

# Continuous Vision-Language-Action Co-Learning with Semantic-Physical Alignment for Behavioral Cloning

Xiuxiu Qi<sup>1,2</sup>, Yu Yang<sup>3</sup>, Jiannong Cao<sup>2\*</sup>, Luyao Bai<sup>2</sup>,  
Chongshan Fan<sup>1</sup>, Chengtai Cao<sup>4</sup>, Hongpeng Wang<sup>1\*</sup>

<sup>1</sup>The College of Artificial Intelligence & Shenzhen Research Institute, Nankai University, Tianjin, China.

<sup>2</sup>Department of Computing, The Hong Kong Polytechnic University, Hong Kong SAR, China.

<sup>3</sup>Centre for Learning, Teaching and Technology, The Education University of Hong Kong, Hong Kong SAR, China.

<sup>4</sup>Department of Computer Science, City University of Hong Kong, Hong Kong SAR, China.

## Abstract

Language-Conditioned Manipulation (LCM) facilitates human-robot interaction via Behavioral Cloning (BC), which learns control policies from human demonstrations and serves as a cornerstone of embodied AI. Overcoming compounding errors in sequential action decisions remains a central challenge to improving BC performance. Existing approaches mitigate compounding errors through data augmentation, expressive representation, or temporal abstraction. However, they suffer from physical discontinuities and semantic-physical misalignment, leading to inaccurate action cloning and intermittent execution. In this paper, we present Continuous vision-language-action Co-Learning with Semantic-Physical Alignment (CCoL), a novel BC framework that ensures temporally consistent execution and fine-grained semantic grounding. It generates robust and smooth action execution trajectories through continuous co-learning across vision, language, and proprioceptive inputs (i.e., robot internal states). Meanwhile, we anchor language semantics to visuomotor representations by a bidirectional cross-attention to learn contextual information for action generation, successfully overcoming the problem of semantic-physical misalignment. Extensive experiments show that CCoL achieves an average 8.0% relative improvement across three simulation suites, with up to 19.2% relative gain in human-demonstrated bimanual insertion tasks. Real-world tests on a 7-DoF robot further confirm CCoL’s generalization under unseen and noisy object states.

## Introduction

Language-Conditioned Manipulation (LCM) is an emerging field of research in embodied AI and a stepping stone toward artificial general intelligence (Stepputtis et al. 2020; Karamcheti et al. 2023; Pan et al. 2025). This domain seeks to develop an embodied agent to understand its surrounding environments and perform complex, goal-directed behaviors based on human instructions. As an imitation learning paradigm, Behavioral Cloning (BC) bridges high-level task instructions with low-level robotic control (Li et al. 2024a; Kelly et al. 2019; Wang et al. 2024b). Traditional BC methods mimic expert demonstrations through state-action mappings, enabling significant robotic applications, from self-driving vehicles (Codevilla et al. 2019) to robot manipulation (Cui

and Trinkle 2021). For LCM tasks, BC extends this paradigm by incorporating language as an additional modality, allowing robots to generalize executable actions via multimodal inputs. This integration is particularly critical in unstructured environments where tasks demand precise alignment between semantic intent (e.g., object affordances) and physical execution (e.g., end-effector trajectories).

Despite its promise, BC suffers from compounding errors caused by sequential error propagation, resulting in a covariate shift (Ingram et al. 2023; Mehta et al. 2025). Semantic-physical misalignment between multimodal inputs amplifies these errors (e.g., misaligned object affordances or trajectory deviations), leading to catastrophic failures in long-horizon LCM tasks as even minor inaccuracies accumulate quadratically over time (Block et al. 2023). To address this challenge, prior works have explored three primary directions: (1) data augmentation techniques such as noise injection and synthetic data generation (Deshpande et al. 2024), which enhance training diversity and improve error approximation, and interactive correction methods that iteratively collect expert feedback during rollout (Hoque et al. 2023); (2) expressive representation spaces that aim to minimize single-step prediction errors by leveraging semantically fused multimodal features (Ke, Gkanatsios, and Fragkiadaki 2025); (3) temporal abstraction methods that reduce decision-making horizon by treating long action sequences as single, high-level primitives (Shi et al. 2023; Lu et al. 2024).

However, existing methods still face two key multimodal grounding challenges impacting action cloning and execution coherence. First, physical discontinuities arise from discrete action modeling paradigms, such as temporal abstraction (Fu, Zhao, and Finn 2024), which disrupt essential motion continuity. For instance, abrupt dual-arm waypoint transitions during peg insertion tasks result in jerky trajectories with non-smooth acceleration profiles, ultimately causing failures in long-horizon tasks due to incoherent action execution. Second, semantic-physical misalignment occurs when high-level semantic goals fail to guide physical actions accurately. Static fusion methods like R3M (Nair et al. 2022) globally align language and vision but overlook stepwise semantic adaptation, resulting in inaccurate action cloning. For example, executing “place the cup on the shelf” requires the robot to dynamically shift attention from the cup during grasping to the shelf during placement.

\*Corresponding Authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

To address these challenges, we present a novel BC framework, **Continuous Vision-Language-Action Co-Learning with Semantic-Physical Alignment (CCoL)**<sup>1</sup>. CCoL introduces continuous co-learning across vision, language, and proprioceptive inputs to ensure temporally consistent latent trajectories, while stepwise anchoring language semantics to visuomotor representations to enable semantic-to-physical correspondence. Specifically, the multimodal continuous co-learning mechanism captures proprioceptive dynamics in latent space by Neural Ordinary Differential Equations (NeuralODEs) (Biloš et al. 2021; Li et al. 2024b) and enables joint fusion across vision, language, and proprioceptive inputs. Meanwhile, we design a cross-modal semantic-physical alignment mechanism that considers language as a flexible target specification and progressively anchors language semantics to visuomotor representations via a bidirectional cross-attention mechanism. Finally, a goal-conditioned decoder generates action sequences from enriched fusion representations, with a hybrid loss jointly optimizing action accuracy and latent state smoothness for stable training and physically feasible outputs. Our contributions are as follows:

- We propose a multimodal continuous co-learning mechanism integrating vision, language, and proprioception to model temporal dependencies in latent space, ensuring smooth and consistent transitions across action states.
- We introduce a cross-modal semantic-physical alignment module that anchors linguistic specifications to visuomotor representations at each step, enabling fine-grained semantic adaptation during task execution.
- We conduct extensive experiments on three simulation suites, showing that CCoL achieves SOTA performance and proves real-world effectiveness with a 7-DoF robot.

## Preliminaries

LCM leverages BC methods to infer executable behaviors from heterogeneous inputs, enabling precise and generalizable motor control (Stepputtis et al. 2020; Karamcheti et al. 2023; Li et al. 2024c). Formally, we assume access to a dataset  $\mathcal{D} = \{l, \tau_i\}_{i=1}^N$  of  $N$  expert demonstrations, where each demonstration has a sequence of state-action pairs over horizon  $H$ ,  $\tau = \{(s_1, a_1), (s_2, a_2), \dots, (s_H, a_H)\}$ , aligned with a task goal  $l \in L$  (e.g., “slide the red cube”). The states  $s_t = (o_t, r_t) \in S$  integrate current visual observation  $o_t$  and robot proprioceptive data  $r_t$ , while actions  $a_t \in A$  represent motor commands like end-effector positions or joint velocities. Our goal is to learn a policy  $\pi_\theta : S \times L \rightarrow A$ , which maps states  $S$  and natural language instructions  $L$  to actions  $A$ . We optimize this policy  $\pi_\theta$  by minimizing the negative log-likelihood of the observed expert actions:

$$\mathcal{L}_{\pi_\theta} = \mathbb{E}_{(\tau_i, l) \sim \mathcal{D}} \left[ \sum_{t=1}^H -\log \pi_\theta(a_t | s_t, l) \right]. \quad (1)$$

While LCM extends traditional BC with language guidance, it is still subject to critical *compound error* limitations.

<sup>1</sup>Project website is available at <https://qhemu.github.io/CCoL/>.

**Single-Step Error Propagation.** Mismatched vision-action mappings induce stepwise prediction errors, causing a covariate shift. Let  $a'_t = \pi_\theta(s_t, l)$  denote predicted action. Stepwise error  $\epsilon = \|a'_t - a_t\|^2$  propagates through the transition dynamics. Since transitions are sequential, the cumulative error grows quadratically with  $H$  rather than linearly:  $H\epsilon + (H-1)\epsilon + \dots + \epsilon \sim O(H^2\epsilon)$  (Xu, Li, and Yu 2020).

**Discontinuous Action Sequences.** Discretized action predictions (e.g., time abstraction) induce *higher-order physical infeasibility* through piecewise-constant control signals. These abrupt transitions violate differential continuity constraints (e.g., non-smooth acceleration profiles, discontinuous contact forces (Righetti et al. 2013)). Our key observation is that physical feasibility requires *continuous temporal coherence* beyond adjacent actions. We propose to regularize the action trajectory through dynamical consistency:

$$\mathcal{E}_{\text{disc}} = \int_t^{t+\Delta t} \rho \left( \frac{da}{dt}, \frac{d^2a}{dt^2} \right) dt, \quad (2)$$

where  $\rho(\cdot)$  encodes physical constraints. To address both single-step inaccuracies and temporal discontinuity, we propose a unified optimization framework:

$$\mathcal{L}_{\pi_\theta} = \mathbb{E}_{(\tau_i, l) \sim \mathcal{D}} \sum_{t=1}^H \left( \underbrace{\|\pi_\theta(s_t, l) - a_t\|_2^2}_{\text{Stepwise Loss}} + \underbrace{\mathcal{E}_{\text{disc}}}_{\text{Disc Penalty}} \right), \quad (3)$$

where the stepwise loss ensures local alignment between learned and expert actions, and the discontinuity penalty regularizes action sequences to ensure physical feasibility.

## Methodology

As depicted in Fig. 1, CCoL primarily consists of two key components: *Multimodal Continuous Co-Learning (MCC)* and *Cross-Modal Semantic-Physical Alignment (CSA)*.

### Context-Aware Representation Learning

We first introduce context-aware representation learning to ensure that each modality is independently encoded while preserving its unique characteristics.

**Vision Encoder.** The vision encoder processes RGB-D video frames  $o_t \in \mathbb{R}^{H \times W \times 3}$  using a Vision Transformer (ViT) (Han et al. 2022). The encoder extracts spatially grounded features  $x_t = \text{ViT}(o_t), x_t \in \mathbb{R}^{d_v}$ , capturing spatial cues critical for grounding language in the visual scene.

**Text Encoder.** To process language instructions, the text encoder employs a RoBERTa model (Liu et al. 2019). Instructional texts are tokenized and encoded into contextualized embeddings:  $\hat{l}_t = \text{RoBERTa}(l)$ , where  $\hat{l}_t \in \mathbb{R}^{d_l}$ .

**Proprioceptive Encoder.** The proprioceptive encoder processes robotic internal states, including position sequences  $r_t \in \mathbb{R}^k$ . Inspired by (Zhao et al. 2023), we implement a Conditional Variational Autoencoder (CVAE) to map these inputs into a latent motion pattern space. The inputs  $r_t$  are first processed through linear transformations and concatenated with a [CLS] token. Sinusoidal positional encodings are added to incorporate sequential order information. The

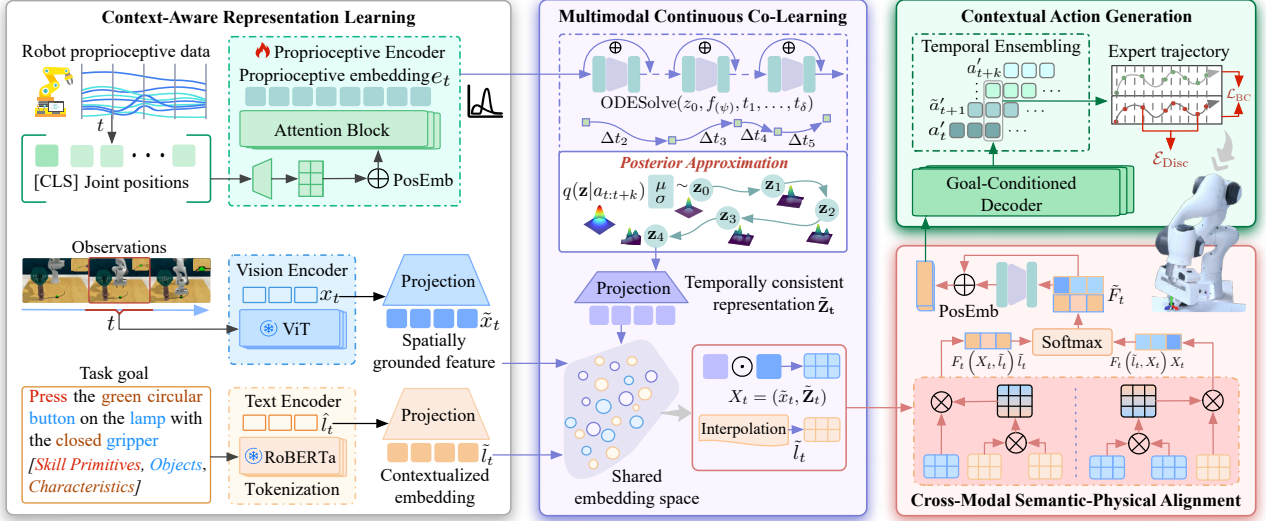


Figure 1: Overview of the CCoL framework. MCC leverages dynamic proprioceptive modeling to capture temporal evolution and maps vision, language, and proprioception into a shared latent space (**purple frame**). CSA synchronizes stepwise semantic information across modalities (**red frame**) to enable the generation of contextually and physically feasible action sequences.

concatenated sequence is then processed by a Transformer (Vaswani et al. 2017) or a Temporal Convolutional Network (TCN) (Lea et al. 2017), producing the proprioceptive embedding  $e_t = \text{CVAE}([\text{CLS}; r_t])$ , where  $e_t \in \mathbb{R}^{d_a}$ .

### Multimodal Continuous Co-Learning

Conventional BC approaches often suffer from temporally fragmented action sequences, where decoupled per-step predictions fail to account for underlying motion dynamics. Such limitations lead to jerky trajectories (e.g., high-jerk robotic arm movements) or kinematically invalid transitions (e.g., sudden end-effector orientation flips). Inspired by the notion of co-learning from (Liu, Chen, and Liu 2024), we introduce a novel *Multimodal Continuous Co-learning* (MCC) that uses NeuralODEs to capture the continuous evolution of proprioceptive embeddings. By integrating multimodal embeddings, we construct a shared latent space that captures temporal coherence and task-specific adaptability. Specifically, we project proprioceptive embedding at [CLS] token to predict the parameters of diagonal Gaussian distribution:

$$\begin{aligned} \mu &= W_\mu e_t [\text{CLS}] + b_\mu, \\ \log \sigma^2 &= W_\sigma e_t [\text{CLS}] + b_\sigma, \end{aligned} \quad (4)$$

where  $W_\mu, W_\sigma$  and  $b_\mu, b_\sigma$  are learnable parameters. The latent state  $z_0$  is initialized via a conditional variational encoder with the reparameterization trick:

$$z_0 = \mu + \varepsilon \cdot \exp\left(\frac{1}{2} \log \sigma^2\right), \quad \varepsilon \sim \mathcal{N}(0, 1), \quad (5)$$

where  $\varepsilon$  is a noise vector sampled from a standard normal distribution  $\mathcal{N}(0, 1)$ . The NeuralODEs models the continuous evolution of this latent state as the solution to an initial value problem defined by a differential equation:

$$z(t_\delta) = z_0 + \int_0^{t_\delta} f(z(t), t; \psi) dt, \quad (6)$$

where  $t_\delta$  represents the elapsed integration time, and  $f(z(t), t; \psi)$  is a residual MLP that learns the derivatives of the latent state with respect to time. Parameters  $\psi$  are trained to ensure that latent trajectories match the underlying proprioceptive dynamics. We solve this differential equation numerically using adaptive-step Dormand-Prince solvers via `odeint`, which computes  $\mathbf{Z}_t = \text{odeint}(f, z_0, t)$  as the latent states over discrete time points  $t$ . These latent trajectories replace stepwise proprioceptive features  $e_t$ , providing temporally consistent representations that mitigate fragmentation and discontinuities in conventional encoders.

To enable joint reasoning across modalities, visual features  $x_t$ , language features  $\hat{l}_t$ , and proprioceptive features  $\mathbf{Z}_t$  are projected into a shared embedding space:

$$\begin{aligned} \tilde{x}_t &= \text{ReLU}(W_v x_t + b_v), \\ \tilde{l}_t &= \text{ReLU}(W_l \hat{l}_t + b_l), \\ \tilde{\mathbf{Z}}_t &= \text{ReLU}(W_a \mathbf{Z}_t + b_a). \end{aligned} \quad (7)$$

Here,  $W_v \in \mathbb{R}^{d_v \times h}$ ,  $W_l \in \mathbb{R}^{d_l \times h}$  and  $W_a \in \mathbb{R}^{d_a \times h}$  map original features to a higher dimension  $h$ , while  $b_v, b_l, b_a \in \mathbb{R}^h$  are learnable biases. Moreover, language embeddings are upsampled to match the resolution of visual features via bilinear interpolation, enabling pixel-wise synchronization for fine-grained cross-modal fusion.

### Cross-Modal Semantic-Physical Alignment

Traditional feature fusion methods fail to dynamically align linguistic instructions with visual observations and proprioceptive states, causing inconsistencies in cross-modal grounding. To this end, we design a *Cross-Modal Semantic-Physical Alignment* (CSA) that anchors linguistic concepts to visuomotor representations at each timestep. As shown in Fig. 2, the normalized attention scores directly decode linguistic grounding into physically constrained pixel/trajectory

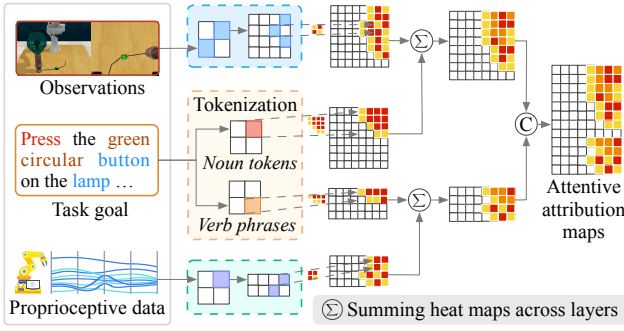


Figure 2: Illustration of attentive attribution maps (e.g., noun grounding and verb-conditioned trajectory adjustment).

attributions, ensuring semantic-to-physical correspondence. Specifically, we design an attentive attribution mapper that anchors lexical semantics to visuomotor representations at each timestep via bidirectional cross-attention. Formally, at each attention head  $\iota \in [1, n]$ , given text embedding  $\tilde{l}_t$  and joint vision-proprioceptive context  $X_t = (\tilde{x}_t, \tilde{Z}_t)$ , the attention score determines the token-wise relevance:

$$F(t)_t(\tilde{l}_t, X_t) = \text{softmax} \left( \frac{\left( W_q^{(\iota)} \tilde{l}_t \right) \left( W_k^{(\iota)} X_t \right)^T}{\sqrt{d_k}} \right),$$

$$F(t)_t(X_t, \tilde{l}_t) = \text{softmax} \left( \frac{\left( W_q^{(\iota)} X_t \right) \left( W_k^{(\iota)} \tilde{l}_t \right)^T}{\sqrt{d_k}} \right), \quad (8)$$

where  $W_q^{(\iota)}, W_k^{(\iota)} \in \mathbb{R}^{h \times d_k}$  are learned projection matrices for queries and keys,  $d_k$  is the dimension of the key vectors. These attention scores determine the alignment between linguistic tokens (e.g., verbs and nouns) and physical features, ensuring semantic-to-physical correspondence. The fused features  $\tilde{F}_t$  are calculated as follows:

$$\tilde{F}_t = \sum_n F(\tilde{l}_t, X_t) X_{t,n} + F(X_t, \tilde{l}_t) \tilde{l}_{t,n}, \quad (9)$$

where each term is broadcast across the sequence length and feature dimensions. To maintain temporal coherence, we incorporate positional encodings  $\mathbf{pos} \in \mathbb{R}^h$  into the self-attention mechanism. Specifically, the queries and keys are generated from these encodings through learned projections  $W_q^{(\iota)}$  and  $W_k^{(\iota)}$  to capture position-dependent relationships:

$$\xi_t = \sum_n \text{softmax} \left( \frac{\left( W_q^{(\iota)} \mathbf{pos}_t \right) \left( W_k^{(\iota)} \mathbf{pos}_t \right)^T}{\sqrt{d_k}} \right) \tilde{F}_{t,n}. \quad (10)$$

Here,  $\xi_t$  represents the final multimodal representation at  $t$ , incorporating both semantic alignment and temporal coherence. CSA overcomes the limitations of static fusion methods by grounding linguistic instructions in visuomotor representations via bidirectional cross-attention. Stepwise semantic anchoring aligns linguistic tokens (e.g., ‘‘button’’, ‘‘press’’) with visual regions and proprioceptive states, while positional encodings ensure temporal coherence.

## Contextual Action Generation

The goal-conditioned decoder predicts future joint positions over the next  $k$  timesteps, ensuring contextual relevance and robustness to distributional shifts. Residual connections, layer normalization, and dropout are used to stabilize training and improve output quality.

$$a'_{t:t+k} = \text{LayerNorm}(\tilde{F}_t + \text{Dropout}(\xi_t)). \quad (11)$$

Each attention block is followed by a position-wise feed-forward network, and the processed features (dimension  $k \times d_k$ ) are down-projected through an MLP into  $k \times d$ , representing the predicted action sequences for the next  $k$  timesteps. During training, the policy  $\pi_\theta$  is optimized via supervised learning on expert demonstrations, minimizing the discrepancy between predicted and expert actions. During inference, only the decoder is active, ensuring deterministic outputs for consistent evaluation. The decoder generates action sequences  $a'$  by conditioning on real-time images and robot state  $s'$  and task instructions  $l'$ , while it is executed under the environment dynamics (e.g., variations in position).

## Posterior Approximation and Optimization

To model the complex posterior distribution of the latent variable  $z$ , we follow the CVAE framework. The true posterior  $p_\psi(z|a_{t:t+k})$  is approximated by  $q_\phi(z|a_{t:t+k})$ , parameterized by a neural network. The optimization objective is to maximize the Evidence Lower Bound (ELBO), comprising reconstruction and KL divergence losses:

$$\log p(a'_{t:t+k}) \geq \underbrace{\mathbb{E}_{q(z|a'_{t:t+k})} [\log p(a'_{t:t+k}|z, s, l)]}_{\text{Reconstruction Likelihood}} - \underbrace{\text{KL}(q(z|a'_{t:t+k}) \| p(z))}_{\text{Kullback-Leibler divergence}}, \quad (12)$$

where  $p(a'_{t:t+k}|z, s, l)$  is the likelihood of the data given the latent variables, and  $p(z)$  is the prior distribution of the latent variables. By maximizing the ELBO, we effectively maximize the log probability of the data while ensuring that the variational distribution  $q(z|a'_{t:t+k})$  is close to the true posterior. The reconstruction loss ensures that decoded trajectories align with expert demonstrations, while the KL divergence regularizes the latent encoding to conform to a standard Gaussian prior,  $\mathcal{L}_{\text{BC}} = \mathcal{L}_{\text{recon}} + \mathcal{L}_{\text{KL}}$ . However, the estimated posterior  $q_\phi(z|a'_{t:t+k})$  in vanilla CVAE is restricted to a normal distribution, which limits its ability to reflect the real trajectory characteristics. By employing NeuralODEs, we enhance the posterior’s adaptability by learning dynamic representations of latent trajectories, better reflecting real-world data dynamics. Additionally, a discontinuity penalty is introduced to ensure the smooth evolution of latent states:

$$\mathcal{E}_{\text{disc}} = \sum_t \int_{t_0}^{t_0+k} \left\| \frac{dz(t)}{dt} - f(z(t), t; \psi) \right\|_2^2 dt, \quad (13)$$

where  $\frac{dz(t)}{dt}$  is the actual rate of change of latent states, and  $f(z(t), t; \psi)$  is the rate of change predicted by NeuralODE. The total loss combines the BC and a discontinuity penalty:

$$\mathcal{L} = \frac{1}{N} \sum \mathcal{L}_{\text{BC}} + \mathcal{E}_{\text{disc}}. \quad (14)$$

## Related Work

**Language-Conditioned Manipulation (LCM).** LCM enables robots to perform tasks via natural language, making it a key paradigm in embodied AI by integrating vision, language, and robotics (Guhur et al. 2023; Chen et al. 2024). A key approach to tackling LCM tasks is Behavioral Cloning (BC), which trains policies for sequential decision-making directly from expert demonstrations (Kober and Peters 2009; Foster, Block, and Misra 2024). Recent work explores vision-based representations and vision-language interaction to enhance policy learning. Voxel-based methods like PerAct (Shridhar, Manuelli, and Fox 2023) and large-scale models like RT-1 (Brohan et al. 2023) enable language grounding but suffer from high computational and parameter costs. Thus, we propose a lightweight framework that predicts actions from RGB and robot states, while preserving fine-grained linguistic-physical alignment.

**Compounding Errors in Behavioral Cloning.** BC models an action distribution given the current state at each timestep (Wang et al. 2024a). However, BC suffers from compounding errors that lead to distributional drift and unrecoverable states over long horizons (Mehta et al. 2025). Prior works include interactive corrections require expert feedback during execution and are impractical for complex teleoperation (Hoque et al. 2023), while data augmentation (Lee et al. 2015; Deshpande et al. 2024) reduces supervision but struggles with fine manipulation. Temporal abstraction methods (Zhao et al. 2023; Lu et al. 2024) discretize actions into phases, but often induce jerky motions. R3M (Nair et al. 2022) attempts to align text with vision but lacks dynamic adaptation to evolving task contexts. In contrast, we jointly optimize stepwise semantic grounding and temporal continuity, mitigating compounding errors and the rigidity in BC.

**Neural Ordinary Differential Equations.** Neural Ordinary Differential Equations (NeuralODEs) extend neural networks to model continuous-time dynamics, evolving from discrete ResNet (He et al. 2016) updates to continuous trajectories (Biloš et al. 2021). ODE-RNNs (Rubanova, Chen, and Duvenaud 2019) are designed to model irregularly sampled time series, while TrajODE (Liang et al. 2021) focuses on enhancing functionalities of recurrent neural networks. Studies (Lin et al. 2021; Li et al. 2024b) show that learnable control in NeuralODEs approximates a nonlinear dynamics-inversion controller. Although NeuralODEs have been utilized in robotic controllers, we harness them for high-level decision-making, leveraging their continuous dynamics to handle the complexities of high-dimensional visual tasks.

## Experiments

We conduct comprehensive evaluations against state-of-the-art behavior cloning baselines and ablate key components of our method. We further examine how stepwise semantic anchoring improves visual-action alignment, and how NeuralODE solver parameters affect trajectory continuity.

### Experimental Setup

**Simulation Environments and Datasets.** We evaluate CCoL on three widely used simulation suites: (1) We con-

Method	Cube Transfer			Bimanual Insertion		
	Scripted	Human	Avg.	Scripted	Human	Avg.
BCCNN	1.0	0.0	0.5	1.0	0.0	0.5
RT-1	2.0	0.0	1.0	1.0	0.0	0.5
VINN	3.0	0.0	1.5	1.0	0.0	0.5
BeT	27.0	1.0	14.0	3.0	0.0	1.5
DP	54.0	4.0	29.0	74.0	0.0	37.0
ACT	86.0	50.0	68.0	32.0	20.0	26.0
AWE	99.0	71.0	85.0	57.0	30.0	43.5
DIC	95.9	78.1	87.0	83.2	30.2	56.7
<b>CCoL</b>	<b>99.0</b>	<b>82.0</b>	<b>90.5</b>	<b>87.0</b>	<b>36.0</b>	<b>61.5</b>

Table 1: Success rate (%) on Aloha MuJoCo.

Method	LampOn	GrillMeat	Phone	OpenBottle	Avg.
BCCNN	4.0	0.0	0.0	0.0	1.0
ACT	76.0	70.3	22.0	42.7	52.8
AWE	85.7	74.3	34.7	46.3	60.3
<b>CCoL</b>	<b>93.7</b>	<b>82.3</b>	<b>44.3</b>	<b>51.7</b>	<b>68.0</b>
HDP	92.7	80.7	42.3	50.7	66.6
3DDiff	89.3	85.0	71.7	69.3	78.8
<b>CCoL<sub>3D</sub></b>	<b>97.3</b>	<b>87.3</b>	<b>76.7</b>	<b>78.3</b>	<b>84.9</b>

Table 2: Comparison on RL Bench tasks (3 seeds).

duct two bimanual collaborative tasks on Aloha MuJoCo (Zhao et al. 2023), using human and scripted demonstrations under varied object positions. (2) We employ RL Bench (James et al. 2020) for multi-scenario evaluation. (3) Following (Zeng et al. 2024), we evaluate CCoL on multi-stage Franka Kitchen tasks to assess long-horizon performance.

**Baselines.** We evaluate CCoL against representative BC methods across three categories. *Temporal modeling baselines* include BCCNN (Jang et al. 2022), RT-1 (Brohan et al. 2023), BeT (Shafiullah et al. 2022), and VINN (Pari et al. 2022), mapping states to actions via CNN, transformers, non-parametric retrieval and weighted k-nearest neighbors. ACT (Zhao et al. 2023) and AWE (Shi et al. 2023) reduce planning horizon via action chunking and waypoint abstraction. *Diffusion-based methods* like DP (Chi et al. 2025) and DIC (Xu et al. 2025) model actions via conditional denoising diffusion. HDP (Ma et al. 2024) and 3DDiff (Ke, Gkanatsios, and Fragkiadaki 2025) leverage 3D scene tokens to decompose policy into joint-level diffusion. *Representation-enhanced policies* like R3M (Nair et al. 2023), Voltron (Karamcheti et al. 2023), and MPI (Zeng et al. 2024) leverage visual representations to improve semantic grounding.

**Metrics and Implementation Details.** The success rate (%) measures task completion under defined conditions, such as transferring a red cube between grippers with 1cm clearance to avoid collisions. To ensure consistency, our model maintains the same settings as AWE (Shi et al. 2023; Ke, Gkanatsios, and Fragkiadaki 2025) for Aloha MuJoCo and RL Bench, and aligns with MPI (Zeng et al. 2024) for the Franka Kitchen setup. The model employs the SGD optimizer, starting with a learning rate of 1e-5, momentum of

Method	Backbone	Single-task				Long-horizon task			
		① Turn Knob	② Open Door	③ Open Microwave	Avg.	①+②	②+③	①+②+③	Avg.
INSUP	ResNet50	28.0	18.0	26.7	24.2	–	–	–	–
CLIP	ResNet50	26.3	13.0	24.7	21.3	–	–	–	–
R3M	ResNet50	53.3	<u>50.7</u>	59.3	54.4	25.1	29.6	15.5	23.4
Voltron	ViT-S	71.7	45.3	40.3	52.4	34.6	22.6	13.0	23.4
MPI	ViT-S	<u>83.3</u>	50.3	<u>59.7</u>	<u>64.4</u>	<u>38.2</u>	<u>31.2</u>	<u>23.3</u>	<u>30.9</u>
<b>CCoL</b>	ViT-S	<b><u>85.3</u></b>	<b><u>56.7</u></b>	<b><u>64.7</u></b>	<b><u>68.9</u></b>	<b><u>43.7</u></b>	<b><u>37.3</u></b>	<b><u>27.7</u></b>	<b><u>36.2</u></b>
MPI	ViT-B	<b>89.0</b>	<u>57.7</u>	54.0	<u>66.9</u>	<u>41.7</u>	34.7	26.3	34.2
<b>CCoL</b>	ViT-B <sub>frozen</sub>	81.7	<u>53.3</u>	<b>62.7</b>	65.9	41.6	<u>35.4</u>	<u>26.5</u>	<u>34.5</u>
<b>CCoL</b>	ViT-B	<u>88.7</u>	<b>58.3</b>	<u>58.7</u>	<b>68.6</b>	<b>45.7</b>	<b>38.5</b>	<b>30.1</b>	<b>38.1</b>

Table 3: Performance comparison on Franka Kitchen. **Bold** and underlined values mark the best and second-best performance.

Method	Cube Transfer		Bimanual Insertion	
	Scripted	Human	Scripted	Human
<b>CCoL</b>	<b>99.0</b>	<b>82.0</b>	<b>87.0</b>	<b>36.0</b>
└ w/o MCC	99.0	74.0	72.0	34.0
└ w/o CSA	99.0	73.0	74.0	35.0
└ w/o $\mathcal{E}_{disc}$	99.0	78.0	76.0	32.0
└ w/ CSA <sub>no att</sub>	98.0	72.0	70.0	28.0
└ w/ TCN	96.0	72.0	74.0	36.0
└ w/ TCN <sub>no MCC</sub>	92.0	58.0	54.0	20.0

Table 4: Ablation study of MCC and CSA variants.




Figure 3: Trajectory smoothness analysis for the shoulder joint during the bimanual insertion scripted task.

0.9, a chunking size of  $k=50$ , and a batch size of 8. The ODEsolver evaluates solutions at two discrete time points.

## Experiment Results

(1) As shown in Table 1, CCoL excels in bimanual coordination tasks, surpassing AWE by +11.8% (absolute) and DIC by +5.8% (relative) in average success rate. In particular, it achieves +19.2% relative gain over DIC in human-demonstrated bimanual insertion, showing robustness to noisy supervision and task complexity. (2) On RL Bench (Table 2), CCoL surpasses AWE by +7.7% in 2D setting, and CCoL<sub>3D</sub> outperforms 3DDiff by +6.1%, highlighting the benefit of multimodal coordination in 3D spatial domains. For a fair comparison, CCoL<sub>3D</sub> aligns its 3D representation resolution with 3DDiff by leveraging RGB-D input, a CLIP-based ResNet-50 encoder, and shared 3D tokens with

relative attention. (3) We evaluate CCoL with multiple vision backbones (Table 3). Under ViT-S (22M), it improves over MPI by +6.9% and +17.2% on single-task and long-horizon settings, respectively. With ViT-B (86M), it achieves +11.4% improvement on the long-horizon setting, and notably reaches 34.5% even with a frozen encoder, showing strong performance without vision tuning.

## Ablation Study

**Quantitative Analysis.** Table 4 reveals that w/o (i.e., without) MCC causes a 15.0% drop on bimanual insertion (scripted) due to disrupted continuity through latent proprioceptive modeling, while w/o CSA leads to a 9.0% average drop on cube transfer (human), highlighting its importance in aligning linguistic with visuomotor context. Removing  $\mathcal{E}_{disc}$  from MCC degrades performance by weakening temporal smoothness. CSA<sub>no att</sub> replacing bidirectional attention with average pooling performed worse than completely w/o CSA. Replacing attention with TCN in the proprioceptive encoder reduces success by 13.0% in the bimanual insertion (scripted) task, while further removing MCC leads to a 16.0% drop on human demos, showing that both temporal modeling and attention are essential for coherent behavior.

**Trajectory Smoothness Analysis.** To assess CCoL’s physical continuity, we compute velocity  $v(t) = \frac{da'_t}{dt} \approx a'_{t+1} - a'_t$  and acceleration  $c(t) = \frac{d^2a'_t}{dt^2} \approx a'_{t+2} - 2a'_{t+1} + a'_t$ . As shown in Fig. 3, CCoL reduces high-frequency oscillations (e.g., peak acceleration) compared to baselines. This smoothness stems from MCC’s latent proprioceptive modeling via NeuralODEs (Eq. 6), which enables continuous transitions. Compared to CCoL<sub>w/o MCC</sub>, CCoL reduces velocity fluctuations by 30.8% and acceleration fluctuations by 32.7%. Remarkably, the minimum acceleration improved 20.2%. These findings validate the ability of MCC to suppress high-frequency jitter and severe negative acceleration (sudden deceleration), resulting in smoother transitions.

**Hyperparameter Analysis.** Fig. 5 shows the effect of NeuralODE solver timesteps (2.0, 1.0, 0.5) on long-horizon task performance. Larger timesteps yield higher success rates and faster convergence, especially in cube transfer (scripted), supporting our design that coarse steps improve efficiency while preserving smooth dynamics. Notably, CCoL(2.0) shows sta-

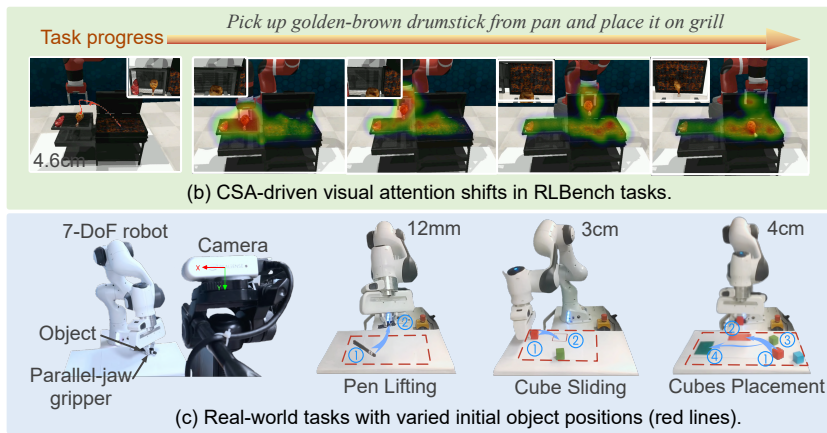
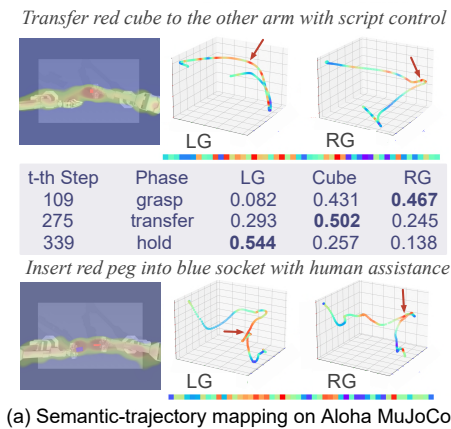


Figure 4: Attentive attribution mapping via CSA and real-world task settings with varied initial object positions.




Figure 5: Hyperparameter impact on insertion and transfer.

ble monotonic improvement, whereas smaller steps (e.g., 0.5) lead to oscillations from sensitivity to transient noise.

**Qualitative Analysis.** We analyze how linguistic cues influence visual attention and trajectory planning through CSA. Fig. 4(a) illustrates that CSA aligns nouns (e.g., “cube”, “socket”) to visual targets while verbs (e.g., “transfer”, “insert”) map to distinct trajectory patterns. We also evaluate CSA’s temporal grounding behavior by analyzing attention shifts across tasks. In cube transfer, quantitative attention scores show a clear progression from the right gripper (grasp) to the red cube (transfer) and then the left gripper (handover). A similar shift is observed in the RL Bench grillmeat task (Fig. 4(b)). This phase-to-phase transition shows that CSA dynamically adjusts attention based on semantic cues, enabling step-aware visuomotor alignment.

**Real-World Experiments.** To validate our framework in practical scenarios, we deploy it on a 7-DoF Franka Emika Panda robot equipped with an Intel RealSense D435i RGB-D camera and a parallel-jaw gripper (Fig. 4(c)). We design three manipulation tasks (pen lifting, cube sliding and cubes placement) to assess generalization under unseen (e.g., varying pen diameters) or noisy states (e.g., vase, pen holder). For training, we collect 50 kinesthetic demonstrations per task with randomized initial conditions. At test time, each task is evaluated over 15 trials under shifted conditions. As shown in Fig. 6, CCoL consistently achieves high success rates (e.g., 86.7% in cubes placement), demonstrating strong real-world generalization. Since failures are mainly caused




Figure 6: Real-world generalization under unseen states. Bars indicate success and four types of failure rates.

by localization errors, contact (grasping) failures, kinematic inconsistencies, and size adaptation issues, we report the average failure rate (FR) of each failure. We found that when picking cubes, placement errors occurred in small color areas, despite successful localization and contact. When picking up a pen, failures were caused by small diameters and insufficient lifting distances. When sliding a cube, failures stem from positioning deviations or incomplete sliding into the target area. Our model trains in 5.3 hours on an RTX 4090 GPU and runs at 0.015s ( $\pm 0.003$ s) per action sequence, meeting real-time requirements ( $\approx 67$ Hz policy frequency).

## Conclusion

We present CCoL, a novel framework that synergizes multimodal continuous co-learning and stepwise semantic-physical alignment to mitigate compounding errors in BC. By integrating NeuralODEs for smooth latent trajectory learning and bidirectional cross-attention for linguistic grounding, CCoL ensures both temporal coherence and semantic precision during task execution. Extensive simulation experiments, alongside real-world robustness, demonstrate its superiority, especially in bimanual collaborative and long-horizon manipulation tasks. Future work will extend CCoL to LLM-based methods and explore its integration with foundation models for open-world manipulation.

## Acknowledgments

This work was supported by National High Level Hospital Clinical Research Funding (Grant 2025-PUMCH-D-005), National Natural Science Foundation of China (Grant 62473213), Sustainable Development Science and Technology Special Project of Shenzhen (Grant KCXFZ20230731100900002), Shenzhen Science and Technology Program (Grant KQTD20210811090143060), and Beijing Tianjin Hebei Basic Research Cooperation Special Project (Grant 24JCZXJC00060). We also acknowledge support in part from HK RGC General Research Fund (No.: PolyU 15235424), “Research on Key Technologies for Systematic Artificial Intelligence Agents” project under China Mobile Innovation and Research Institute (No.: R24114H7), and Research Institute for Artificial Intelligence of Things, The Hong Kong Polytechnic University. This work was partially conducted at the Robotics and Embodied Intelligence (REI) Lab, The Education University of Hong Kong.

## References

- Biloš, M.; Sommer, J.; Rangapuram, S. S.; Januschowski, T.; and Günnemann, S. 2021. Neural flows: Efficient alternative to neural ODEs. *Advances in neural information processing systems*, 34: 21325–21337.
- Block, A.; Jadbabaie, A.; Pfrommer, D.; Simchowitz, M.; and Tedrake, R. 2023. Provable guarantees for generative behavior cloning: Bridging low-level stability and high-level behavior. *Advances in Neural Information Processing Systems*, 36: 48534–48547.
- Brohan, A.; Brown, N.; Carbajal, J.; Chebotar, Y.; Dabis, J.; Finn, C.; Gopalakrishnan, K.; Hausman, K.; Herzog, A.; Hsu, J.; et al. 2023. RT-1: Robotics Transformer for Real-World Control at Scale. In *Robotics: Science and Systems XIX, Daegu, Republic of Korea, July 10-14, 2023*.
- Chen, B.; Kang, J.; Zhong, P.; Liang, Y.; Sheng, Y.; and Wang, J. 2024. Embodied Contrastive Learning with Geometric Consistency and Behavioral Awareness for Object Navigation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 4776–4785.
- Chi, C.; Xu, Z.; Feng, S.; Cousineau, E.; Du, Y.; Burchfiel, B.; Tedrake, R.; and Song, S. 2025. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 44(10-11): 1684–1704.
- Codevilla, F.; Santana, E.; López, A. M.; and Gaidon, A. 2019. Exploring the limitations of behavior cloning for autonomous driving. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9329–9338.
- Cui, J.; and Trinkle, J. 2021. Toward next-generation learned robot manipulation. *Science robotics*, 6(54): eabd9461.
- Deshpande, A.; Ke, L.; Pfeifer, Q.; Gupta, A.; and Srinivasa, S. S. 2024. Data efficient behavior cloning for fine manipulation via continuity-based corrective labels. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- Foster, D. J.; Block, A.; and Misra, D. 2024. Is behavior cloning all you need? understanding horizon in imitation learning. *Advances in Neural Information Processing Systems*, 37: 120602–120666.
- Fu, Z.; Zhao, T. Z.; and Finn, C. 2024. Mobile ALOHA: Learning bimanual mobile manipulation using low-cost whole-body teleoperation. *Conference on Robot Learning, 6-9 November 2024, Munich, Germany*, 270: 4066–4083.
- Guhur, P.-L.; Chen, S.; Pinel, R. G.; Tapaswi, M.; Laptev, I.; and Schmid, C. 2023. Instruction-driven history-aware policies for robotic manipulations. In *Conference on Robot Learning*, 175–187. PMLR.
- Han, K.; Wang, Y.; Chen, H.; Chen, X.; Guo, J.; Liu, Z.; Tang, Y.; Xiao, A.; Xu, C.; Xu, Y.; et al. 2022. A survey on vision transformer. *IEEE transactions on pattern analysis and machine intelligence*, 45(1): 87–110.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition, CVPR 2016, Las Vegas, USA*, 770–778.
- Hoque, R.; Chen, L. Y.; Sharma, S.; Dharmarajan, K.; Thananjeyan, B.; Abbeel, P.; and Goldberg, K. 2023. Fleet-dagger: Interactive robot fleet learning with scalable human supervision. In *Conference on Robot Learning*, 368–380.
- Ingram, B.; Van Alten, C.; Klein, R.; and Rosman, B. 2023. Creating diverse play-style-centric agents through behavioural cloning. In *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, volume 19, 255–265.
- James, S.; Ma, Z.; Arrojo, D. R.; and Davison, A. J. 2020. Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters (RAL)*, 5(2): 3019–3026.
- Jang, E.; Irpan, A.; Khansari, M.; Kappler, D.; Ebert, F.; Lynch, C.; Levine, S.; and Finn, C. 2022. Bc-z: Zero-shot task generalization with robotic imitation learning. In *Conference on Robot Learning*, 991–1002. PMLR.
- Karamcheti, S.; Nair, S.; Chen, A. S.; Kollar, T.; Finn, C.; Sadigh, D.; and Liang, P. 2023. Language-Driven Representation Learning for Robotics. In *Robotics: Science and Systems XIX, Daegu, Republic of Korea, July 10-14, 2023*.
- Ke, T.-W.; Gkanatsios, N.; and Fragkiadaki, K. 2025. 3D Diffuser Actor: Policy Diffusion with 3D Scene Representations. *Conference on Robot Learning*, 1949–1974.
- Kelly, M.; Sidrane, C.; Driggs-Campbell, K.; and Kochenderfer, M. J. 2019. Hg-dagger: Interactive imitation learning with human experts. In *2019 International Conference on Robotics and Automation (ICRA)*, 8077–8083.
- Kober, J.; and Peters, J. 2009. Learning motor primitives for robotics. In *2009 IEEE International Conference on Robotics and Automation*, 2112–2118.
- Lea, C.; Flynn, M. D.; Vidal, R.; Reiter, A.; and Hager, G. D. 2017. Temporal convolutional networks for action segmentation and detection. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, July 21-26, 2017*, 1003–1012.
- Lee, A. X.; Lu, H.; Gupta, A.; Levine, S.; and Abbeel, P. 2015. Learning force-based manipulation of deformable objects from multiple demonstrations. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 177–184.

- Li, L.; Liang, S.; Zhu, Z.; Ding, C.; Zha, H.; and Wu, B. 2024a. Learning to optimize permutation flow shop scheduling via graph-based imitation learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 20185–20193.
- Li, M.; Bian, W.; Chen, L.; and Liu, M. 2024b. HiDeS: a higher-order-derivative-supervised neural ordinary differential equation for multi-robot systems and opinion dynamics. *Frontiers in Neurorobotics*, 18: 1382305.
- Li, X.; Liu, M.; Zhang, H.; Yu, C.; Xu, J.; Wu, H.; Cheang, C.; Jing, Y.; Zhang, W.; Liu, H.; et al. 2024c. Vision-Language Foundation Models as Effective Robot Imitators. In *International Conference on Learning Representations*.
- Liang, Y.; Ouyang, K.; Yan, H.; Wang, Y.; Tong, Z.; and Zimmermann, R. 2021. Modeling Trajectories with Neural Ordinary Differential Equations. In *Proceedings of International Joint Conference on Artificial Intelligence*, 1498–1504.
- Lin, H.; Li, B.; Zhou, X.; Wang, J.; and Meng, M. Q.-H. 2021. No need for interactions: Robust model-based imitation learning using neural ode. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 11088–11094.
- Liu, J.; Chen, C.; and Liu, M. 2024. Multi-modality co-learning for efficient skeleton-based action recognition. In *Proceedings of the 32nd ACM international conference on multimedia*, 4909–4918.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *CoRR*, abs/1907.11692.
- Lu, J.; Xia, W.; Wang, D.; Wang, Z.; Zhao, B.; Hu, D.; and Li, X. 2024. KOI: Accelerating Online Imitation Learning via Hybrid Key-state Guidance. In *Conference on Robot Learning, 6-9 November 2024, Munich, Germany*, volume 270, 3847–3865. PMLR.
- Ma, X.; Patidar, S.; Houghton, I.; and James, S. 2024. Hierarchical diffusion policy for kinematics-aware multi-task robotic manipulation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18081–18090.
- Mehta, S. A.; Ciftci, Y. U.; Ramachandran, B.; Bansal, S.; and Losey, D. P. 2025. Stable-BC: Controlling covariate shift with stable behavior cloning. *IEEE Robotics and Automation Letters*, 10: 1952–1959.
- Nair, S.; Rajeswaran, A.; Kumar, V.; Finn, C.; and Gupta, A. 2022. R3M: A Universal Visual Representation for Robot Manipulation. In *Conference on Robot Learning, CoRL 2022, 14-18 December 2022, Auckland, New Zealand*, volume 205, 892–909. PMLR.
- Pan, S.; Xu, Y.; Xu, R.; Zhou, Z.; Wu, S.; and Yu, Z. 2025. Self-Correcting Robot Manipulation via Gaussian-Splatted Foresight. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 26642–26650.
- Pari, J.; Shafiullah, N. M.; Arunachalam, S. P.; and Pinto, L. 2022. The Surprising Effectiveness of Representation Learning for Visual Imitation. In *Robotics: Science and Systems XVIII, New York City, USA, June 27 - July 1, 2022*.
- Righetti, L.; Buchli, J.; Mistry, M.; Kalakrishnan, M.; and Schaal, S. 2013. Optimal distribution of contact forces with inverse-dynamics control. *The International Journal of Robotics Research*, 32(3): 280–298.
- Rubanova, Y.; Chen, R. T.; and Duvenaud, D. K. 2019. Latent ordinary differential equations for irregularly-sampled time series. *Advances in neural information processing systems*.
- Shafiullah, N. M.; Cui, Z.; Altanzaya, A. A.; and Pinto, L. 2022. Behavior transformers: Cloning  $k$  modes with one stone. *Advances in neural information processing systems*, 35: 22955–22968.
- Shi, L. X.; Sharma, A.; Zhao, T. Z.; and Finn, C. 2023. Waypoint-Based Imitation Learning for Robotic Manipulation. In *Conference on Robot Learning*, 2195–2209. PMLR.
- Shridhar, M.; Manuelli, L.; and Fox, D. 2023. Perceiver-actor: A multi-task transformer for robotic manipulation. In *Conference on Robot Learning*, 785–799. PMLR.
- Stepputtis, S.; Campbell, J.; Phielipp, M.; Lee, S.; Baral, C.; and Ben Amor, H. 2020. Language-conditioned imitation learning for robot manipulation tasks. *Advances in Neural Information Processing Systems*, 33: 13139–13150.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, B.; Wu, G.; Pang, T.; Zhang, Y.; and Yin, Y. 2024a. Diffail: Diffusion adversarial imitation learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 15447–15455.
- Wang, Z.; Zhuang, H.; Li, L.; Zhang, Y.; Zhong, J.; Chen, J.; Yang, Y.; Tang, B.; and Wu, Z. 2024b. Explore 3d dance generation via reward model from automatically-ranked demonstrations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 301–309.
- Xu, H.; Ding, J.; Xu, J.; Wang, R.; Chen, J.; Mai, J.; Fu, Y.; Ghanem, B.; Xu, F.; and Elhoseiny, M. 2025. Diffusion-based imaginative coordination for bimanual manipulation. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 11469–11479.
- Xu, T.; Li, Z.; and Yu, Y. 2020. Error bounds of imitating policies and environments. *Advances in Neural Information Processing Systems*, 33: 15737–15749.
- Zeng, J.; Bu, Q.; Wang, B.; Xia, W.; Chen, L.; Dong, H.; Song, H.; Wang, D.; Hu, D.; Luo, P.; et al. 2024. Learning Manipulation by Predicting Interaction. In *Robotics: Science and Systems, The Netherlands, July 15-19, 2024*.
- Zhao, T.; Kumar, V.; Levine, S.; and Finn, C. 2023. Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. In *Robotics: Science and Systems XIX, Daegu, Republic of Korea, July 10-14, 2023*.