Drawer*, a short-horizon task with modest precision demands, where exploration-based RL methods occasionally succeed.

Relative to pure VLM-Reasoning, CrossInstruct demonstrates clear advantages in cluttered or low-contrast scenes, such as *Basketball-in-Hoop* and *Play Jenga*, where semantic sketching alone is insufficient to localize exact geometric keypoints. Interestingly, for *square block on peg*, the reasoning-only variant performs comparably, as the bright, simple visual cues make peg-hole delineation achievable. However, when there exists multiple objects with similar colors, such as the setup illustrated in fig. 10, the absence of the smaller pointing model can lead to erroneous results. The VLM-reasoning baseline occasionally fails due to the spatial coordinates being inaccurate, and the robot not reaching desired poses,



Fig. 11: Generalizing Cross-modal instructions to new setups, with the *Place Cups* task shown in the left and the *Saw Block* task on the right. An image with an instruction is overlaid, and we observe that the setups that the instructions are defined over are visually different from the execution setup, highlighting the generalization exhibited. We can also provide instructions such as “repeat 3x”.




Fig. 12: We sample trajectories from CrossInstruct to initialize policies for RL training. Returns against the number of steps shown (left: Jenga; right: Peg). As the Jenga task uses a sparse binary reward, training from scratch gives no nonzero returns. For the Peg task, we illustrate the performance of baseline RL methods trained from scratch against our TD3+BC approach. We observe that CrossInstruct provides an effective generative model to initialize RL training.

with notable examples in fig. 9. These findings suggest that precision coupling is most valuable when object boundaries are ambiguous, small, or occluded.




Fig. 10: Trajectories inferred with VLM-reasoning, without precision coupling can often mistakenly spatially-ground the objects in similar colored objects. On the right, both the peg and the square are blue, causing an erroneous trajectory (red). This does not happen in the left figure where blue pegs are not present.

the “repeat 3x” instruction, maintaining orientation along the constrained sawing path. Both tasks are completed without depth sensing or closed-loop replanning, highlighting CrossInstruct’s ability to bind abstract intent to real-world dynamics. Example cross-modal instructions, along with the generalized setup, which differs from the environment in which the instructions were given, are shown in fig. 11. Overall, the results demonstrate that CrossInstruct effectively combines high-level semantic reasoning with low-level geometric precision, yielding robust instruction-following policies that generalize across tasks, seeds, and embodiments.

F. CrossInstruct as a Generative Model For Reinforcement Learning Initialization

While CrossInstruct produces executable trajectories, the broader advantage lies in its ability to serve as a *generative*

b) Real World:

CrossInstruct generalizes robustly to physical setups, successfully executing mixed sketch-text instructions despite significant differences in visual appearance and manipulator kinematics. In *Place Cups*, the system correctly infers the color-matching rule, plans grasp-friendly motions, and executes without collisions. In *Saw Block*, it respects

model to initialize and improve downstream reinforcement learning (RL). By sampling from the trajectory distribution $p(\tau | \mathcal{I}, \mathcal{V}, \mathcal{P})$, we obtain diverse, yet semantically consistent, rollouts that can be directly integrated into a reinforcement learning (RL) pipeline. We then follow the TD3+BC approach outlined in section IV-C, allowing the trajectories to provide additional regularization.

We evaluated this generative initialization on two representative RLbench tasks: *Play Jenga* and *Insert Square Block on Peg*. Both of these tasks are difficult and require precise motions. We run 40k RL steps for the *Jenga* task, and 400k for the *Peg* tasks. We additionally train policies on these tasks from scratch using comparison approaches TD3 [35], SAC [38], and PPO [39]. In the Jenga task, the reward is designed as a binary “success” or “failure” signal, and our CrossInstruct-initialized policies reliably learn to dislodge the correct block, converging to a success rate around 90%, while training from scratch cannot produce a nonzero return. In the peg-insertion task, CrossInstruct significantly accelerates learning, whereas the baseline methods cannot reach comparable returns, even after 400k steps. The returns for both the Jenga and Peg tasks are illustrated in fig. 12. This shows that CrossInstruct not only produces executable plans, but also acts as a low-cost trajectory generator to reduce sample complexity in reinforcement learning. Crucially, stochastic sampling from $p(\tau)$ injects natural variability in the demonstrations, improving the robustness of the policy.

Experiments Summary: In eight RLbench tasks, CrossInstruct substantially outperforms both a VLM-reasoning baseline (without precision coupling) and pure RL (SAC/TD3) trained from scratch, confirming the benefit of hierarchical precision coupling and multi-view lifting. In both simulated and real-world settings, CrossInstruct executes mixed sketch-text instructions without fine-tuning, despite substantial differences in appearance, correctly binding abstract intent, such as color matching and “repeat 3x” tool use, to robust motions.

VI. CONCLUSIONS AND FUTURE WORK

In this work, we introduce *CrossInstruct*, a new paradigm for teaching robots through *cross-modal instructions*. These are rough sketches and textual labels provided directly on images of an environment, which can be different from the workspace of the robot. This representation avoids the high cost of collecting kinesthetic or teleoperated demonstrations while still conveying rich information about the intended behavior. Our framework achieves this by introducing a hierarchical precision coupling module that integrates a large

reasoning vision–language model with a lightweight, fine-tuned pointing model. Beyond execution, we showed that CrossInstruct acts as a generative model for robot trajectories, enabling effective downstream reinforcement learning.

A key avenue for future inquiry is to enable humans to provide cross-modal information to correct the robot’s behavior after executing motions. Here, the human will be in the loop viewing videos of robot execution and providing cross-modal instructions to iteratively shape and correct the robot’s behavior. Additionally, integrating large pre-trained models, such as DUST3R [40], VGER [41], or VGGT [42], to enable synthetic camera views from novel camera poses, thereby reducing the burden to physically position the cameras.

REFERENCES

- [1] H. Ravichandar, A. S. Polydoros, S. Chernova, and A. Billard, “Recent advances in robot learning from demonstration,” *Annual review of control, robotics, and autonomous systems*, 2020.
- [2] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, R. Tedrake, and S. Song, “Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots,” in *Proceedings of Robotics: Science and Systems (RSS)*, 2024.
- [3] W. Zhi, T. Zhang, and M. Johnson-Roberson, “Instructing robots by sketching: Learning from demonstration via probabilistic diagrammatic teaching,” in *IEEE International Conference on Robotics and Automation*, 2024.
- [4] P. Yu, A. Bhaskar, A. Singh, Z. Mahammad, and P. Tokekar, “Sketch-to-skill: Bootstrapping robot learning with human drawn trajectory sketches,” *arXiv preprint arXiv:2503.11918*, 2025.
- [5] S. A. Mehta, H. Nemlekar, H. Sumant, and D. P. Losey, “L2d2: Robot learning from 2d drawings,” *arXiv preprint arXiv:2505.12072*, 2025.
- [6] A. Singh, A. Bhaskar, P. Yu, S. Chakraborty, R. Dasyam, A. Bedi, and P. Tokekar, “Varp: Reinforcement learning from vision-language model feedback with agent regularized preferences,” *arXiv preprint arXiv:2503.13817*, 2025.
- [7] Q. Sun, W. Zhi, T. Zhang, and M. Johnson-Roberson, “Teaching robots where to go and how to act with human sketches via spatial diagrammatic instructions,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- [8] Q. Dong, L. Li, D. Dai, C. Zheng, J. Ma, R. Li, H. Xia, J. Xu, Z. Wu, B. Chang, X. Sun, L. Li, and Z. Sui, “A survey on in-context learning,” in *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2024.
- [9] S. M. LaValle, *Planning Algorithms*. USA: Cambridge University Press, 2006.
- [10] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, “Chomp: Gradient optimization techniques for efficient motion planning,” in *IEEE International Conference on Robotics and Automation*, 2009.
- [11] T. Lai, W. Zhi, T. Hermans, and F. Ramos, “Parallelised diffeomorphic sampling-based motion planning,” in *Conference on Robot Learning (CoRL)*, 2021.
- [12] A. Paraschos, C. Daniel, J. Peters, and G. Neumann, “Probabilistic movement primitives,” in *Proceedings of the 26th International Conference on Neural Information Processing Systems*, 2013.
- [13] P. Florence, C. Lynch, A. Zeng, O. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson, “Implicit behavioral cloning,” in *Conference on Robot Learning (CoRL)*, 2021.
- [14] W. Zhi, I. Akinola, K. van Wyk, N. Ratliff, and F. Ramos, “Global and reactive motion generation with geometric fabric command sequences,” in *IEEE International Conference on Robotics and Automation, ICRA*, 2023.
- [15] A. Y. Ng and S. J. Russell, “Algorithms for inverse reinforcement learning,” in *Proceedings of the Seventeenth International Conference on Machine Learning*, 2000.
- [16] K. Tanada, Y. Iwanaga, M. Tsuchinaga, Y. Nakamura, T. Mori, R. Sakai, and T. Yamamoto, “Sketch-moma: Teleoperation for mobile manipulator via interpretation of hand-drawn sketches,” *arXiv*, 2024.
- [17] W. Zhi, H. Tang, T. Zhang, and M. Johnson-Roberson, “Teaching periodic stable robot motion generation via sketch,” *IEEE Robotics and Automation Letters*, 2025.
- [18] P. Yu, A. Bhaskar, A. Singh, Z. Mahammad, and P. Tokekar, “Sketch-to-skill: Bootstrapping robot learning with human drawn trajectory sketches,” in *Proceedings of Robotics: Science and Systems (RSS)*, 2025.
- [19] P. Sundaresan, Q. Vuong, J. Gu, P. Xu, T. Xiao, S. Kirmani, T. Yu, M. Stark, A. Jain, K. Hausman, D. Sadigh, J. Bohg, and S. Schaal, “Rt-sketch: Goal-conditioned imitation learning from hand-drawn sketches,” in *Conference on Robot Learning*.
- [20] Z. Yang, C. Garrett, T. Lozano-Pérez, L. P. Kaelbling, and D. Fox, “Guiding long-horizon task and motion planning with vision language models,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2025.
- [21] M. Shridhar, L. Manuelli, and D. Fox, “Perceiver-actor: A multi-task transformer for robotic manipulation,” in *Proceedings of the 6th Conference on Robot Learning (CoRL)*, 2022.
- [22] M. Ahn *et al.*, “Do as i can, not as i say: Grounding language in robotic affordances,” in *Conference on Robot Learning*, 2022.
- [23] W. Huang, C. Wang, R. Zhang, Y. Li, J. Wu, and L. Fei-Fei, “Voxposer: Composable 3d value maps for robotic manipulation with language models,” in *Proceedings of The 7th Conference on Robot Learning*, 2023.
- [24] Y. Li, Y. Deng, J. Zhang, J. Jang, M. Memmel, R. Yu, C. Garrett, F. Ramos, D. Fox, A. Li, A. Gupta, and A. Goyal, “Hamster: Hierarchical action models for open-world robot manipulation,” *arXiv preprint arXiv:2502.05485*, 2025.
- [25] K. Fang, F. Liu, P. Abbeel, and S. Levine, “Moka: Open-world robotic manipulation through mark-based visual prompting,” *Robotics: Science and Systems (RSS)*, 2024.
- [26] W. Yuan, J. Duan, V. Blukis, W. Pumacay, R. Krishna, A. Murali, A. Mousavian, and D. Fox, “Robopoint: A vision-language model for spatial affordance prediction in robotics,” in *Proceedings of The 8th Conference on Robot Learning*, 2025.
- [27] S. James, Z. Ma, D. R. Arrojo, and A. J. Davison, “Rlbench: The robot learning benchmark & learning environment,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, 2020.
- [28] OpenAI, “Openai o3 and o4-mini system card,” OpenAI, Tech. Rep., 2025.
- [29] M. D. *et al.*, “Molmo and pixmo: Open weights and open data for state-of-the-art multimodal models,” *arXiv preprint arXiv:2409.17146*, 2024.
- [30] W. Zhi, H. Tang, T. Zhang, and M. Johnson-Roberson, “Unifying representation and calibration with 3d foundation models,” *IEEE Robotics and Automation Letters*, 2024.
- [31] R. Horaud and F. Dornaika, “Hand-eye calibration,” *I. J. Robotic Res.*, 1995.
- [32] W. Zhi, H. Tang, T. Zhang, and M. Johnson-Roberson, “Simultaneous geometry and pose estimation of held objects via 3d foundation models,” *IEEE Robotics and Automation Letters*, 2024.
- [33] R. Senanayake, “The role of predictive uncertainty and diversity in embodied ai and robot learning,” *arXiv preprint arXiv:2405.03164*, 2024.
- [34] J. Fong, M. Wrenninge, C. Kulla, and R. Habel, “Production volume rendering: Siggraph 2017 course,” in *ACM SIGGRAPH 2017 Courses*, ser. SIGGRAPH ’17, 2017.
- [35] S. Fujimoto, H. van Hoof, and D. Meger, “Addressing function approximation error in actor-critic methods,” in *International Conference on Machine Learning*, 2018.
- [36] N. D. Ratliff, J. Issac, D. Kappler, S. Birchfield, and D. Fox, “Riemannian motion policies,” 2018.
- [37] W. Zhi, T. Lai, L. Ott, E. V. Bonilla, and F. Ramos, “Learning efficient and robust ordinary differential equations via invertible neural networks,” in *International Conference on Machine Learning, ICML*, 2022.
- [38] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *International conference on machine learning*, 2018.
- [39] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *CoRR*, 2017.
- [40] S. Wang, V. Leroy, Y. Cabon, B. Chidlovskii, and J. Revaud, “Dust3r: Geometric 3d vision made easy,” in *CVPR*, 2024.
- [41] W. Zhi, Z. Ma, T. Zhang, and M. Johnson-Roberson, “From single images to motion policies via video-generation environment representations,” *arXiv preprint arXiv:2505.19306*, 2025.
- [42] J. Wang, M. Chen, N. Karaev, A. Vedaldi, C. Rupprecht, and D. Novotny, “Vgggt: Visual geometry grounded transformer,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.