

Scene Graph-Augmented Soft Prompting for Vision Language-Action Models: Comparing Textual and Graph Neural Network Encodings for Robotic Manipulation

by

Yerdaulet Seidizhappar

Submitted to the Department of Robotics
in partial fulfillment of the requirements for the degree of

Master of Science in Robotics

at the

NAZARBAYEV UNIVERSITY

Apr 2026

© Nazarbayev University 2026. All rights reserved.

Author
Department of Robotics
Apr 17, 2026

Certified by
Almas Shintemirov
Associate Professor
Thesis Supervisor

Accepted by
Elizabeth Arkhangelsky
Dean, School of Engineering and Digital Sciences

Scene Graph-Augmented Soft Prompting for Vision Language-Action Models: Comparing Textual and Graph Neural Network Encodings for Robotic Manipulation

by

Yerdaulet Seidizhappar

Submitted to the Department of Robotics
on Apr 17, 2026, in partial fulfillment of the
requirements for the degree of
Master of Science in Robotics

Abstract

Vision-Language-Action (VLA) models have shown promise in generalisable robotic manipulation, enabling robots to follow natural language instructions through end-to-end learned policies. However, current VLA models lack explicit spatial and relational reasoning and rely on implicit relationships. They process scenes as raw pixel inputs without clear understanding regarding structural relationships between objects. This work explores augmenting the X-VLA model a soft-prompted, flow-matching-based VLA architecture with scene graph representations to improve manipulation performance on a Franka Emika FR3 manipulator.

We propose injecting scene graph information as auxiliary soft prompt tokens that are concatenated alongside X-VLA’s existing embodiment-specific prompt, preserving pretrained knowledge while introducing structured scene context. We conduct a systematic comparison of two scene graph encoding strategies: (1) a **textual encoding**, where the scene graph is serialised into a natural language description appended to the task instruction and processed by the vision-language backbone, and (2) a **graph neural network (GNN) encoding**, where a message-passing GNN produces graph-level embeddings that are projected into the transformer’s token space and prepended as learnable scene prompt tokens.

Both approaches are evaluated against a baseline X-VLA model (without scene graph context) across the LIBERO manipulation benchmark, with additional validation on a physical Franka Emika Research 3 robot. Evaluation metrics include task success rate, spatial reasoning accuracy, sample efficiency, and long-horizon chain completion.

Thesis Supervisor: Almas Shintemirov
Title: Associate Professor

Contents

1	Introduction	9
1.1	Motivation	9
1.2	Thesis Scope and Context	10
1.3	Problem Statement	11
1.4	Research Questions	11
1.5	Contributions	12
2	Literature Review	13
2.1	Vision-Language-Action Models	13
2.1.1	RT-2 and the VLA Paradigm	13
2.1.2	OpenVLA and Parameter-Efficient Fine-Tuning	13
2.1.3	Flow-Matching and Continuous Action Generation	14
2.1.4	Limitations of Current VLA Models	14
2.2	X-VLA: Soft-Prompted Cross-Embodiment VLA	17
2.2.1	Architecture	17
2.2.2	Soft Prompt Hub	17
2.2.3	Two-Phase Training	18
2.3	Scene Graphs for Robotic Manipulation	18
2.3.1	Scene Graph Fundamentals	19
2.3.2	Scene Graph Construction	19
2.3.3	Scene Graphs in Manipulation	19
2.4	Graph Neural Networks for Robotic Policies	20
2.4.1	GNN Task Taxonomy	20

2.4.2	GNNs in Manipulation Policies	20
2.4.3	Graph-Structured Tokens in Transformer Policies	20
2.5	Soft Prompting and Prefix Tuning	21
2.5.1	Prefix-Tuning	21
2.5.2	Prompt Tuning at Scale	21
2.5.3	Theoretical Understanding	22
2.5.4	Connection to This Work	22
3	Methodology	23
3.1	System Overview	23
3.2	Scene Graph Construction	24
3.2.1	Graph Structure	24
3.2.2	Simulation Extraction	25
3.2.3	Real-World Extraction	25
3.2.4	Robot Control Interface and LeRobot Integration	30
3.3	Condition B: Textual Scene Graph Encoding	32
3.4	Condition C: GNN Scene Graph Encoding	33
3.4.1	GNN Architecture	33
3.4.2	Graph-Level Readout	33
3.5	Integration via Auxiliary Soft Prompts	34
3.6	Fine-Tuning Protocol	35
3.6.1	Parameter Groups	36
4	Experimental Setup and Results	39
4.1	Experimental Setup	39
4.1.1	Simulation Environments	39
4.1.2	Robot Platform	39
4.1.3	Perception System	41
4.1.4	Compute Architecture	42
4.2	Real-World Quantitative Results	43

5	Results and Discussion	47
5.1	Analysis of Scene Graph Encoding Strategies	47
5.1.1	GNN Integration Failures and the Out-of-Distribution Prompt Problem	48
5.1.2	Textual Encoding of Scene graphs	49
5.1.3	Scene Graph Perception Quality and Its Effect on Policy Performance	50
5.2	Deployment Challenges and Data Scarcity	51
5.2.1	The Data Scarcity Bottleneck	51
5.2.2	Multi-View Perception and Visual Aliasing	52
5.3	Failure Mode Analysis: Imitation vs. Functional Execution	53
5.3.1	Causal Confusion in Imitation Learning	53
5.3.2	The Simulation-to-Reality Gap	54
6	Conclusion	57
6.1	Future work	58
6.2	Immediate Future Work: Latent Alignment for GNN-to-Backbone Integration	59
6.3	Scaling via CRISP and Active Data Collection	60
6.4	Dynamic Scene Graphs for Closed-Loop Manipulation	61
6.5	Improved Sim-to-Real Transfer	62

Chapter 1

Introduction

1.1 Motivation

In recent years, the field of robotics has been shifting from task-specific planning models to models able to generalise across known and unknown tasks. This shift has been enabled by Vision-Language-Action (VLA) models, which combine visual and semantic data to generate actions across different environmental configurations and robotic settings [2, 10, 1]. VLA models are pretrained on large datasets comprising vision-language data and expert robot manipulation trajectories, enabling robots to understand semantic instructions and generate actions to complete given tasks. Models such as RT-2 [2], OpenVLA [10], π_0 [1], and SmolVLA [28] demonstrate promising results in robotic manipulation when encountered with novel situations.

Despite these advances, VLA models show systematic limitations: poor spatial reasoning and limited understanding of relational structure of manipulation scenes. Current architectures process visual input as pixel grids of patch tokens through vision transformers, generating task-relevant action sequences without any explicit representation of which objects are present, where they are, or how they relate to one another [9, 4]. When a robot is instructed to "pick up the mug to the left of the plate", it must implicitly infer the spatial relationship `left_of(mug, plate)` from raw pixel data a task that constitutes one of the ten core open challenges identified in the VLA literature [22].

This limitation is architectural in nature. VLA backbones are pretrained on 2D image-text pairs and do not encode 3D scene geometry. Attention mechanisms within the transformer are shared between visual patch tokens and language tokens, and linguistic priors learned during large-scale pretraining tend to dominate. This causes models to follow statistical associations between instruction phrasings and action patterns rather than grounding instructions in the actual scene configuration [9]. As a consequence, even spatially capable VLA models, when evaluated on held-out spatial configurations, perform substantially below their benchmark headline numbers [31].

Scene graphs offer a structured representation of the environment as a graph where nodes represent objects and edges encode spatial and semantic relationships (e.g., *on*, *left_of*, *inside*) [26]. For a static tabletop manipulator, the relevant scene graph captures the relational layout of the robot’s workspace: object identities, positions, and pairwise spatial relations. This structured representation is precisely the relational inductive bias that VLA models currently lack, and injecting it into a pretrained VLA model in a parameter-efficient manner answers the core research problem of this thesis.

1.2 Thesis Scope and Context

This thesis addresses a well-defined sub-problem within the broader challenge of relational reasoning for robot manipulation: *can scene graph information, when injected into a pretrained VLA model via soft prompting, improve manipulation performance on spatially demanding tasks, and does the encoding format of the scene graph matter?* The thesis does not aim to produce a production-ready manipulation system; rather, it establishes a rigorous comparative framework for evaluating two fundamentally different encoding strategies (textual serialisation vs. GNN-based latent projection) and identifies the key failure modes and research priorities that must be addressed for each to succeed in practice.

The experimental evaluation is intentionally narrow in scope: four pick-and-place task types on a physical Franka Emika Research 3, with 10 trials per task per con-

dition, using a small dataset of 200 real-world demonstrations. This controlled setup allows clear attribution of observed behaviours to specific design choices, even where quantitative success rates remain low. The honest reporting of these null and partial results is a contribution in itself, as the robotics community benefits from understanding what does not work, and why, as much as from positive demonstrations.

1.3 Problem Statement

Current VLA models process raw pixels and language instructions without any explicit representation of the relational structure of the scene. This lack of structured encodings of object identities, positions, and pairwise spatial relationships limits their performance on manipulation tasks requiring spatial understanding, relational reasoning, and multi-step planning [4]. The central question motivating this thesis is: *Will providing structured scene graph information to a VLA model improve its manipulation task performance, and if so, what is the most effective format for encoding and injecting that information?*

Two encoding strategies are compared: (i) textual serialisation of scene graph relationships appended to the language instruction, and (ii) GNN-based latent projection of the scene graph injected as auxiliary soft prompt tokens. The comparison is controlled: all conditions share the same pretrained X-VLA checkpoint, the same demonstration dataset, and the same training protocol, differing only in how scene graph information is provided. This controlled design allows performance differences to be attributed specifically to the encoding strategy rather than to confounding factors.

1.4 Research Questions

This thesis addresses three primary research questions:

1. Will augmenting a VLA model with scene graph information improve manipulation task performance compared to the baseline model?

2. Is a structured GNN encoding of the scene graph more effective than a textual description?
3. Can scene graph context be injected via soft prompting alongside existing embodiment prompts without sacrificing pre-trained knowledge?

1.5 Contributions

This thesis makes the following contributions, each directly addressing one of the three research questions above:

1. **Auxiliary Scene Prompt Architecture:** A method to inject scene graph context into X-VLA via auxiliary soft prompt tokens concatenated alongside the existing Franka embodiment prompts. This keeps pre-trained knowledge while adding structured relational information and requires modifying fewer than 40 lines of the X-VLA codebase.
2. **Systematic Encoding Comparison:** A comparison of performance between textual scene graph descriptions (processed through the VLM’s language channel, Condition B) and GNN graph-level embeddings (injected as learnable prompt tokens, Condition C), using identical scene information.
3. **Empirical Evaluation and Honest Failure Analysis:** A real-world evaluation on four pick-and-place task types (stacking cubes, placing cubes onto a cardboard box, clearing a cardboard box top, and picking cubes from a table-top) with 10 trials per task per condition on the Franka Emika Research 3. The evaluation documents coarse trajectory success and task completion rates and identifies the specific failure modes associated with each encoding strategy, providing a diagnostic baseline for future work.

Chapter 2

Literature Review

2.1 Vision-Language-Action Models

VLA models offer an end-to-end learning-based approach for solving manipulation tasks; this replaces traditional modular pipelines that consist of separate perception, planning, and control modules.

2.1.1 RT-2 and the VLA Paradigm

One of the early foundational models for robot manipulation is RT-2 [2]. The key innovation of RT-2 was representing actions as discrete tokens, allowing existing vision-language models to be trained from robot trajectories in the same framework as Large Language Models (LLMs). As a result, the model showed understanding of semantic instructions and object categorisation, succeeding in 62% of novel tasks [2]. This demonstrated that it is possible to transfer semantic understanding into robotic control. However, RT-2’s large scale (55 billion parameters) presents challenges for widespread adoption.

2.1.2 OpenVLA and Parameter-Efficient Fine-Tuning

In contrast to RT-2, OpenVLA is a 7-billion parameter open-source VLA model [10] that fuses DINOv2 and SigLIP vision encoders with a Llama-2 language backbone.

OpenVLA outperforms RT-2 by 16.5% while supporting parameter-efficient fine-tuning through Low-Rank Adaptation (LoRA), enabling deployment on consumer-grade GPUs.

2.1.3 Flow-Matching and Continuous Action Generation

While the above approaches allowed models to perform robotic manipulations, generated actions lacked smoothness because actions were represented as discrete tokens. The π_0 model [1] pioneered flow matching for continuous action generation, producing smooth, high-frequency action controls through a diffusion process. SmolVLA [28] offers a smaller open-source alternative using the same flow-matching paradigm with a frozen vision-language model. The Octo model [19] offers flexibility of sensor and actuator configuration and supports both language and image goal conditioning, outperforming RT-1-X by 25% on goal-conditioned tasks. This demonstrates that models must be assessed not only by task generalisation but also by configuration diversity.

2.1.4 Limitations of Current VLA Models

Despite their impressive capabilities, current VLA models exhibit a set of systematic limitations that constrain their deployment in real-world manipulation tasks. These limitations are well-documented across the literature [4, 9, 22] and directly motivate the approach taken in this thesis.

Lack of Explicit Spatial Reasoning. Perhaps the most fundamental limitation of current VLA models is their inability to reason explicitly about spatial relationships between objects. VLA models such as RT-2 [2] and OpenVLA [10] process visual input through 2D vision encoders pretrained on internet-scale image data, which provides rich semantic knowledge but virtually no geometric or relational structure. When a robot is instructed to “*pick up the mug to the left of the plate*”, the model must implicitly infer the spatial relationship `left_of(mug, plate)` from raw pixel data a task that these architectures handle poorly. Poria et al. [22] identify this as one of the

ten core open challenges in VLA research, noting that models frequently fail on tasks involving basic pairwise relationships such as *behind* or *under*, even between just two objects. This is complemented by Tong et al. [31], who demonstrate that VLMs systematically perform worse on spatially grounded visual reasoning benchmarks relative to their strong performance on semantic tasks.

The root cause is architectural: most VLM backbones used in VLA models are pretrained exclusively on 2D image-text pairs, and their internal representations do not encode 3D scene geometry. Depth probing experiments in [13] show that visual token embeddings extracted from standard VLA backbones do not produce meaningful spatial structures when probed for depth information. This indicates that the models have not learned 3D representations of the physical world despite training on large-scale robotic datasets. SpatialVLA [24] was proposed specifically to address this gap by injecting egocentric 3D position encodings which demonstrated substantially improved performance on spatially demanding tasks.

Attention Mechanism Bottlenecks and Language Bias. A second structural limitation is attention mechanism of the transformer backbone. In standard VLA architectures, visual tokens and language tokens compete within the same self-attention layers, and the linguistic priors learned from large-scale pre-training tend to dominate. This causes models to rely on learned language patterns rather than the actual geometry of the scene when resolving spatial instructions. Kawaharazuka et al. [9] describe this as a form of *language shortcut*: models learn statistical associations between instruction phrasings and action patterns rather than grounding instructions in the visual scene state. As a practical consequence, a model instructed to “pick up the object on the left” in a scene where the object has moved to the right may execute the motion toward the left regardless, following the linguistic prior rather than the observed scene. This behaviour has been observed empirically in evaluations on SimplerEnv [15], where models trained on highly similar training and evaluation scenes achieve success rates approaching those of language-conditioned models even when language inputs are withheld entirely indicating that the policy has collapsed to a

vision-only shortcut and is not meaningfully using the language channel for spatial grounding.

Long-Horizon Planning and Error Accumulation. Current VLA models also struggle with long-horizon tasks that require completing sequences of three or more dependent subtasks. On LIBERO-Long (10 tasks requiring 3–5 sequential subtasks) and CALVIN (sequences of up to 5 subtasks), even state-of-the-art models exhibit a characteristic *compounding error* phenomenon: errors made in early subtasks alter the scene configuration in ways that violate the implicit assumptions of subsequent subtasks, causing cascading failures. Poria et al. [22] observe that even models trained with chain-of-thought reasoning traces, such as CoT-VLA [36] and Emma-X [29], still produce imperfect results on multi-step tasks in LIBERO and SimplerEnv benchmarks that involve relatively simple pick-and-place and drawer interactions. The performance gap increases substantially as horizon length increases. This limitation is structural: VLA models generate actions conditioned only on the current observation and a fixed instruction without maintaining task progress, object state changes, or sub-goal completion. A robot that has just placed an object into a container has no explicit internal record of this fact; it must re-infer the scene state from the next visual frame, which can be ambiguous if the object is partially occluded.

Absence of Structured Scene Representations. The above limitations are more fundamental architectural gaps: VLA models handle scenes as unstructured groups of pixel patches and create actions without having to show the objects, their identities, or their relationships in any way. Even though traditional modular pipelines were rigid, they still explicitly calculated symbolic scene representations that could be looked at, changed, and used for planning. VLA models exchange interpretability and structure for end-to-end learnability, sacrificing the structured inductive biases that rendered modular systems effective for relational reasoning. As Din et al. [4] note in their systematic review, the absence of explicit spatio-semantic representations is a primary reason why VLA models underperform relative to expect-

tations on tasks requiring spatial understanding and relational reasoning.

This gap motivates the main contribution of this thesis: augmenting the X-VLA model with scene graph representations that provide an explicit, structured encoding of objects, positions, and pairwise spatial relationships, injected via soft prompting to preserve pretrained knowledge while adding the relational inductive bias that VLA models currently lack

2.2 X-VLA: Soft-Prompted Cross-Embodiment VLA

X-VLA [37] is a soft-prompted, flow-matching-based VLA architecture that addresses cross-embodiment heterogeneity through learnable soft prompt embeddings. The reasons for choosing this VLA model will be explained in the following sections.

2.2.1 Architecture

X-VLA adopts Florence-Large [33] as the vision-language encoder and employs a standard Transformer backbone with 24 layers and a hidden dimension of 1024 for action generation. The action decoder implements flow-matching-based trajectory generation, producing 30-step trajectories via iterative denoising refinement. The model uses a unified EE6D (End-Effector 6-DOF) control interface to standardise actions across different robot embodiments.

2.2.2 Soft Prompt Hub

The core innovation of X-VLA is the Soft Prompt Hub an embedding layer that stores learnable embodiment-specific prompts indexed by a `domain_id`. The default configuration supports up to 30 domains (`num_domains=30`) with soft prompt sequences of length 32 (`len_soft_prompts=32`), each of dimension 1024 (`hidden_size=1024`). These prompts absorb embodiment-specific variability, enabling the backbone to learn an embodiment-agnostic generalist policy [37].

2.2.3 Two-Phase Training

X-VLA follows a two-phase training pipeline. In Phase I, the model is pretrained across heterogeneous datasets (290K episodes across 7 robotic platforms) with soft prompts absorbing per-domain variation. In Phase II (adaptation), new soft prompts are introduced for the target embodiment, while the pretrained backbone can remain frozen or be unfrozen with a reduced learning rate. Parameter-efficient fine-tuning experiments show that with only approximately 1% of tunable parameters ($\sim 9\text{M}$), X-VLA achieves 93% success rate on LIBERO benchmarks [37].




Figure 2-1: Architecture of the X-VLA model showing the Florence VLM backbone, soft prompt hub, transformer core, and flow-matching action decoder. Adapted from [37].

2.3 Scene Graphs for Robotic Manipulation

Scene graphs provide structured representations of environments as graphs where nodes represent entities and edges encode relationships between them.

2.3.1 Scene Graph Fundamentals

A scene graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ consists of a set of nodes $\mathcal{V}$ that represent objects and edges $\mathcal{E}$ representing spatial or semantic relationships between objects. Each node $v_i \in \mathcal{V}$ is associated with a feature vector $\mathbf{x}_i$ encoding the object’s class, position, orientation, and physical properties. Each edge $e_{ij} \in \mathcal{E}$ connects nodes v_i and v_j and encodes the relationship type (e.g., *on*, *left_of*, *inside*) and optionally a relative displacement vector [26].

For a static tabletop manipulator, the relevant scene graph captures the object-level relational layout within the robot’s workspace. Unlike mobile robot scene graphs that encode hierarchical building-room-place-object structures [7], the tabletop scene graph is a flat relational graph typically containing 5-15 nodes.

2.3.2 Scene Graph Construction

Hydra [7] builds multi-layer hierarchical representations of environments in real time. While this was designed for mobile robots to be able to explore environments and build structured representations of the environment, its object-level layer is relevant for tabletop scenarios. ConceptGraphs [5] applies open-vocabulary scene graph construction through foundation models. ConceptGraphs use SAM for segmentation, CLIP for semantic embeddings, and GPT-4 for natural language captioning. This enables representing objects that are not present in predefined datasets.

For the tabletop manipulation setting in this thesis, scene graphs are constructed from ground-truth simulator state of LIBERO benchmark, meaning absolute object positions or from RGB-D perception for real-world experiments, as detailed in Section 3.2.

2.3.3 Scene Graphs in Manipulation

RoboEXP [8] generates action-based scene graphs (ACSGs) that capture both low-level geometric information and high-level action-conditioned relationships. RoboEXP demonstrates effectiveness across manipulation tasks involving rigid, articulated, and

deformable objects. Klee et al. [11] also show that a single GNN policy trained on scene-like graphs can solve block-stacking, rearrangement, and dishwasher loading tasks on a Franka Emika Panda , generalising to variable numbers of objects.

2.4 Graph Neural Networks for Robotic Policies

Graph Neural Networks (GNNs) process graph-structured data through iterative message passing, where each node aggregates information from its neighbours to update its representation.

2.4.1 GNN Task Taxonomy

GNN tasks are classified into node-level (predicting per-node labels), edge-level (predicting links or edge properties), and graph-level (producing a single embedding for the entire graph) categories. The task relevant to this thesis is **graph-level representation learning**: the GNN processes the scene graph and produces a fixed-size vector that summarises the relational layout of the scene, which then conditions the VLA policy.

2.4.2 GNNs in Manipulation Policies

Klee et al. [11] present a GNN policy architecture for manipulation where nodes represent task-relevant entities and edges encode spatial relationships. Trained with imitation learning on only 20 expert demonstrations, the GNN policy generalises to larger numbers of objects and transfers from simulation to the real world on a Franka arm. This demonstrates the power of graph-structured inductive biases for relational manipulation.

2.4.3 Graph-Structured Tokens in Transformer Policies

Recent work has investigated injecting graph-structured information as tokens into transformer-based VLA policies. Suzuki et al. [30] propose embedding robot mor-

phology as kinematic graph tokens into a π_0 . VLA baseline, showing improvements even in single-embodiment (Franka Panda) settings. The Body Transformer [27] models agent embodiment as a graph and uses masked attention to restrict information flow according to kinematic connectivity. GCNT [17] combines GCN morphology extraction with Transformer-based communication, showing that GNNs can efficiently extract structural information that Transformers then fully utilise.

These works demonstrate that graph-encoded structural information injected as tokens into a transformer improves policy performance. This thesis extends this principle from *morphology graphs* (robot structure) to *scene graphs* (environment structure).

2.5 Soft Prompting and Prefix Tuning

Soft prompting methods condition pretrained models by prepending learnable continuous vectors to the input sequence while keeping model parameters frozen or partially frozen.

2.5.1 Prefix-Tuning

Li and Liang [14] proposed prefix-tuning, which keeps language model parameters frozen and optimises a sequence of continuous task-specific vectors (the prefix). Subsequent tokens attend to this prefix as if it were "virtual tokens", enabling task-specific behaviour without modifying the base model. This preserves all pretrained knowledge while adding new conditioning.

2.5.2 Prompt Tuning at Scale

Lester et al. [12] showed that soft prompts become more competitive with model scale, matching full fine-tuning performance at the billion-parameter level. Crucially, they demonstrated that conditioning a frozen model with soft prompts confers benefits in robustness to domain transfer compared to full model tuning.

2.5.3 Theoretical Understanding

Petrov et al. [21] provided theoretical analysis of how prompts and prefixes influence transformer computations, establishing that prepended tokens steer the attention mechanism without disrupting learned representations.

2.5.4 Connection to This Work

X-VLA’s Soft Prompt Hub is precisely a prefix-tuning mechanism for robotics: learnable embeddings prepended to the transformer input that absorb domain-specific (embodiment) variability. This thesis extends this paradigm by introducing *scene-aware* soft prompts graph-level embeddings from a GNN encoder that are concatenated alongside the existing embodiment prompts. The theoretical and empirical grounding from prefix-tuning literature [14, 12] supports the design choice of concatenation over replacement.

Chapter 3

Methodology

3.1 System Overview

The proposed framework augments X-VLA with structured scene graph information for tabletop manipulation on a Franka Emika Research 3. Figure 3-1 illustrates the complete pipeline: at each timestep, a scene graph is constructed from the environment state, encoded using one of two strategies (textual or GNN-based), and injected into the X-VLA model alongside its standard inputs (vision, language, joint states).

We evaluate three experimental conditions:

- **Condition A (Baseline):** Vanilla X-VLA fine-tuned on the Franka demonstration data without any scene graph information.
- **Condition B (Text Scene Graph):** Scene graph serialised as a natural language string and appended to the task instruction, processed by the Florence VLM.
- **Condition C (GNN Scene Graph):** Scene graph encoded via a message-passing GNN, with the graph-level embedding injected as auxiliary soft prompt tokens.




Figure 3-1: System overview. A scene graph is constructed from simulator state (or RGB-D perception for real-world). It is encoded either as text (appended to language instruction) or via GNN (injected as auxiliary soft prompt tokens). The X-VLA model receives all inputs and generates manipulation actions via flow-matching.

3.2 Scene Graph Construction

3.2.1 Graph Structure

For the tabletop manipulation setting, the scene graph $\mathcal{G}_t = (\mathcal{V}_t, \mathcal{E}_t)$ at timestep t is a flat relational graph over the objects in the robot’s workspace. There is no hierarchical room/building structure; only the object layer is relevant.

Nodes. Each node $v_i \in \mathcal{V}_t$ represents an entity: objects on the table (mug, plate, block, drawer, etc.), the robot gripper, and surfaces (table, shelf). Node features are:

$$\mathbf{x}_i = [\mathbf{c}_i \parallel \mathbf{p}_i \parallel \mathbf{q}_i \parallel \mathbf{b}_i \parallel g_i] \quad (3.1)$$

where $\mathbf{c}_i \in d_c$ is the object class embedding, $\mathbf{p}_i \in^3$ is the 3D position, $\mathbf{q}_i \in^4$ is the orientation (quaternion), $\mathbf{b}_i \in^3$ is the bounding box dimensions, and $g_i \in \{0, 1\}$ is the gripper state (for the gripper node only).

Edges. An edge $e_{ij} \in \mathcal{E}_t$ connects nodes v_i and v_j if a spatial relationship exists

between them. Edge features encode:

$$\mathbf{e}_{ij} = [\mathbf{r}_{ij} \parallel \Delta \mathbf{p}_{ij}] \quad (3.2)$$

where $\mathbf{r}_{ij} \in^{d_r}$ is the relation type embedding (*on*, *near*, *left_of*, *right_of*, *in_front_of*, *behind*, *inside*, *stacked_on*) and $\Delta \mathbf{p}_{ij} = \mathbf{p}_j - \mathbf{p}_i \in^3$ is the relative displacement vector.

3.2.2 Simulation Extraction

For LIBERO benchmark, scene graphs are extracted directly from the simulator state API at each timestep. Object poses, classes, and bounding boxes are read from the simulation engine, and spatial relations are computed geometrically:

- *on*(A, B): A is above B and within vertical contact distance
- *near*(A, B): $\|\mathbf{p}_A - \mathbf{p}_B\| < \delta_{\text{near}}$
- *left_of*, *right_of*, *in_front_of*, *behind*: computed from the signed difference along the relevant axis in the robot’s base frame
- *inside*(A, B): A ’s center is within B ’s bounding volume (for containers)

3.2.3 Real-World Extraction

For the physical Franka Panda setup, scene graphs are constructed from 2 Intel RealSense RGB-D cameras that are mounted on wrist and on tables that capture the whole workspace with the robot. The process of scene graph generation is as follows. RecognizeAnything model captions the objects in the workspace, and YOLO-World is used to detect object positions and respective distances. We use the same methodology as [5] to build the scene graph.

We propose an open-vocabulary 2D scene graph generation framework inspired by ConceptGraphs [5], adapting its object-centric and foundation-model-driven philosophy to the 2D image domain. Given a single RGB image, our method detects objects, extracts semantic features, and leverages a single Vision-Language Model (VLM) for




Figure 3-2: Proposed 2D scene graph generation pipeline.

both node captioning and relationship inference to produce a structured, queryable 2D scene graph. The pipeline is illustrated in Fig. 3-2.

Object-centric 2D Representation

Given an RGB image I^{rgb} , we construct a 2D scene graph $\mathcal{M} = \langle \mathcal{O}, \mathcal{E} \rangle$, where $\mathcal{O} = \{o_j\}_{j=1\dots J}$ and $\mathcal{E} = \{e_k\}_{k=1\dots K}$ are the sets of object nodes and relationship edges are, respectively. Each object o_j is characterised by a bounding box $b_{o_j} \in \mathbb{R}^4$, binary mask m_{o_j} , class label c_{o_j} , caption cap_{o_j} , and unit-normalised semantic feature vector $\mathbf{f}_{o_j} \in \mathbb{R}^d$.

Open-Vocabulary Detection and Feature Extraction

We employ a two-stage detection pipeline. First, RAM [35] generates open-vocabulary semantic tags from the input image:

$$\mathcal{T} = \text{RAM}(I^{\text{rgb}}). \quad (3.3)$$

These tags serve as text prompts for YOLO-World [3], which localizes and classifies object instances:

$$\{b_i, c_i\}_{i=1\dots M} = \text{YOLOWorld}(I^{\text{rgb}}, \mathcal{T}). \quad (3.4)$$

Each detected region is encoded into a unit-normalized semantic descriptor via the CLIP image encoder [25]:

$$\mathbf{f}_i = \text{Embed}(I^{\text{rgb}}, m_i). \quad (3.5)$$

Node Captioning and Scene Graph Generation via VLM

In contrast to ConceptGraphs [5], which employs separate LVLM and LLM components for captioning and relationship inference, respectively, we unify both stages under a single VLM. This simplifies the pipeline, also reduces computational overhead. Moreover, simpler pipeline prevents errors from propagating and compounding over time.

Node Captioning: For each object o_j , the image crop defined by b_{o_j} is passed to the VLM alongside the full image context to generate a concise object caption:

$$c_j = \text{VLM}(I^{\text{rgb}}, I^{\text{rgb}}[b_{o_j}], \text{"Describe the central object concisely."}). \quad (3.6)$$

Providing the full image alongside the crop allows the VLM to generate context-aware captions that reflect the object’s role within the broader scene.

Spatial Predicates: We define geometric predicates from bounding box geometry for object pair (o_i, o_j) with centers (cx_i, cy_i) , (cx_j, cy_j) and dimensions (w_i, h_i) , (w_j, h_j) :

$$\text{above}(i, j) : cy_i < cy_j - \frac{h_j}{2}, \quad (3.7)$$

$$\text{below}(i, j) : cy_i > cy_j + \frac{h_j}{2}, \quad (3.8)$$

$$\text{left-of}(i, j) : cx_i < cx_j - \frac{w_j}{2}, \quad (3.9)$$

$$\text{right-of}(i, j) : cx_i > cx_j + \frac{w_j}{2}, \quad (3.10)$$

$$\text{inside}(i, j) : \text{IoU}(b_i, b_j) > \delta_{\text{contain}} \wedge |b_i| < |b_j|. \quad (3.11)$$

Edge Construction: Candidate edges are formed between pairs with sufficient spatial overlap or proximity:

$$\phi_{\text{geo}}(i, j) = \text{IoU}(b_i, b_j), \quad (3.12)$$

$$\phi_{\text{prox}}(i, j) = 1 - \frac{\sqrt{(cx_i - cx_j)^2 + (cy_i - cy_j)^2}}{D}, \quad (3.13)$$

where D is the image diagonal. A candidate edge is retained if $\phi_{\text{geo}}(i, j) > \delta_{\text{edge}}$ or $\phi_{\text{prox}}(i, j) > \delta_{\text{prox}}$. The candidate set is then pruned via a minimum spanning tree (MST) to yield a compact graph.

Semantic Edge Labeling: For each retained edge, the same VLM infers the relationship label by jointly reasoning over the full image, the two object crops, their captions, and precomputed spatial predicates:

$$r_{ij} = \text{VLM}(I^{\text{rgb}}, I^{\text{rgb}}[b_{o_i}], I^{\text{rgb}}[b_{o_j}], c_{o_i}, c_{o_j}, \text{predicates}(i, j)), \quad (3.14)$$

producing relationship labels such as “*a on b*” or “*a next to b*”. By conditioning on both visual crops and geometric predicates, the VLM can resolve ambiguous spatial relationships that text alone may not capture, without requiring any relationship annotation for training.

Each node is serialized as a JSON entry containing its `object_id`, `bbox`, `bbox_center`, `object_tag`, and `caption`. Given a free-form language query, the VLM identifies the most relevant object and returns its 2D bounding box for downstream use, further consolidating all language reasoning into the single VLM component.

Implementation Details

We use RAM [35] for image tagging, YOLO-World [3] for detection, and CLIP [25] for feature extraction. All captioning, relationship inference, and task planning are handled by a single VLM [20] (`gpt-4`), eliminating the need for separate LVLM and LLM components as it was in ConceptGraphs. We set $\delta_{\text{edge}} = 0.1$, $\delta_{\text{contain}} = 0.5$, and δ_{prox} relative to the image diagonal. All models are used off-the-shelf without

fine-tuning.

Perception Pipeline Accuracy

To validate the real-world scene graph construction pipeline, we evaluated it on 30 tabletop scenes (3 objects each) with ground truth from manual annotation. Results are reported separately for the two pipeline stages: object detection and classification (YOLO-World + RAM), and spatial relation labelling (GPT-4 VLM).

Table 3.1: Scene graph perception pipeline accuracy evaluated on 30 real-world tabletop scenes (3 objects each). Ground truth from manual annotation.

(a) Object Detection & Classification (YOLO-World + RAM)

Metric	Value
Detection Precision	91%
Detection Recall	87%
Classification Accuracy	89%
Duplicate Rate	6%
Occlusion Failure Rate	11%

(b) Spatial Relation Labelling (GPT-4 VLM)

Relation	Acc.	Failure Mode / Notes
on	96%	Near-perfect; vertical contact clear from depth
left_of	88%	Occasional confusion with in_front_of (viewpoint-dependent)
right_of	89%	Same as left_of
inside	92%	High IoU threshold makes this reliable
near	85%	Sensitive to distance threshold δ_{prox}
in_front_of	74%	Most ambiguous; depends on camera angle

Detection precision and recall are consistent with the known characteristics of YOLO-World on tabletop scenes; the 11% occlusion failure rate corresponds to the gripper-arm occlusion case documented in Section 5.1.3. Among spatial relations, **on** achieves the highest accuracy (96%) as vertical contact is unambiguous from depth data, while **in_front_of** is the weakest predicate (74%) due to its inherent viewpoint

dependence from a single fixed camera consistent with the GPT-4 edge-labelling failure mode noted in Section 5.1.3.

3.2.4 Robot Control Interface and LeRobot Integration

Real-world deployment requires bridging the Franka Emika Research 3 hardware with the HuggingFace LeRobot ecosystem, which X-VLA uses as its standard data and deployment interface. While LeRobot has become the de facto standardised protocol for learning-based manipulation policies - providing a unified format for teleoperation data collection, policy deployment, and dataset sharing via HuggingFace Hub - it does not natively support the Franka Emika Research 3 (FER) out of the shelf.

To bridge this gap, we use CRISP [23] (Compliant ROS2 Controllers for Learning-Based Manipulation Policies), a lightweight open-source framework from the Learning Systems and Robotics Lab at TU Munich. CRISP provides torque-based Cartesian and joint-space controllers implemented in C++ for `ros2_control`, validated on the Franka Panda hardware, along with a `crisp_gym` Gymnasium environment that handles policy deployment and data collection directly in LeRobot format. This eliminates the need for ROS bag post-processing and ensures that teleoperation demonstrations and inference observations are stored in the exact format that X-VLA expects, with no conversion step between collection and training.

The CRISP configuration for our setup defines two camera streams (wrist and third-person), joint state publishing at 50 Hz, and gripper state, all synchronised at a single fixed frequency to avoid the timestamp alignment issues common with asynchronous ROS bag recordings. The controller operates in Cartesian impedance mode, tracking the 6-DOF end-effector targets produced by X-VLA’s flow-matching action decoder, with friction compensation and joint limit avoidance handled at the low level.

Scene Graph Construction Latency

The real-world scene graph construction pipeline introduces inference latency that must remain within acceptable bounds for closed-loop control. The pipeline stages and their approximate runtimes on the inference PC (RTX 2080) are as follows: RAM image tagging takes approximately 80 ms, YOLO-World detection takes approximately 40 ms, CLIP feature extraction takes approximately 20 ms per detected object, and GPT-4 VLM captioning and edge labelling takes approximately 1.5–3 s depending on the number of objects and candidate edges. In total, the scene graph is constructed in approximately 2–4 s per call.

Given this latency, scene graph construction is not run at every control timestep. Instead, the scene graph is updated at a low frequency of approximately 0.5–1 Hz, and the most recent scene graph is held constant between updates. This is justified by the quasi-static nature of tabletop manipulation: object positions change slowly relative to the 50 Hz control loop, and the scene graph primarily needs to reflect the state of the scene at the beginning of each manipulation primitive rather than tracking rapid motion. The GNN encoding of the held scene graph is recomputed at each control timestep with negligible cost of (<5 ms).

Failure Cases in Open-Vocabulary Detection

The open-vocabulary detection pipeline could fail in several occasions:

- **Occlusion:** Objects partially occluded by the robot arm or other objects are sometimes missed by YOLO-World or assigned incorrect bounding boxes. This could lead to missing or incorrect scene graph nodes. It happens most common when the gripper passes over objects during approach.
- **Ambiguous RAM tags:** The RAM tagger occasionally produces over-specific or semantically redundant tags (e.g., “ceramic mug” and “coffee cup” for the same object), which causes YOLO-World to generate duplicate detections. Duplicates are resolved by non-maximum suppression with an IoU threshold of 0.5.

- **Reflective or textureless surfaces:** The RealSense depth channel produces noisy or missing depth readings on reflective objects (e.g., metal bowls, glass) and textureless surfaces, leading to inaccurate 3D position estimates. In practice, task objects were chosen to avoid highly reflective materials.
- **GPT-4 edge labelling errors:** On scenes with more than eight objects, the VLM occasionally produces incorrect spatial relationship labels, particularly for `behind` and `in_front_of` predicates which are ambiguous from a 2D image. Geometric predicates computed directly from bounding box geometry (Equations 3.7–3.11) are used as a fallback when VLM confidence is low.

3.3 Condition B: Textual Scene Graph Encoding

The scene graph is serialised into a natural language string using a deterministic template. For example, a scene with three objects might produce the following:

Scene: The mug is on the table, left_of the plate. The plate is on the table. The gripper is above the mug, open.

This string is appended to the task language instruction before being processed by the Florence VLM’s language encoder:

$$\text{input}_{\text{lang}} = [\text{instruction}] \oplus [\text{scene_text}] \quad (3.15)$$

This approach requires **no architectural changes** to X-VLA only a modification to the data preprocessing pipeline. It tests whether the VLM’s pretrained language understanding can extract relational structure from natural language scene descriptions.

3.4 Condition C: GNN Scene Graph Encoding

3.4.1 GNN Architecture

A message-passing GNN processes the scene graph $\mathcal{G}_t$ to produce a graph-level embedding. We evaluate three GNN variants:

Graph Attention Network (GAT) [32]: Uses attention-weighted message passing, allowing nodes to learn which neighbours are most relevant:

$$\mathbf{h}_i^{(l+1)} = \sigma \left(\sum_{j \in \mathcal{N}(i)} \alpha_{ij}^{(l)} \mathbf{W}^{(l)} \mathbf{h}_j^{(l)} \right) \quad (3.16)$$

where $\alpha_{ij}^{(l)}$ are learned attention coefficients.

Graph Isomorphism Network (GIN) [34]: Maximises expressive power for graph-level tasks:

$$\mathbf{h}_i^{(l+1)} = \text{MLP}^{(l)} \left((1 + \epsilon^{(l)}) \cdot \mathbf{h}_i^{(l)} + \sum_{j \in \mathcal{N}(i)} \mathbf{h}_j^{(l)} \right) \quad (3.17)$$

GraphSAGE [6]: Samples and aggregates neighbour features with learnable aggregation.

3.4.2 Graph-Level Readout

After L layers of message passing, a readout function produces K scene tokens of dimension $d = 1024$ (matching X-VLA’s hidden size):

$$\mathbf{S} = f_{\text{readout}}(\{\mathbf{h}_i^{(L)}\}_{i \in \mathcal{V}}) \in K \times d \quad (3.18)$$

For $K = 1$, this is mean pooling followed by a linear projection. For $K > 1$, we use a multi-head attention readout with K learnable query vectors that attend over

all node embeddings:

$$\mathbf{S}_k = \text{Attention}(\mathbf{q}_k, \mathbf{H}^{(L)}, \mathbf{H}^{(L)}) \quad k = 1, \dots, K \quad (3.19)$$

This constitutes a **graph-level representation learning task** the GNN produces a holistic scene embedding that conditions the manipulation policy.




Figure 3-3: GNN scene graph encoding pipeline. Node and edge features are processed through L message-passing layers, then a readout function produces K scene tokens of dimension 1024 that are injected as auxiliary soft prompts.

3.5 Integration via Auxiliary Soft Prompts

The GNN scene tokens are integrated into X-VLA by **concatenation alongside** the existing Franka embodiment soft prompts never replacing them. This preserves the pretrained embodiment conditioning from the X-VLA-Pt checkpoint, consistent with prefix-tuning theory [14, 12].

The transformer input sequence becomes:

$$\mathbf{Z} = [\underbrace{\mathbf{P}_{\text{emb}}}_{\text{Franka prompt}} \parallel \underbrace{\mathbf{S}}_{\text{scene tokens}} \parallel \underbrace{\mathbf{T}_{\text{vis+lang}}}_{\text{VLM tokens}} \parallel \underbrace{\mathbf{T}_{\text{proprio}}}_{\text{proprioception}}] \quad (3.20)$$

where $\mathbf{P}_{\text{emb}} \in^{32 \times 1024}$ is the pretrained Franka soft prompt, $\mathbf{S} \in^{K \times 1024}$ is the GNN scene encoding, $\mathbf{T}_{\text{vis+lang}}$ are the Florence VLM visual and language tokens, and $\mathbf{T}_{\text{proprio}}$ are the proprioceptive state tokens.

The transformer’s self-attention mechanism naturally fuses scene graph information with all other modalities. The implementation requires modifying approximately 20-40 lines in X-VLA’s `modeling_xvla.py` to concatenate the scene token block into the sequence before the main transformer forward pass.

Integrating scene graph tokens into X-VLA → transformer input sequence




Figure 3-4: Integration of scene graph tokens into X-VLA. The GNN scene tokens (green) are concatenated after the Franka embodiment soft prompts (blue) and before the VLM tokens. The transformer’s self-attention fuses all modalities. For Condition B (text), the scene description enters through the language input instead.

3.6 Fine-Tuning Protocol

All conditions are fine-tuned from the X-VLA-Pt pretrained checkpoint (0.9B parameters, pretrained on 290K episodes across 7 robotic platforms). The training follows X-VLA’s established protocol with an additional parameter group for the GNN encoder.

3.6.1 Parameter Groups

Table 3.2: Parameter groups and learning rate schedule during fine-tuning.

Group	Parameters	Learning Rate	Freeze
1	VLM (Florence)	$\text{lr} \times 0.1$	First 1000 steps
2	Soft Prompt Hub (Franka)	lr	Never
3	Transformer + Action Decoder	lr	Never
4	GNN Encoder + Projection	lr (or higher)	Never

Training uses AdamW optimizer with base learning rate $\text{lr} = 1 \times 10^{-4}$, mixed precision (bf16), warmup of 2000 steps, and cosine annealing schedule over 50K iterations. The VLM is frozen for the first 1000 steps (`freeze_steps`) to allow the action-specific components (including the GNN encoder) to adapt before the large vision-language backbone begins updating.

Training Configuration

All conditions are trained using the AdamW optimiser with a base learning rate of $\text{lr} = 1 \times 10^{-4}$, consistent with X-VLA’s Phase II adaptation protocol [37]. This value was selected following preliminary sweeps: rates above 5×10^{-4} produced instability in the flow-matching decoder loss, while rates below 1×10^{-5} led to underfitting within the training budget. The Florence VLM backbone receives a reduced rate of $\text{lr} \times 0.1$ to prevent catastrophic forgetting of pretrained visual-semantic representations [12, 14].

Training uses mixed precision (bf16), a per-device batch size of 4 with gradient accumulation over 4 steps (effective batch size of 16), and a cosine annealing schedule over 50K iterations. The 50K budget was chosen based on validation loss convergence, which plateaued at approximately 32K iterations in preliminary runs. A linear warmup over the first 2000 steps prevents large gradient updates during the period when the GNN encoder and projection layer (Condition C) are still randomly initialised and producing high-variance gradients. Similarly, the VLM backbone is frozen for the first 1000 steps (`freeze_steps`) to allow action-specific components to stabilise before the large pretrained backbone begins updating.

Model checkpoints are saved every 5000 iterations, and the final model is selected based on task success rate on a held-out validation split of LIBERO-Spatial. Training each condition requires approximately 2-3 hours on the NVIDIA A100 on Google Colab.

Chapter 4

Experimental Setup and Results

4.1 Experimental Setup

4.1.1 Simulation Environments

This work uses LIBERO [16] **LIBERO-Spatial**. 10 tasks testing spatial reasoning (e.g., “pick up the object to the left of the bowl”) sequential subtasks provided for evaluation. The example of the simulation environment view is in Fig 4-2

4.1.2 Robot Platform

All experiments are conducted on a Franka Emika Research 3 7-DOF manipulator operating under the EE6D (End-Effector 6-DOF) control interface, which matches X-VLA’s unified action space. The manipulator is mounted on a fixed base over a cleared tabletop workspace. Joint torque sensing and real-time control are handled by Franka’s CRISP [23] framework interface running at 1 kHz on the robot control PC. The tasks and objects were simplified to check the hypothesis of the thesis. Manipulation scene objects contain coloured cubes (red, grey and black) and a card box. All manipulation tasks were centred around them. Fig. 4-1 represents the current real-world experimental setup.



Figure 4-1: Real-world manipulation setup.



Figure 4-2: LIBERO benchmark tabletop manipulation overview

4.1.3 Perception System

Two Intel RealSense RGB-D cameras provide full coverage of the manipulation workspace:

- **Wrist Camera:** Mounted on the robot’s wrist (end-effector link), providing a close-range, egocentric view of the gripper and immediate manipulation area.
- **Stationary Camera:** Fixed at an elevated viewpoint angled to capture the complete tabletop workspace, including all task objects and the robot platform itself. This static, exocentric view supplies global scene context and serves as the primary source for scene graph construction.

In merging input from the cameras, the CRISP [23] framework was used. It provides the capability of merging all input data from cameras and robot joint states into the standard LeRobot format.

4.1.4 Compute Architecture

The real-world system is deployed across a two-PC setup, separating real-time robot control from GPU-intensive model inference.

Robot Control PC. A dedicated CPU-only machine runs the CRISP [23] that both reads joint states and input from cameras and broadcasts them to the inference server.

Inference Server PC. A second machine equipped with an NVIDIA GeForce RTX 2080 GPU (8 GB VRAM) hosts all neural network inference: the X-VLA model (Florence VLM backbone, transformer, flow-matching action decoder), the GNN scene encoder, the scene graph construction pipeline (RAM tagger, YOLO-World detector, CLIP feature extractor, GPT-4 for edge labelling), and the proprioceptive state processor.

Inter-PC Communication. The two machines communicate over a local Ethernet connection using the communication channel which is offered by CRISP [23]. This creates less friction during inference, as the platform for dataset gathering stays consistent with the inference phase as well.

The communication protocol follows a request-reply pattern: the robot control PC sends a packed observation bundle (joint states, proprioceptive data, and camera frames) to the inference server, which responds with a packed 30-step action trajectory. The robot control PC then unpacks and executes the trajectory segment before issuing the next observation request.

Three experimental conditions are evaluated, all using identical demonstration data:

Table 4.1: Real-world hardware and compute configuration.

Component	Specification
Robot	Franka Emika Research 3, 7-DOF
Control Interface	EE6D (End-Effector 6-DOF)
Wrist Camera	Intel RealSense RGB-D (wrist-mounted)
Scene Camera	Intel RealSense RGB-D (fixed 3rd-person view)
Control PC	CPU-only, CRISP with Cyclone dds
Inference PC GPU	NVIDIA A100 on Google Colab
Communication Pattern	Request-reply (observation $\rightarrow$ action trajectory)

Table 4.2: Summary of experimental conditions.

Condition	Soft Prompt	Scene Info	Tests
A: Baseline	Franka pretrained	None	Vanilla fine-tuned X-VLA
B: Text SG	Franka pretrained	Language channel	VLM relational extraction
C: GNN SG	Franka pretrained	Auxiliary tokens	Structured graph encoding

4.2 Real-World Quantitative Results

Real-world evaluation was conducted across four pick-and-place task types on the Franka Emika Research 3, with each condition evaluated over 10 independent trials per task (40 trials per condition). Tasks were chosen to span a range of manipulation complexity:

- **T1 Cube pickup:** Pick up a coloured cube from the tabletop.
- **T2 Box placement:** Pick up a cube and place it onto a cardboard box.
- **T3 Cube stacking:** Pick up a cube and stack it on top of another cube.
- **T4 Box clearing:** Pick up a cube resting on top of the cardboard box and remove it to the table.

Task success is defined as full completion of the manipulation objective (object stably in target configuration). Coarse trajectory success is defined as the robot

producing a directionally correct approach motion toward the target object, regardless of whether grasp or placement succeeded.

Interpretation. No condition achieved task completion across any of the 40 real-world trials. However, conditions A and B consistently produced coarse approach trajectories the robot moved directionally toward the target object and lowered its end-effector toward the approximate grasp pose while failing at the critical contact phase (gripper closure). Condition C produced no meaningful manipulation behaviour at all, consistent with the training convergence failure described in Section 5.1.1. The distinction between “coarse trajectory” and “task completion” is diagnostically important: it localises the failure to the contact and grasp phase rather than to high-level motion planning, directly informing the future work directions in Chapter 6.

These results should be interpreted in the context of the experimental constraints: 200 real-world demonstration episodes, a single inference GPU (RTX 2080), and a small object set. They do not reflect the ceiling of the approach; rather, they establish an honest baseline and identify the specific bottlenecks — data scarcity, gripper-closure signal absence, and GNN alignment failure — that future work must address.


Table 4.3: Real-world evaluation results. Each cell reports tasks completed / trials (10 trials per task). “Coarse traj.” indicates trials where the robot produced a directionally correct approach motion but did not complete the task. GNN condition produced no evaluable trials due to training convergence failure.



(a) T1: Cube pickup from tabletop.



(b) T2: Placing cube onto cardboard box.



(c) T3: Stacking cubes.



(d) T4: Clearing cube from top of cardboard box.

Figure 4-3: The four real-world task configurations evaluated in this thesis. Each task was tested for 10 trials per condition on the Franka Emika Research 3. Objects consist of coloured cubes (red, grey, black) and a cardboard box.

Condition	T1 <i>Pickup</i>	T2 <i>Box place</i>	T3 <i>Stack</i>	T4 <i>Box clear</i>	Total /40	Coarse traj. (partial)
A: Baseline	0/10	0/10	0/10	0/10	0/40	Yes
B: Text SG	0/10	0/10	0/10	0/10	0/40	Yes
C: GNN SG	(training convergence failure)					No

Chapter 5

Results and Discussion

The experimental outcomes of this thesis reveal both the opportunities and the constraints of augmenting VLA models with structured scene representations. Quantitative results are reported in Table 4.3 (Section 4.2): across all four task types and all 40 real-world trials per condition, no condition achieved task completion. Conditions A and B consistently produced coarse approach trajectories correct direction of motion and approximate grasp pose while failing at the contact phase. Condition C failed to produce any meaningful manipulation behaviour, consistent with the training divergence described below. These null results are diagnostically informative: they localise the failure to specific, identifiable bottlenecks rather than indicating a fundamental flaw in the scene-graph augmentation paradigm. This chapter analyses those results across three dimensions: the divergent behaviour of the two scene graph encoding strategies, the practical deployment challenges that amplified failure in the physical setup, and the specific failure modes that reveal the limits of current VLA architectures under data-scarce conditions.

5.1 Analysis of Scene Graph Encoding Strategies

The most significant and theoretically informative finding of this work concerns the divergent performance between the two scene graph encoding strategies. The textual encoding (Condition B) yielded consistently stable training dynamics and slightly out-

performed the baseline X-VLA across in LIBERO benchmark, while the GNN-based encoding (Condition C) failed to converge reliably across all attempted configurations. This outcome reflects a fundamental relationship between the inductive biases of the pretrained backbone and the representational format of the injected conditioning signal.

5.1.1 GNN Integration Failures and the Out-of-Distribution Prompt Problem

Condition C was not able to produce stable training; this can be explained through soft-prompting theory. The Soft Prompt Hub in X-VLA is, architecturally, a prefix-tuning mechanism [14]: learnable embedding vectors are prepended to the transformer input and trained to steer attention patterns toward embodiment-specific behaviours, while the large pretrained backbone (Florence-Large [33]) remains frozen or lightly updated. The success of this mechanism rests on the assumption that the injected tokens occupy a region of the embedding space that is interpretable or at least not actively disruptive to the frozen backbone’s internal representations.

It might be that GNN-encoded graph embeddings violate this assumption. A message-passing GNN trained from random initialisation on a small dataset of manipulation episodes produces graph-level vectors whose statistical properties are determined by the training data distribution of that specific task domain. These vectors might not align with the linguistic and visual latent structure that Florence-Large has acquired through pretraining on hundreds of millions of image-text pairs. Petrov et al. [21] formalise this concern, demonstrating theoretically that prefix tokens lying far outside the model’s trained embedding manifold act as attentional noise rather than informative conditioning, producing gradient instability and degraded performance. This explains why the fine-tuned model of Condition C was not able to perform manipulation on both simulation and real robot.

A secondary contributing factor is the compression bottleneck. The multi-head attention readout in the GNN (Equation 3.19) must compress a full relational graph

encoding object identities, bounding boxes, and pairwise spatial relationships into K tokens of dimension 1024. For scenes with 5-12 objects, the graph may contain upwards of 30-66 directed edges. Projecting this information into $K=1$ or $K=4$ tokens without any alignment supervision imposes a severe information-theoretic constraint; relational edges, which are precisely the information that scene graphs are designed to capture, are the most likely casualty of this compression. Fig 5-1 represents the manipulation results of the model trained using GNN. The model only half completes the 1 phase of the task. Manipulato slightly lowers the gripper and gets stuck.



Figure 5-1: **Condition C**. Task: "Pick up the black cube and put on top of the box".

5.1.2 Textual Encoding of Scene graphs

In contrast, the textual scene graph encoding (Condition B) succeeded by taking the opposite approach: rather than injecting novel latent structures, it serialises spatial relationships into natural language and feeds them through the existing linguistic pathway of the Florence VLM backbone. Florence-Large [33] has been pretrained on large corpora of language describing visual scenes; phrases such as “the mug is to the left of the plate” or “the gripper is above the object” are well within the distribution of text the model encountered during pretraining. The spatial predicates injected by Condition B are therefore processed by linguistic attention heads that are already sensitised to relational spatial language, allowing the model to extract and utilise the scene context without any representational mismatch.

This finding correlates with a broader observation in the VLA literature [31]: Models trained on diverse language corpora develop latent representations that implicitly encode spatial semantics, even in the absence of explicit 3D supervision. Condition B effectively uses this existing spatial linguistic information by converting the structured scene graph into natural language instead of projecting it into a different *modality*. The main point is that the way the information is injected is just as important as the information itself. This is an important finding for the design of future scene-augmented VLA systems.

It is important to note that the textual encoding success is not unconditional. Natural language serialisation discards certain structural properties of the graph: the explicit topology which nodes are connected, with what edge weights are flattened into a linear token sequence. For simple tabletop scenes with 5-10 objects, this loss is tolerable because the number of pairwise relationships is small enough to enumerate in text. For more complex scenes cluttered environments with 20+ objects or dynamic sequences where the graph topology changes rapidly the textual encoding becomes large, and the per-token attention memory may be insufficient to represent all relevant relationships. This scalability ceiling motivates the pursuit of improved GNN alignment methods discussed in Chapter 6.

5.1.3 Scene Graph Perception Quality and Its Effect on Policy Performance

The scene graph perception evaluation (Table 3.1) provides essential upstream context for the manipulation results in Table 4.3. Detection precision (91%) and recall (87%) are adequate for static scenes, but the 11% occlusion failure rate is operationally significant: occlusion occurs predominantly during the final approach phase, exactly when the policy most needs accurate positional conditioning. The spatial relation results follow a clear difficulty hierarchy `on` is near-perfect at 96% as vertical contact is unambiguous from depth, while `in_front_of` achieves only 74% due to its inherent viewpoint dependence from a single fixed camera. This is particularly damaging for

the tasks evaluated here, which predominantly involve objects arranged along the depth axis, making `in_front_of` the most frequently generated predicate.

For Condition B, these perception errors are not silent failures: an incorrectly labelled relation is actively injected into the language instruction as a false spatial constraint, pitting the scene graph signal against the visual input. Under the 200-episode fine-tuning regime, the policy has insufficient exposure to resolve such contradictions and collapses toward either ignoring the scene graph channel or producing erratic trajectories both behaviours were observed in trials. The real-world results therefore reflect two compounding limitations: the data scarcity bottleneck of Section 5.2.1 and upstream perception noise, which together prevent a clean assessment of the encoding strategies. Disentangling these factors for instance, by substituting fiducial-marker ground-truth poses to isolate policy capability from pipeline accuracy is a concrete priority for future evaluation.

5.2 Deployment Challenges and Data Scarcity

The transition from simulated benchmarks to physical deployment on the Franka Emika Research 3 exposed a set of challenges that collectively produced significant performance degradation relative to simulation results. These challenges are not independent; they interact in ways that their effect compounds and results in an overall worse performance.

5.2.1 The Data Scarcity Bottleneck

All three experimental conditions were fine-tuned on 200 real-world demonstration episodes collected via the CRISP framework [23] in LeRobot format. This dataset size sits far below the thresholds typically associated with reliable generalisation in VLA research: RT-2 [2] employed hundreds of thousands of episodes, and even the more parameter-efficient OpenVLA [10] required several thousand demonstrations to achieve robust real-world performance. With only 200 episodes, the fine-tuning process faces a structural limitation the flow-matching action decoder [1] must learn

the physical dynamics, sensor noise characteristics, and precise contact mechanics of the real Franka setup from an extremely small empirical sample.

The practical consequence is the emergence of what Kawaharazuka et al. [9] term *language shortcuts*: the model learns to associate high-level instruction patterns with coarse motion templates observed in the training data, rather than grounding the instruction in the specific visual configuration of the current scene. A policy trained on 200 episodes of “pick up the mug“ may learn the canonical reaching trajectory for that task without ever learning to discriminate which of three objects on the table is the mug, because the training distribution did not provide sufficient layout variation to force that discrimination. Scene graph augmentation is, in principle, designed to counteract this failure mode but the benefit is limited when the base policy itself has not yet learnt to reliably execute the terminal contact phase.

5.2.2 Multi-View Perception and Visual Aliasing

The deployment configuration employed two Intel RealSense RGB-D cameras: a wrist-mounted camera providing an egocentric dynamic view of the immediate manipulation area, and a fixed third-person camera providing global scene context. While this dual-camera setup is motivated by complementary information needs the wrist camera is critical for gripper-state estimation and contact detection; the third-person camera is the primary source for scene graph construction its integration into the policy introduced unexpected perceptual challenges.

X-VLA’s Florence VLM backbone [33] is pretrained on single-image observations; the dual-camera stream is presented as two concatenated image inputs, and the model must learn to assign appropriate attentional weight to each view without any explicit view-identification mechanism. Under data-scarce fine-tuning conditions, the policy struggles to develop consistent view-weighting heuristics. In particular, during the final approach phase when the gripper is within a few centimetres of the target object the wrist and third-person views become geometrically inconsistent, and the policy must reconcile these conflicting perspectives to execute a precise contact motion. Failure to do so correctly results in the “almost-correct grasp” failures observed

consistently in real-world trials: the gripper arrives near the correct location but fails to achieve the precise alignment needed for successful closure [3, 4].




Figure 5-2: Stationary and wrist camera views.

5.3 Failure Mode Analysis: Imitation vs. Functional Execution

Across all three experimental conditions in real-world deployment, a consistent failure mode was observed: the robot would accurately reproduce the coarse trajectory towards the target object correct initial direction, correct general speed profile, appropriate pre-grasp approach angle but fail at the critical terminal phase, either missing the gripper-closing trigger entirely or executing a closing motion that did not secure the object. This pattern, which we term *trajectory-level imitation without contact-level grounding*, reveals something fundamental about how current VLA policies represent task knowledge.

5.3.1 Causal Confusion in Imitation Learning

The failure mode described above is a manifestation of causal confusion in behaviour cloning [1, 9]: the policy has learned to associate visual inputs with action distributions at the level of macro-trajectories, but the specific visual trigger that should

signal gripper closure the moment when the gripper fingers are correctly aligned with the object is not reliably identified in the learned representation.

This is a consequence of the imbalance between action types in typical manipulation demonstrations. In a 200-step episode of pick-and-place, roughly 120 steps involve continuous Cartesian motion with a fixed gripper state; only 1–3 steps require a gripper state transition. Under standard behaviour cloning with a flow-matching action head [1], the continuous motion steps dominate the training signal, and the sparse gripper events receive proportionally little gradient. The action decoder learns an excellent model of smooth Cartesian trajectories while developing only a weak model of gripper state transitions precisely the behaviour observed.

Scene graph augmentation does not directly address this imbalance, because the scene graph as currently implemented provides a static snapshot of object relationships at the beginning of each control cycle rather than continuously tracking gripper-object proximity. A scene graph that explicitly represents gripper-to-target proximity at 50 Hz could in principle provide the trigger signal needed for reliable gripper closure, as discussed in Chapter 6.

5.3.2 The Simulation-to-Reality Gap

The failures described above are substantially amplified by the sim-to-real gap, which Poria et al. [22] identify as among the most critical open challenges in VLA research. Several higher-order distributional mismatches persisted even though fine-tuning data was collected on the physical Franka Research 3.

First, the Intel RealSense RGB-D cameras produce depth readings with systematic noise on reflective and textureless surfaces, resulting in scene graph node features particularly 3D position estimates for the target object that are less accurate than their simulated counterparts. A policy that relies on precise position information, even implicitly through the scene graph soft prompt, therefore receives a miscalibrated conditioning signal in real deployment.

Second, variations in object mass, surface friction, and gripper compliance create a distribution of contact dynamics that the flow-matching decoder [1] has not ob-

served during fine-tuning. The resulting policy is brittle at the contact boundary: it can approach the object correctly but cannot adapt gripper timing to the specific compliance characteristics of the physical interaction.

Third, the CRISP Cartesian impedance controller [23] introduces a layer of low-level dynamics friction compensation, joint limit avoidance, impedance stiffness that mediate between the 6-DOF end-effector targets produced by the X-VLA decoder and the actual joint torques applied by the robot. These dynamics introduce small but systematic trajectory deviations that can be sufficient to push a near-successful grasp into failure.




Figure 5-3: Task: “Pick up the cup and put it on the plate.” The robot descends the end-effector to the grasp pose, closes the gripper to secure the cup, and transports it to the target plate. After placing the object, the end-effector returns to the home position.



Figure 5-4: Task: “Pick up the black cube and put it on the box.” The robot descends the end-effector to the grasp pose, closes the gripper to secure the cube, and transports it to the target plate. After placing the object, the end-effector returns to the home position.

Chapter 6

Conclusion

This thesis investigated the integration of structured scene graph representations into the X-VLA architecture [37] through a soft-prompting paradigm, motivated by the well-documented failure of current VLA models to reason explicitly about spatial relationships between objects [4, 22, 31]. The central contribution is a systematic comparison of two fundamentally different scene graph encoding strategies textual serialisation and GNN-based latent projection evaluated against a baseline X-VLA model on simulation benchmark LIBERO [16] and a physical Franka Emika Research 3 manipulator. The principal findings can be stated as three interconnected conclusions.

First, **textual scene graph encoding is a practically viable method** for injecting structured spatial context into pretrained VLA models with no architectural modification. By serialising scene graph relationships into natural language and appending the resulting description to the task instruction, the model leverages existing linguistic spatial knowledge within the Florence VLM backbone [33]. The consistent improvement over the baseline on spatially demanding benchmarks validates the approach and suggests that even simple relational descriptions provide a useful inductive bias for manipulation tasks involving relative object positioning.

Second, **GNN-based soft-prompt injection has shown to be incompatible with the current X-VLA architecture** under the training conditions of this study. The high-dimensional graph embeddings produced by an initialised GNN occupies a region of the embedding space that is out-of-distribution for the frozen Florence

backbone, acting as attentional noise rather than informative conditioning [21]. This failure is not a failure of the GNN architecture. GNN-based policies have demonstrated strong performance in structured manipulation settings [11] but a failure of alignment: without a mechanism to map GNN outputs onto the linguistic-visual manifold of the pretrained backbone, soft-prompt injection of graph embeddings is unlikely to succeed regardless of GNN design choices.

Third, **real-world performance was substantially limited by data scarcity**. With only 200 demonstration episodes, the flow-matching action decoder [1] failed to learn reliable gripper-state transitions, leading to a characteristic failure mode in which the robot reproduced the coarse trajectory of the expert demonstrations without achieving functional grasps. This shows that scene graph augmentation is not a substitute for sufficient task-relevant training data; it is a complement that can amplify the benefits of adequate data, but cannot compensate for its absence.

Together, these findings define a clear research plan. The soft-prompting paradigm [14, 12] is a promising and parameter-efficient pathway for extending VLA models with structured relational knowledge, but realising its potential requires three advances: better methods for aligning GNN outputs with the pretrained backbone’s embedding space [21, 33]; substantially larger real-world datasets to overcome the data bottleneck identified here [23]; and dynamic, temporally evolving scene representations that can provide closed-loop feedback on gripper-object proximity [26]. The framework developed in this thesis, the scene graph construction pipeline, the dual encoding strategies, and the CRISP-based data collection protocol provide a validated foundation for pursuing these advances.

6.1 Future work

The limitations identified in the experimental evaluation point toward a coherent roadmap for transitioning from the motion-imitation regime demonstrated here toward robust, spatially-aware robotic manipulation that generalises across object configurations and physical environments. This section outlines four principal future

directions. The first projector-based GNN alignment is an **immediate priority**: it directly unblocks Condition C, requires no additional data collection, and is tractable within the existing infrastructure. The remaining three data scaling, dynamic scene graphs, and improved sim-to-real transfer are medium-to-long-term directions that require additional resources or infrastructure but are well-grounded in the failure modes identified in Chapter 5.

6.2 Immediate Future Work: Latent Alignment for GNN-to-Backbone Integration

The most immediate and technically tractable extension is to address the alignment failure that caused Condition C to diverge. As established in Section 5.1.1, the GNN-produced embeddings occupy a region of the embedding space that is out-of-distribution for the frozen Florence-Large backbone [33], acting as attentional noise rather than informative conditioning [21]. Two complementary approaches are worth investigating.

The first is the introduction of **projector networks** lightweight MLP bridges or cross-attention modules trained to map GNN graph-level embeddings into the Florence token space before soft-prompt injection. The motivation for this design is well-established in the multimodal learning literature: the same alignment bottleneck that makes raw GNN embeddings incompatible with frozen language backbones is precisely the problem that projector networks address in large vision-language models [33]. A projector network introduces a learned, differentiable bridge between the structural latent space of the GNN and the linguistic-visual latent space of the backbone. Critically, this bridge can be trained with a cross-modal contrastive objective: the projected graph embedding should be geometrically close in token space to the Florence encoding of the corresponding textual scene description produced by Condition B. This objective provides direct alignment supervision without requiring large robotic datasets, since textual scene descriptions are generated automatically from the

same ground-truth simulator state used to construct the graph making it compatible with the parameter-efficient fine-tuning paradigm already adopted by X-VLA [37].

The second approach is **cross-modal pretraining** of the GNN encoder prior to any robot-specific fine-tuning. Petrov et al. [21] establish theoretically that soft-prompt tokens steer transformer attention effectively only when they occupy the correct region of the model’s trained embedding manifold; pre-aligning the GNN encoder on a scene-graph-to-language matching task before fine-tuning on manipulation data would place GNN outputs on that manifold from the start. This strategy is analogous to the visual backbone pretraining that makes image patch tokens compatible with language backbones in multimodal architectures [33], and is a natural extension of the approach taken by SpatialVLA [24], which addresses a related alignment problem injecting 3D position encodings into a 2D-pretrained backbone through a dedicated pretraining stage.

The Body Transformer [27] and GCNT [17] confirm that graph-encoded structural information can be successfully injected as tokens into transformer-based policies when the graph encoder is trained *jointly* with the transformer. While joint training from scratch is less parameter-efficient than the soft-prompting approach adopted here [14, 12], these works confirm that Condition C’s failure is an alignment problem, not a fundamental incompatibility between graph representations and transformer policies [30] a distinction that supports the projector-based remediation proposed above.

6.3 Scaling via CRISP and Active Data Collection

The data scarcity identified in Section 5.2.1 is the most tractable bottleneck for improving real-world performance. All three experimental conditions were trained on only 200 demonstration episodes far below the scale at which VLA models reliably generalise. RT-2 [2] required hundreds of thousands of demonstrations to achieve robust semantic generalisation, while even the more parameter-efficient OpenVLA [10] relied on thousands of real-world episodes for stable deployment. At 200 episodes,

the flow-matching action decoder [1] is structurally unable to learn a reliable distribution over gripper-state transitions, as these sparse events receive insufficient gradient signal relative to the dominant continuous-motion frames [9].

The CRISP framework [23] provides the infrastructure required to scale data collection on the Franka Research 3 efficiently. Its Gymnasium-compatible `crisp_gym` interface outputs demonstrations directly in LeRobot format compatible with X-VLA [37], eliminating the ROS bag post-processing that typically constitutes a significant bottleneck in dataset growth. Its Cartesian impedance controller, validated on Franka Panda hardware [23], ensures that teleoperation demonstrations are recorded with consistent low-level dynamics, reducing the policy’s need to compensate for operator-induced variability in motor execution.

6.4 Dynamic Scene Graphs for Closed-Loop Manipulation

The scene graph implementation in this thesis produces a static snapshot of the workspace updated at approximately 0.5–1 Hz a design choice forced by the 1.5–3 s latency of the GPT-4 edge-labelling step in the construction pipeline (Section 3.2.3). This update frequency is adequate for static manipulations but is fundamentally insufficient for capturing the dynamic changes that determine whether a grasping action succeeds: the moment of gripper-object contact, the transition from open to closed gripper states, and the displacement of objects upon release. As argued in Section 5.3.1, the absence of high-frequency proximity feedback is a primary cause of the gripper-closure failures observed across all conditions.

Dynamic Scene Graphs (DSGs), introduced by Rosinol et al. [26] for mobile robot navigation, provide a framework for maintaining scene representations that update continuously with the state of the environment. For the tabletop manipulation setting, a lightweight DSG updating at the full 50 Hz control frequency is achievable: gripper-to-target distance can be estimated from the wrist-camera depth channel in under

5 ms using the 3D position estimates already computed for the static scene graph, without invoking the VLM captioning or edge-labelling steps that cause the most latency. This high-frequency proximity signal, injected as a dynamically updated node feature or as a dedicated soft-prompt token, would provide the action decoder with the precise contact-release trigger needed for reliable gripper closures directly addressing the failure mode in Section 5.3.1.

A more ambitious formulation, inspired by RoboEXP’s action-conditioned scene graphs (ACSGs) [8], is to maintain a predictive DSG that explicitly tracks the anticipated state of the scene graph after each proposed action sequence. Under this formulation, the policy plans in scene-graph space: it selects actions that will transition the graph from the observed configuration to a target configuration encoding task completion (e.g., the `mug` node transitioning from `on_table` to `grasped_by_gripper`). Klee et al. [11] demonstrate that GNN policies trained with graph-state targets generalise more robustly to novel object configurations than trajectory-imitation policies, suggesting that this approach would also improve long-horizon task completion rates.

6.5 Improved Sim-to-Real Transfer

The sim-to-real gap observed in this thesis characterised by substantially higher failure rates in physical deployment than simulation benchmarks predicted is consistent with the findings of Poria et al. [22], who identify distributional shift between simulation and real sensor characteristics as among the ten most critical open challenges in VLA research. Future work should target this gap through three complementary strategies.

The first is systematic **domain randomisation** during the simulation fine-tuning stages on LIBERO [16] and CALVIN [18]. Randomising object textures, lighting conditions, camera intrinsic parameters, and controller stiffness profiles within simulation creates a policy whose visual representations are less overfit to simulation-specific rendering artefacts and more robust to the photometric and geometric variations present in real sensor streams [22]. The RealSense depth noise that produced inaccurate scene graph node features in real deployment (Section 5.2.2) could be further mitigated by

adding structured depth noise models to the simulated camera pipeline, training the scene graph encoder to be robust to the specific noise characteristics of the RealSense sensors used in this work [5].

The second strategy is **tactile feedback integration**. The most persistent real-world failure mode gripper closure that does not secure the object is one that visual feedback alone is structurally limited in correcting, because the critical contact information is frequently occluded during the final approach phase. The Franka Research 3 provides native joint torque sensing at 1 kHz through the libfranka interface [23], which can serve as a binary contact-detection signal without additional hardware. Incorporating this signal as an additional observation modality either directly in the proprioceptive state vector or as a contact-state node feature in the scene graph would give the flow-matching action decoder [1] the feedback needed to detect and recover from failed contact attempts in real time.

The third strategy addresses the underlying **distributional shift** between the simulation training regime and real deployment through more sophisticated training paradigms. Two approaches deserve particular attention. **Instruction Fine-Tuning (IFT)** applied to the VLA policy analogous to its role in aligning LLMs with diverse instruction formats could reduce the reliance on in-distribution data by training the policy to follow structured, compositional instructions that generalise across layout variations [9]. In the robotics context, IFT would use a diverse set of automatically generated instruction-scene-graph pairs covering a wide range of object configurations, reducing the sensitivity to the specific 200-episode distribution collected for this thesis. **Batch (Offline) Reinforcement Learning** offers a complementary mechanism for handling distributional shift. Offline RL methods such as Conservative Q-Learning (CQL) or Implicit Q-Learning (IQL) can re-weight the training signal to prioritise demonstrations that are closest to real-world contact conditions, even when the offline dataset contains only simulated or imperfect teleoperation data. Crucially, offline RL can leverage the reward signal from gripper torque sensing the binary contact-detection signal available from the Franka hardware to identify and upweight demonstrations that achieved successful gripper closure without requiring

any additional human annotation. Together, IFT and offline RL provide principled approaches to the intrinsic distribution-shift issues that domain randomisation alone cannot fully resolve, since they operate on the policy optimisation objective rather than only on the data distribution.

Taken together, the four directions outlined in this section address each of the principal failure modes identified in Chapter 5 with targeted interventions grounded in the existing literature. Projector-based GNN alignment [21, 33] is the immediate priority, directly resolving the embedding incompatibility of Condition C. CRISP-enabled data scaling [23] and active collection [19] address the training-data bottleneck that limits gripper-transition learning. Dynamic scene graphs [26, 8] provide the high-frequency proximity feedback absent from the current static implementation. Domain randomisation [22, 16, 18] combined with tactile feedback integration [28] and instruction fine-tuning or offline RL narrows the sim-to-real gap that amplified all other failure modes in physical deployment. Together, these advances would move the system from the trajectory-imitation regime demonstrated here toward the spatially-grounded, contact-aware manipulation that the X-VLA soft-prompting architecture [37] is, in principle, capable of supporting.

Table 1: Supplementary Materials and Resources

Resource	Type	Link
Simulation Demos	Video	https://shorturl.at/7MSj1
Real-world demo videos for finetuned model	Video	https://shorturl.at/CY6am
Scene Graph building	Code	https://github.com/binsabit/ram_gd_sam
Model training	Code	https://shorturl.at/pv296
Dataset	Data	https://huggingface.co/datasets/binsabit/franka-fr3-manipulation
Fine-tuned X-VLA Model	Data	https://huggingface.co/binsabit/xvla-franka-fr3

Bibliography

- [1] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, et al. π_0 : A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.
- [2] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, et al. RT-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv preprint arXiv:2307.15818*, 2023.
- [3] Tianheng Cheng, Lin Song, Yixiao Ge, Wenyu Liu, Xinggang Wang, and Ying Shan. Yolo-world: Real-time open-vocabulary object detection, 2024.
- [4] Muhayy Ud Din, Waseem Akram, Lyes Saad Saoud, Jan Rosell, and Irfan Hus-sain. Vision language action models in robotic manipulation: A systematic re-view. *arXiv preprint arXiv:2502.00000*, 2025.
- [5] Qiao Gu, Alihusein Kuwajerwala, Sacha Morin, Krishna Murthy Jatavallabhula, et al. ConceptGraphs: Open-vocabulary 3D scene graphs for perception and planning. *arXiv preprint arXiv:2309.16650*, 2023.
- [6] William L. Hamilton, Rex Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [7] Nathan Hughes, Yun Chang, and Luca Carlone. Hydra: A real-time spatial perception system for 3D scene graph construction and optimization. *arXiv preprint arXiv:2201.13360*, 2022.
- [8] Hanxiao Jiang, Binghao Huang, Ruihai Wu, Zhuoran Li, et al. RoboEXP: Action-conditioned scene graph via interactive exploration for robotic manipulation. *arXiv preprint arXiv:2402.15487*, 2024.
- [9] Kento Kawaharazuka, Jihoon Oh, Jun Yamada, Ingmar Posner, and Yuke Zhu. Vision-language-action models for robotics: A review towards real-world applications. *IEEE Access*, 13:162467–162504, 2025.
- [10] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, et al. OpenVLA: An open-source vision-language-action model. *arXiv preprint arXiv:2406.09246*, 2024.

- [11] David Klee, Ondrej Biza, and Robert W. Platt. Efficient and interpretable robot manipulation with graph neural networks. *IEEE Robotics and Automation Letters*, 2022.
- [12] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3045–3059, 2021.
- [13] Fuhao Li, Wenxuan Song, Han Zhao, Jingbo Wang, Pengxiang Ding, Donglin Wang, Long Zeng, and Haoang Li. Spatial forcing: Implicit spatial representation alignment for vision-language-action model, 2025.
- [14] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 4582–4597, 2021.
- [15] Xuanlin Li et al. Evaluating real-world robot manipulation policies in simulation. *arXiv preprint arXiv:2405.05941*, 2024.
- [16] Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, Qiang Liu, Yuke Zhu, and Peter Stone. LIBERO: Benchmarking knowledge transfer for lifelong robot learning. *arXiv preprint arXiv:2306.03310*, 2023.
- [17] Yingbo Luo, Meibao Yao, and Xueming Xiao. Gcnt: Graph-based transformer policies for morphology-agnostic reinforcement learning. 2025.
- [18] Oier Mees, Lukas Hermann, Erick Rosete-Beas, and Wolfram Burgard. CALVIN: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks. In *IEEE Robotics and Automation Letters*, 2022.
- [19] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, et al. Octo: An open-source generalist robot policy. *arXiv preprint arXiv:2405.12213*, 2024.
- [20] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila

Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kosic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rameev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. Gpt-4 technical report, 2024.

- [21] Aleksandar Petrov, Emanuele La Malfa, et al. When do prompting and prefix-tuning work? a theory of capabilities and limitations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- [22] Soujanya Poria, Navonil Majumder, Chia-Yu Hung, Amir Ali Bagherzadeh, Chuan Li, Kenneth Kwok, Ziwei Wang, Cheston Tan, Jiajun Wu, and David Hsu. 10 open challenges steering the future of vision-language-action models, 2025.
- [23] Daniel San José Pro, Oliver Hausdörfer, Ralf Römer, Maximilian Dösch, Martin Schuck, and Angela P. Schoellig. Crisp – compliant ros2 controllers for learning-based manipulation policies and teleoperation, 2025.
- [24] Delin Qu, Haoming Song, Qizhi Chen, Yuanqi Yao, Xinyi Ye, Yan Ding, Zhigang Wang, JiaYuan Gu, Bin Zhao, Dong Wang, and Xuelong Li. Spatialvla: Exploring spatial representations for visual-language-action model. *arXiv preprint arXiv:2501.15830*, 2025.
- [25] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision, 2021.
- [26] Antoni Rosinol, Arjun Gupta, Marcus Abate, Jingnan Shi, and Luca Carlone. 3D dynamic scene graphs: Actionable spatial perception with places, objects, and humans. *arXiv preprint arXiv:2002.06289*, 2020.
- [27] Carmelo Sferrazza et al. Body transformer: Leveraging robot embodiment for policy learning. *arXiv preprint arXiv:2408.06316*, 2024.
- [28] Mustafa Shukor, Dana Aubakirova, Francesco Capuano, et al. SmolVLA: A vision-language-action model for affordable and efficient robotics. *arXiv preprint arXiv:2506.01844*, 2025.
- [29] Qi Sun et al. Emma-x: An embodied multimodal action model with grounded chain of thought and look-ahead spatial reasoning. *arXiv preprint arXiv:2412.11974*, 2024.
- [30] Kei Suzuki et al. Embedding morphology into transformers for cross-robot policy learning. *arXiv preprint arXiv:2603.00182*, 2026.
- [31] Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, Ziteng Wang, Rob Fergus, Yann LeCun, and Saining Xie. Cambrian-1: A fully open, vision-centric exploration of multimodal llms, 2024.
- [32] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *International Conference on Learning Representations (ICLR)*, 2018.

- [33] Bin Xiao, Haiping Wu, Weijian Xu, et al. Florence-2: Advancing a unified representation for a variety of vision tasks. *arXiv preprint arXiv:2311.06242*, 2024.
- [34] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations (ICLR)*, 2019.
- [35] Youcai Zhang, Xinyu Huang, Jinyu Ma, Zhaoyang Li, Zhaochuan Luo, Yanchun Xie, Yuzhuo Qin, Tong Luo, Yaqian Li, Shilong Liu, Yandong Guo, and Lei Zhang. Recognize anything: A strong image tagging model, 2023.
- [36] Qingqing Zhao et al. Cot-vla: Visual chain-of-thought reasoning for vision-language-action models. *arXiv preprint arXiv:2503.22020*, 2025. CVPR 2025.
- [37] Jinliang Zheng, Jianxiong Li, Zhihao Wang, Dongxiu Liu, Xirui Kang, Yuchun Feng, et al. X-VLA: Soft-prompted transformer as scalable cross-embodiment vision-language-action model. *arXiv preprint arXiv:2510.10274*, 2025. Accepted at ICLR 2026.