

# AutoSpatial: Visual-Language Reasoning for Social Robot Navigation through Efficient Spatial Reasoning Learning

Yangzhe Kong<sup>1</sup>, Daeun Song<sup>1</sup>, Jing Liang<sup>2</sup>, Dinesh Manocha<sup>2</sup>, Ziyu Yao<sup>1</sup>, and Xuesu Xiao<sup>1</sup>

**Abstract**—We present a novel method, AutoSpatial, an efficient approach with structured spatial grounding to enhance VLMs’ spatial reasoning. By combining minimal manual supervision with large-scale Visual Question-Answering (VQA) pairs auto-labeling, our approach tackles the challenge of VLMs’ limited spatial understanding in social navigation tasks. By applying a hierarchical two-round VQA strategy during training, AutoSpatial achieves both global and detailed understanding of scenarios, demonstrating more accurate spatial perception, movement prediction, Chain of Thought (CoT) reasoning, final action, and explanation compared to other SOTA approaches. These five components are essential for comprehensive social navigation reasoning. Our approach was evaluated using both expert systems (GPT-4o, Gemini 2.0 Flash, and Claude 3.5 Sonnet) that provided cross-validation scores and human evaluators who assigned relative rankings to compare model performances across four key aspects. Augmented by the enhanced spatial reasoning capabilities, AutoSpatial demonstrates substantial improvements by averaged cross-validation score from expert systems in: perception & prediction (up to 10.71%), reasoning (up to 16.26%), action (up to 20.50%), and explanation (up to 18.73%) compared to baseline models trained only on manually annotated data.

## I. INTRODUCTION

Social navigation poses unique challenges for robots in human environments, where avoiding obstacles is insufficient—robots must also interpret and respond to nuanced human behaviors [1], [2]. While traditional and learning-based methods have improved navigation [3], [4], they often fail in dynamic social scenarios due to a lack of explicit spatial reasoning [5].

Vision-Language Models (VLMs) offer promising reasoning capabilities [6], but their application to social navigation is hindered by two major limitations: (1) poor spatial grounding, as they lack explicit metrics like relative positions or movement trajectories [7], [8]; and (2) insufficient domain-specific data, since most annotated datasets focus on driving, not pedestrian interaction [9], [10].

**Main Results:** To address these limitations, we introduce AutoSpatial, an efficient approach that systematically enhances spatial reasoning capabilities in VLMs for social navigation reasoning, including human’s motion, relative position to the robot, and social interactions. AutoSpatial employs a structured question-answer paradigm to provide ground truth text labels for various visual observations and to train VLMs in perception, prediction, Chain-of-Thought (CoT) reasoning, final action, and explanation, enabling them to

<sup>1</sup>George Mason University { ykong7, dsong26, ziyuyao, xiao } @gmu.edu <sup>2</sup>University of Maryland, College Park { jingl, dmanocha}@umd.edu

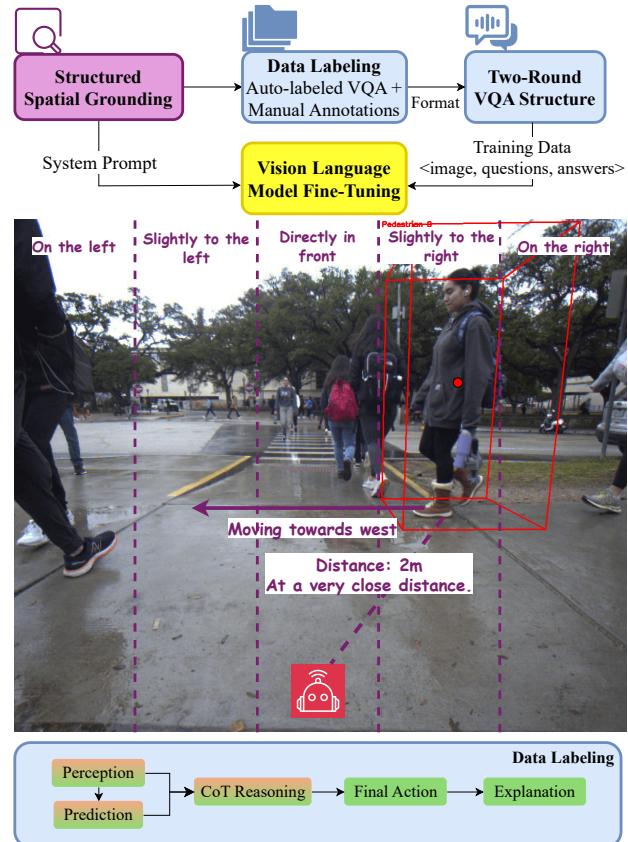


Fig. 1: Overview of the AutoSpatial approach. The approach consists of two primary components: **Structured Spatial Grounding** (purple), which standardizes positional (e.g., ‘Slightly to the right’ and ‘At a very close distance’ for the pedestrian in the bounding box) and directional descriptions (e.g., ‘Moving towards west’), and **Data Labeling** (blue), which combines auto-labeled Visual Question-Answer (VQA) pairs that focus on **perception**, **prediction**, and **CoT reasoning**, with manual annotations that focus on **perception**, **prediction**, **CoT reasoning**, **final action**, and **explanation**, and further formats them into a two-round VQA structure (see details in Fig. 2).

better interpret spatial relationships and social dynamics. Our approach leverages limited manual annotations in combination with large-scale auto-labeled data to significantly reduce annotation costs while maintaining high data quality. We conduct extensive evaluations on the SNEI benchmark [11] for spatial reasoning and social action estimation (Fig. 3 right).

- We propose a VLM-based approach, AutoSpatial, for enhancing spatial reasoning in social navigation;
- We introduce a novel automatic data labeling strategy with a two-round Visual Question-Answer (VQA) structure, including individual and group social motion understanding, that efficiently addresses domain-specific data scarcity while hierarchically enhancing spatial understanding;
- We evaluate our approach using averaged cross-validation scores from expert systems (GPT-4o, Gemini 2.0 Flash, Claude 3.5 Sonnet) and relative rankings from human evaluators. AutoSpatial outperforms other approaches in perception & prediction (up to 10.71%), reasoning (up to 16.26%), action (up to 20.50%), and explanation (up to 18.73%) in averaged cross validation of expert models.

## II. RELATED WORK

VLMs, which combine visual perception with natural language understanding, have demonstrated significant potential across various robotics tasks. We review related work in social robot navigation and VLMs for robotics tasks, with particular focus on their spatial reasoning capabilities and limitations in dynamic social environments.

### A. Social Robot Navigation

Robust social navigation is essential for robots in human-populated spaces [1], [2], [12]. Early approaches used hand-crafted rules, such as the Social Force Model [3] and proxemics-based methods [13], which simulate human-robot interactions via spatial constraints. Learning-based methods [4], [14]–[17] improve adaptability by leveraging large datasets to learn socially compliant behaviors. However, these models often lack explicit reasoning about nuanced human-robot dynamics shaped by unwritten social norms [11], [18].

### B. Spatial Reasoning Limitations of Robotics VLMs

Recent VLMs enable task and motion planning [19]–[21] and support human-robot interaction [16], [22]. Yet, two core challenges hinder their performance in social navigation:

**Lack of spatial reasoning:** VLMs lack explicit spatial encodings (e.g., trajectories, cost maps) [2], and struggle with ambiguous relative positioning [8], [18], [23].

**Insufficient domain-specific data:** Unlike other domains [24], social navigation lacks large-scale annotated datasets reflecting complex, implicit crowd behavior [9]–[11].

We address these challenges with AutoSpatial, which augments VLMs through structured spatial VQA training, combining minimal manual annotation and large-scale auto-labeling.

## III. METHODOLOGY

We present AutoSpatial, a systematic approach for enhancing VLMs' spatial reasoning capabilities in social navigation. As illustrated in Fig. 1, our approach is built upon

the structured spatial grounding that standardizes positional and directional descriptions, providing a consistent spatial terminology foundation. Guided by this, we label our training data through two complementary approaches: auto-labeled VQA pairs for large-scale spatial understanding and targeted manual annotations for complex social scenarios. These data are further formatted into a structured two-round VQA architecture (see details in Fig. 2), which enables hierarchical learning from basic spatial perception to complex reasoning. This methodology addresses our key observation that existing VLMs often struggle with fundamental spatial reasoning, leading to incorrect actions even when high-level reasoning appears sound.

### A. Structured Spatial Grounding

Our key innovation is to decompose spatial relationships into two clear, independent components: (1) positional description of humans, including angular position and distance, and (2) description of human-moving directions. This structured representation makes it easier to precisely locate and track pedestrians in relation to the robot. Traditional approaches, which describe spatial relationships between humans and robots in navigation, often lead to ambiguity, especially when addressing multiple agents [7].

For angular position of pedestrians, we divide the robot's field of view into five distinct zones: 'on the left', 'slightly to the left', 'directly in front', 'slightly to the right', and 'on the right' of the robot. This discretization provides a clear way to describe a pedestrian's relative position w.r.t. the robot, while maintaining sufficient granularity for navigation tasks. Each pedestrian is assigned to a zone based on their center point's horizontal pixel location in the image.

Distance estimation follows a five-level classification: 'very close' ( $< 3m$ ), 'close' ( $3 - 6m$ ), 'moderate' ( $6 - 9m$ ), 'far' ( $9 - 12m$ ), and 'very far' ( $> 12m$ ). By combining these two components, our approach can generate precise spatial descriptions such as "the pedestrian is slightly to the left of the robot at a moderate distance."

For movement directions, we adopt a unified coordinate system relative to the robot's frame, where the robot's forward direction is defined as north. This allows us to describe pedestrian movements using eight cardinal and intercardinal directions (N, NE, E, SE, S, SW, W, NW) plus a stationary state. This standardization eliminates the ambiguity often found in relative directional descriptions and provides a consistent reference frame for perception and prediction.

The structured spatial grounding plays a critical role in both our auto-labeling process and manual annotation. During auto-labeling, it provides a consistent framework for generating spatial VQA pairs, ensuring that perception, prediction, and reasoning outputs are systematically structured. For manual annotations, it serves as a guideline for annotators to maintain consistency in descriptions, reducing subjective variations.

### B. Data Labeling

Our training data combines auto-labeled VQA pairs and manual annotations to achieve both breadth in spatial reason-

### Input Image



### Round 1

**Question:** In this scene, there are one or multiple pedestrians marked with bounding boxes. Pedestrian 0 is marked in Red.

**[Task: Perception]** Where is the pedestrian in the bounding box located relative to the robot? In which direction is the pedestrian in the bounding box moving relative to the robot?

**[Task: Prediction]** Assess the potential future movement of agent(s) within the scene.

**[Task: Chain of Thoughts Reasoning]** Given the robot's perception and prediction, formulates a high-level natural language instruction.

**Answer:**

**[Task: Perception]**

Pedestrian 0

The pedestrian is slightly to the right of the robot, at a very close distance. The pedestrian is moving towards west.

**[Task: Prediction]**

Pedestrian 0

The pedestrian will continue moving towards west, and move across the robot's path from east to west.

**[Task: Chain of Thoughts Reasoning]**

Pedestrian 0

Given the close proximity of the person crossing the robot's path from east to west, the robot should slow down and stop if necessary, wait for the person to pass, and then continue.

### Round 2

#### Question:

Given separate descriptions of individual pedestrians from bounding boxes (Pedestrian 0, 1, etc.), synthesize them into a comprehensive scene description.

*[Descriptions of first 3 tasks are the same with Round 1]*

#### [Task: Final Action]

Comprise high-level natural language action commands.

#### [Task: Explanation]

Explain generally what the robot sees, what the robot does, and why.

#### Answer:

#### [Task: Perception]

The scene shows an active intersection of pathways with moderate pedestrian traffic. A person in a gray hoodie and black pants is walking towards west at a very close distance. Several other pedestrians are visible in the background, including a person with a red backpack directly in front of the robot at a medium distance moving towards north. The surface appears wet from recent rain.

#### [Task: Prediction]

The pedestrians will continue navigating through this intersection following an implicit order based on arrival time and proximity, with the person in gray likely to complete their crossing first due to their current position. The pedestrians in the background, including the one with the red backpack, will continue their current trajectories, mainly towards north.

#### [Task: Chain of Thoughts Reasoning]

This intersection shows typical pedestrian behavior where people navigate in multiple directions while maintaining an informal "first come, first cross" pattern. Given that the person in gray hoodie arrived at the intersection before the robot and is already in the process of crossing, social navigation norms and safe navigation in potentially slippery conditions suggest yielding to their movement.

#### [Task: Final Action]

Wait for the established crossing sequence to finish before proceeding.

#### [Task: Explanation]

I'm observing a typical pedestrian intersection where people are crossing in multiple directions. Since the person in gray arrived first and is already crossing, I'll wait my turn according to the natural flow of pedestrian traffic. This respects the unwritten "first arrive, first cross" norm that pedestrians naturally follow at such intersections.

Fig. 2: An example of the two-round VQA structure, where training data of each round follows the same format of <image, question, ground-truth answer>. **Round 1** is auto-labeled, focusing on perception, prediction, and CoT reasoning for individual pedestrians, while **Round 2** is optional and only available when the image input is manually annotated. It refines individual descriptions into a comprehensive scene-level understanding, incorporating group dynamics and higher-level spatial reasoning. The structured spatial grounding, though not explicitly shown here, is incorporated into the system prompt as part of the model input. Note that bounding boxes and color-pedestrian correspondence are only provided to the model during training.

ing and depth in scene understanding. As shown in Fig. 2, the upper part illustrates the auto-labeled VQA pairs, while the lower part represents the manually annotations.

1) *Auto-labeled VQA Pairs*: We develop a systematic approach to generate VQA pairs based on the CODA dataset [25]. This dataset provides a rich collection of real-world pedestrian interactions in urban environments captured from a robot's perspective, offering the necessary trajectory information and bounding box annotations required by our

approach. Note that this process relies solely on rule-based heuristics rather than complex learning-based models such as VLMs. For each frame, our approach processes three key aspects:

**Spatial Perception**: Spatial perception consists of two components: positional description and movement direction description. Precise spatial descriptions are generated using our structured spatial grounding discussed in Sec. III-A.

**Movement Prediction**: Our approach analyzes short-

term trajectories using a multi-frame sliding window. After smoothing velocity vectors to reduce noise, it categorizes movement patterns into five primary types: started motion, continued motion, left turns, right turns, and transitions to stationary states. Predictions are generated based on velocity and direction changes exceeding predefined thresholds.

**Interaction Description:** Our approach identifies, classifies, and describes various interaction patterns: 1. Trajectory conflicts: Detected by analyzing relative positions and velocities; 2. Path crossing: Categorized as west-to-east or east-to-west movements; 3. Pass-by scenarios: Classified based on relative trajectories; 4. Others: Categorized as scenarios without risks of collision.

2) *Manual Annotations:* We carefully select 72 challenging scenarios from the CODA dataset, focusing on cases where auto-labeled VQAs show limitations. These scenarios typically involve high pedestrian density environments, complex social grouping patterns, and situations requiring abstraction from individual to group-level understanding, where auto-labeled VQAs tend to exhibit significant errors or ambiguities.

Following the SNEI protocol [11], we also annotate data using the structured format of 5 tasks: Perception, Prediction, CoT Reasoning, Final Action, and Explanation. However, our approach differs in that we provide our structured spatial grounding guideline for annotators, as described in III-A. This ensures stricter spatial terminology and reduced ambiguity, which not only improves annotation consistency but also provides clearer supervision.

### C. Two-Round VQA Structure

On top of the generated data, we use a two-round VQA structure (Fig. 2) for scenes with manual annotations to facilitate hierarchical learning. For each selected scene, the first round consists of auto-labeled VQAs focusing on individual pedestrians, while the second round incorporates human annotations to refine group dynamics, social relationships, and collective movement patterns. This hierarchical approach enhances spatial reasoning by first establishing precise spatial-temporal understanding through structured auto-labeling, followed by higher-level scene interpretation via human annotations.

Specifically, we structure our data into consistent formats. For scenes with only auto-labeled data, we use a single-round format of `<image, round 1 question, round 1 ground-truth answer>`. For scenes with manual annotations, we structure the data as a two-round conversation: `<image, round 1 question, round 1 ground-truth answer>` followed by `<image, round 2 question, round 2 ground-truth answer>`.

This progressive structure yields synergistic effects beyond simple data augmentation. To explore this interplay, we conduct experiments (see Section IV-C) to assess how spatial reasoning capability acquisition and social context comprehension mutually reinforce each other.

## IV. EXPERIMENTS

We conduct extensive experiments to evaluate the effectiveness of AutoSpatial in enhancing spatial reasoning

for social navigation. Our evaluation focuses on assessing improvements in perception & prediction, reasoning, action, and explanation, comparing our approach against baseline models.

### A. Experimental Setup

We build our model based on LLaVA-1.6 OneVision-7B and conduct all experiments using a computing node with four A100 80GB GPUs. We follow the default instruction tuning procedure and hyperparameters from LLaVA [26] without any special loss weighting between different data sources. Our focus is on evaluating the effectiveness of our data labeling and structured VQA framework rather than sophisticated training techniques. For each training sample, we provide the model with an input image along with corresponding VQAs following our structured spatial grounding guideline. We use the AdamW optimizer with a learning rate of  $2e-5$  and batch size of 8 with gradient accumulation steps of 4. To comprehensively evaluate our approach, we establish a two-fold evaluation framework that assesses both basic spatial reasoning capabilities and high-level scene understanding.

For evaluating basic spatial reasoning, we employ an automated metric on the auto-labeled VQA pairs from CODA. Specifically, we first perform data balancing to mitigate directional bias by ensuring diverse pedestrian movement directions, then partition the dataset into training and test sets with a 90-10 split ratio. This results in 6,205 training frames (containing 16,530 pedestrian instances) and 573 test frames (containing 1,542 pedestrian instances). The test set evaluation focuses on two fundamental aspects of spatial understanding that can be objectively assessed: movement direction prediction and relative position estimation. Given the inherent variability in text generation, we implement a keyword-based evaluation metric, where model outputs are matched against predefined categorical labels. The scoring system awards 1.0 point for exactly correct answers and 0.5 points for approximately correct answers (e.g., predicting “northeast” when the ground truth is “north”).

For comprehensive scene understanding evaluation, we use the SNEI benchmark with both expert systems, such as GPT-4o, Gemini 2.0 Flash, and Claude 3.5 Sonnet, and 7 human evaluators, most of whom are graduate computer science students with backgrounds in computer vision and robotics. Our human evaluation focused on a randomly sampled subset of 49 scenarios from the SNEI dataset. Both expert systems and human evaluators assess four key aspects—Perception & Prediction, Reasoning, Action, and Explanation. Expert systems assign scores on a scale of 1-10, while human evaluators focus on our best method alongside two baseline models and provide relative rankings for each key aspect. We chose relative rankings for human evaluation rather than absolute scores because we believe it can mitigate potential subjective biases. For a fair comparison, we ensure all models receive identical input formats and evaluation criteria.

To analyze different aspects of our approach, we compare several model configurations. Our baseline models include

TABLE I: Cross-Validation Scores across Models for Different Aspects (Higher is Better).

Model ID <sup>1</sup>	CODA Benchmark	Gemini 2.0 flash				GPT-4o				Claude 3.5 Sonnet			
		P&P	R	A	E	P&P	R	A	E	P&P	R	A	E
AS-A72D1	<b>0.710</b>	6.824	<b>7.079</b>	7.292	7.258	6.838	<b>6.338</b>	6.290	<b>6.471</b>	6.652	<b>5.831</b>	<b>6.385</b>	<b>6.065</b>
AS-A72D2	0.582	<b>6.883</b>	7.070	7.272	7.264	<b>6.862</b>	6.221	<b>6.302</b>	6.413	<b>6.677</b>	5.671	6.351	6.000
AS-A72D3	0.573	6.775	7.048	<b>7.360</b>	<b>7.310</b>	6.813	6.077	6.250	6.269	6.674	5.591	6.292	5.957
AS-A72D5	0.483	6.627	6.878	7.280	7.270	6.739	5.947	6.114	6.139	6.504	5.490	6.151	5.831
AS-A36D1	0.686	6.661	6.996	7.277	7.227	6.764	6.135	6.167	6.247	6.584	5.613	6.234	5.936
AS-A0D1	0.680	6.484	6.660	6.859	6.756	6.582	5.517	5.733	5.920	6.551	5.415	4.831	5.375
LLaVA-M	0.404	6.246	6.089	6.108	6.157	6.662	5.815	5.791	6.351	6.031	5.065	5.342	5.302
Vanilla-LLaVA	0.376	6.243	6.083	6.102	6.154	6.486	5.492	5.458	6.092	5.809	4.815	4.803	5.409

<sup>1</sup> The model IDs follow the format ‘AS-AxxDy’, where ‘AS’ denotes AutoSpatial, ‘Axx’ indicates the number of manual annotations used (e.g., A72 for 72 annotations), and ‘Dy’ represents the downsampling ratio applied to the auto-labeled data (e.g., D1 for no additional downsampling).

TABLE II: Ablation Study

Model ID	CODA Benchmark	Gemini 2.0 flash				GPT-4o				Claude 3.5 Sonnet			
		P&P	R	A	E	P&P	R	A	E	P&P	R	A	E
AS-A72D1-P&P-VQA	0.709	<b>6.781</b>	7.014	7.203	7.153	6.832	6.194	6.185	6.004	<b>6.652</b>	5.717	6.375	6.010
AS-A72D1-R-VQA-1	0.518	6.717	7.008	7.246	7.206	6.601	5.541	5.767	5.542	6.621	5.684	6.323	5.982
AS-A72D1-R-VQA-2	0.520	6.751	<b>7.126</b>	<b>7.287</b>	<b>7.246</b>	<b>6.835</b>	<b>6.231</b>	<b>6.287</b>	<b>6.010</b>	6.584	<b>5.853</b>	<b>6.397</b>	<b>6.050</b>
AS-A72D1	0.710	6.824	7.079	7.292	7.258	6.838	6.338	6.290	6.065	6.652	5.831	6.385	6.065

TABLE III: Human Evaluator Rankings for Different Models.

Model	P&P	R	A	E
AS-A72D1	<b>1.19</b>	<b>1.70</b>	<b>1.43</b>	<b>1.55</b>
LLaVA-M	2.89	2.02	2.03	1.73
Vanilla LLaVA	2.93	2.03	2.55	2.53

the original LLaVA-1.6 OneVision-7B and LLaVA finetuned on only 72 manual annotations (referred to as LLaVA-M, where ‘M’ denotes ‘Manual annotations’). We deliberately focus on comparing with these 7B-parameter models to ensure a fair comparison, as larger proprietary models with hundreds of billions of parameters naturally demonstrate superior performance but operate at a fundamentally different scale, which prohibits inference onboard a mobile robot. We then evaluate variations of our approach with different training data compositions: models trained with the full dataset combining 72 manual annotations and CODA VQA data (18,615 pairs in total), and with different overall downsampling ratios (AS-A72D1 as 1, AS-A72D2 as 2, AS-A72D3 as 3, and AS-A72D5 as 5, see Tab. I) of CODA VQA data, to examine the impact of different data mixing ratios. To investigate the impact of manual annotation quantity, we also train models with reduced manual annotations (36 examples, see AS-A36D1 in Tab. I) and with no manual annotations (AS-A0D1). Additionally, we conduct ablation studies using only perception & prediction VQAs (AS-A72D1-P&P-VQA) and only reasoning VQAs with manual annotations (AS-A72D1-R-VQA-1 and AS-A72D1-R-VQA-2) to understand

their individual contributions.

### B. Results

Tab. I and Tab. III present a comprehensive evaluation of our approach through both automated metrics and human assessment. We evaluate all model variants on the CODA benchmark for basic spatial reasoning and using three state-of-the-art VLMs (Gemini 2.0 flash, GPT-4o, and Claude 3.5 Sonnet) for expert assessment on SNEI dataset, focusing on four aspects: Perception & Prediction (P&P), Reasoning (R), Action (A), and Explanation (E). Additionally, we conduct focused human evaluation on the three most representative models: our best-performing variant (AS-A72D1, AutoSpatial with 72 human Annotations and Downsampling ratio 1), LLaVA-M, and vanilla LLaVA.

Our proposed method (AS-A72D1) demonstrates consistent superior performance across all evaluation metrics. In the CODA benchmark focusing on spatial reasoning, AS-A72D1 achieves a score of 0.710, significantly outperforming both LLaVA-M (0.404) and vanilla LLaVA (0.376). This substantial improvement in spatial reasoning capabilities is further validated by both expert system evaluations and human assessment (Tab. III), where AS-A72D1 consistently outperforms the baseline models, particularly in perception & prediction and reasoning capabilities.

A key finding is that manual annotations and auto-labeled data complement each other, and combining them yields a synergetic effect, surpassing the performance of models trained on either data source alone. While LLaVA-M outperforms Vanilla-LLaVA, the gap remains substantial, indicating that manual annotations alone are insufficient to fully address

	<b>Perception</b> The robot's visual inputs show a busy street scene with multiple pedestrians. The most prominent pedestrian is a man in an orange shirt and white shoes, <b>walking towards the robot</b> . He is holding a white object, possibly a phone or a ticket. There are several other pedestrians in the background, including a woman in a brown sweater and blue jeans who is slightly to the right of the man in orange. The crowd density is moderate, with people walking in various directions. The pedestrians are dressed in casual clothing, and the scene appears to be during the day with clear skies.	<b>Prediction</b> The man in the orange shirt is likely to continue <b>walking towards the robot</b> , as he is already in motion and holding an object that suggests he is focused on his destination. The woman in the brown sweater and blue jeans is also likely to continue <b>walking in the same direction</b> , as she is not interacting with the man in orange and appears to be moving independently.	<b>Chain-of-Thought Reasoning</b> Given the man in the orange shirt's direction and the woman in the brown sweater's direction, the robot should <b>maintain a safe distance and avoid obstructing their paths</b> . The robot should also be aware of the surrounding pedestrians and <b>adjust its movement accordingly</b> to ensure safe navigation.	<b>Final Action</b> The robot should <b>continue to navigate the street while maintaining a safe distance from the pedestrians</b> . It should be aware of the man in the orange shirt and the woman in the brown sweater and <b>adjust its path as necessary to avoid any collisions or obstructions</b> .	<b>Explanation</b> The robot sees a busy street scene with multiple pedestrians. It perceives a man in an orange shirt and white shoes <b>walking towards it</b> , holding a white object. The robot also sees a woman in a brown sweater and blue jeans slightly to the right of the man in orange. The crowd density is moderate, and the pedestrians are dressed in casual clothing. The robot predicts that the man in the orange shirt will <b>continue walking towards it</b> , and the woman in the brown sweater will continue walking in the same direction. The robot <b>maintains a safe distance and adjusts its path as necessary to ensure safe navigation</b> .	The scene shows a busy pedestrian area with a high density of people. A person in an orange shirt is <b>crossing the street directly in front of the robot, moving towards west</b> . The person is at a close distance and is the main obstacle in the robot's path. Several pedestrians are visible on both sides of the street, with some walking towards east and others towards west, creating a busy street scene.
	The person in the orange shirt will continue crossing the street towards west. Other pedestrians will continue their current movements, with some walking towards east and others towards west.	Given the close proximity of the person in the orange shirt crossing directly in front of the robot, the robot must <b>stop and wait for the person to clear the path</b> . Once the person has crossed, the robot can proceed cautiously, monitoring the movements of other pedestrians on the left and right sides to ensure safe navigation.	<b>Stop and wait for the person in the orange shirt to clear the path</b>			

(a) Input Image

(b) Output comparison: LLaVA-M (left) v.s. AutoSpatial (right)

Fig. 3: While LLaVA-M suffers from **faulty spatial reasoning, leading to ambiguous or ineffective navigation decisions**, AutoSpatial exhibits **improved pedestrian identification and reasoning**, when augmented with auto-labeled VQA pairs.

the spatial reasoning deficiencies in VLMs. Notably, training on only auto-labeled VQA data (AS-A0D1) achieves only a marginal performance improvement over Vanilla-LLaVA and, in some cases, even underperforms, demonstrating that human supervision is still essential for grounding social interactions.

The robustness of our approach is demonstrated through several key findings. First, varying the downsampling ratio (AS-A72D1 through AS-A72D5) results in only minor performance differences, indicating that our method is relatively insensitive to the exact mixing proportions of the training data. Second, reducing manual annotations by half (AS-A36D1) leads to minimal performance degradation, suggesting that our approach can achieve strong results even with limited manual supervision.

### C. Ablation Studies

Tab. II examines the individual contributions of different components in our framework. Training with only perception and prediction VQA tasks (AS-A72D1-P&P-VQA) shows strong performance in the P&P metric but limited improvement in other aspects. Conversely, using only reasoning VQA tasks without the ground truth of perception and prediction VQA tasks (AS-A72D1-R-VQA-1) improves reasoning, action, and explanation scores but struggles with basic spatial understanding. Even with the ground truth of the perception and prediction VQA tasks provided in the question (AS-A72D1-R-VQA-2), the improvement in P&P remains limited.

These ablation results reveal a crucial insight: the synergistic effect of combining different types of VQAs with manual annotations. While individual VQA types can improve

specific aspects of performance, only the full combination achieves optimal results across all metrics.

### D. Qualitative Analysis

To better understand our model's capabilities and limitations, we analyze one representative case that highlights the differences between our approach and baseline models. In the scenario shown in Fig. 3, featuring a crowded sidewalk with multiple pedestrians, our model demonstrates superior spatial awareness by clearly identifying key actors and their relative positions using standardized spatial terminology (e.g., “moving towards west” instead of “walking towards/away from the camera”). More importantly, while baseline models generate generic advice like “continue to navigate the street while maintaining a safe distance”, our model provides specific, socially-aware instructions—recommending to wait for pedestrians to cross before proceeding, which better aligns with social norms.

However, our analysis also reveals areas for improvement. While the model demonstrates strong spatial reasoning, it occasionally misinterprets subtle human cues, such as gaze direction and body orientation, which are crucial for understanding pedestrian intent. Additionally, it sometimes reverts to ambiguous camera-relative descriptions used intensively by LLaVA instead of consistently applying the structured spatial grounding, likely due to data limitations during training.

This qualitative analysis validates our approach's effectiveness in improving spatial reasoning while also identifying clear directions for future enhancement. The improved spatial reasoning capabilities directly contribute to more specific and socially appropriate navigation decisions, even though op-

portunities remain for further refinement in human behavior understanding.

## V. CONCLUSIONS

In this paper, we present an efficient approach to enhance VLMs' spatial reasoning capabilities through a combination of auto-labeled data and minimal manual annotations. Our method demonstrates that with just 72 carefully selected manual annotations combined with structured auto-labeled data, we can achieve significant improvements in both basic spatial understanding and high-level scene comprehension.

Our experimental results reveal several key findings. First, the combination of auto-labeled VQAs focusing on individual pedestrians and manual annotations emphasizing group dynamics creates a synergistic effect that exceeds the performance of either approach alone. Second, our method shows remarkable robustness to various training configurations, maintaining strong performance even with reduced manual annotations. Third, qualitative analysis demonstrates that enhanced spatial reasoning capabilities directly contribute to more specific and socially appropriate navigation decisions.

## ACKNOWLEDGMENT

This work has taken place in the RobotiXX Laboratory at George Mason University. RobotiXX research is supported by National Science Foundation (NSF, 2350352), Army Research Office (ARO, W911NF2320004, W911NF2420027, W911NF2520011), Air Force Research Laboratory (AFRL), US Air Forces Central (AFCENT), Google DeepMind (GDM), Clearpath Robotics, Raytheon Technologies (RTX), Tangenta, Mason Innovation Exchange (MIX), and Walmart.

## REFERENCES

- [1] R. Mirsky, X. Xiao, J. Hart, and P. Stone, "Conflict avoidance in social navigation—a survey," *ACM Transactions on Human-Robot Interaction*, vol. 13, no. 1, pp. 1–36, 2024.
- [2] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh, "Core challenges of social robot navigation: A survey," *ACM Transactions on Human-Robot Interaction*, vol. 12, no. 3, pp. 1–39, 2023.
- [3] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Physical review E*, vol. 51, no. 5, p. 4282, 1995.
- [4] N. Hirose, D. Shah, A. Sridhar, and S. Levine, "Sacscon: Scalable autonomous control for social navigation," *IEEE Robotics and Automation Letters*, 2023.
- [5] R. A. Epstein, E. Z. Patai, J. B. Julian, and H. J. Spiers, "The cognitive map in humans: spatial navigation and beyond," *Nature neuroscience*, vol. 20, no. 11, pp. 1504–1513, 2017.
- [6] H. Liu, C. Li, Q. Wu, and Y. J. Lee, "Visual instruction tuning," *Advances in neural information processing systems*, vol. 36, 2024.
- [7] Z. Zhang, F. Hu, J. Lee, F. Shi, P. Kordjamshidi, J. Chai, and Z. Ma, "Do vision-language models represent space and how? evaluating spatial frame of reference under ambiguities," *arXiv preprint arXiv:2410.17385*, 2024.
- [8] A. Kamath, J. Hessel, and K.-W. Chang, "What's" up" with vision-language models? investigating their struggle with spatial reasoning," *arXiv preprint arXiv:2310.19785*, 2023.
- [9] C. Sima, K. Renz, K. Chitta, L. Chen, H. Zhang, C. Xie, J. Beßwenger, P. Luo, A. Geiger, and H. Li, "Drivelm: Driving with graph visual question answering," in *European Conference on Computer Vision*. Springer, 2024, pp. 256–274.
- [10] N. Nejatishahidin, M. R. Vongala, and J. Kosecka, "Structured spatial reasoning with open vocabulary object detectors," *arXiv preprint arXiv:2410.07394*, 2024.
- [11] A. Payandeh, D. Song, M. Nazeri, J. Liang, P. Mukherjee, A. H. Raj, Y. Kong, D. Manocha, and X. Xiao, "Social-llava: Enhancing robot navigation through human-language reasoning in social spaces," *arXiv preprint arXiv:2501.09024*, 2024.
- [12] A. Francis, C. Pérez-d'Arpino, C. Li, F. Xia, A. Alahi, R. Alami, A. Bera, A. Biswas, J. Biswas, R. Chandra, *et al.*, "Principles and guidelines for evaluating social robot navigation algorithms," *arXiv preprint arXiv:2306.16740*, 2023.
- [13] M. Daza, D. Barrios-Aranibar, J. Diaz-Amado, Y. Cardinale, and J. Vilasboas, "An approach of social navigation based on proxemics for crowded environments of humans and robots," *Micromachines*, vol. 12, no. 2, p. 193, 2021.
- [14] M. Luber, L. Spinello, J. Silva, and K. O. Arras, "Socially-aware robot navigation: A learning approach," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 902–907.
- [15] X. Xiao, B. Liu, G. Warnell, J. Fink, and P. Stone, "Appld: Adaptive planner parameter learning from demonstration," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4541–4547, 2020.
- [16] D. Song, J. Liang, A. Payandeh, A. H. Raj, X. Xiao, and D. Manocha, "Vlm-social-nav: Socially aware robot navigation through scoring using vision-language models," *IEEE Robotics and Automation Letters*, 2024.
- [17] A. H. Raj, Z. Hu, H. Karnan, R. Chandra, A. Payandeh, L. Mao, P. Stone, J. Biswas, and X. Xiao, "Rethinking social robot navigation: Leveraging the best of two worlds," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 16330–16337.
- [18] B. Chen, Z. Xu, S. Kirmani, B. Ichter, D. Sadigh, L. Guibas, and F. Xia, "Spatialvlm: Endowing vision-language models with spatial reasoning capabilities," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 14455–14465.
- [19] A. Brohan, Y. Chebotar, C. Finn, K. Hausman, A. Herzog, D. Ho, J. Ibarz, A. Irpan, E. Jang, R. Julian, *et al.*, "Do as i can, not as i say: Grounding language in robotic affordances," in *Conference on robot learning*. PMLR, 2023, pp. 287–318.
- [20] D. Song, J. Liang, X. Xiao, and D. Manocha, "Tgs: Trajectory generation and selection using vision language models in mapless outdoor environments," *arXiv preprint arXiv:2408.02454*, 2024.
- [21] S. Nasiriany, F. Xia, W. Yu, T. Xiao, J. Liang, I. Dasgupta, A. Xie, D. Driess, A. Wahid, Z. Xu, *et al.*, "Pivot: Iterative visual prompting elicits actionable knowledge for vlms," *arXiv preprint arXiv:2402.07872*, 2024.
- [22] S. Narasimhan, A. H. Tan, D. Choi, and G. Nejat, "Olivia-nav: An online lifelong vision language approach for mobile robot social navigation," *arXiv preprint arXiv:2409.13675*, 2024.
- [23] J. Wang, Y. Ming, Z. Shi, V. Vineet, X. Wang, Y. Li, and N. Joshi, "Is a picture worth a thousand words? delving into spatial reasoning for vision language models," *arXiv preprint arXiv:2406.14852*, 2024.
- [24] Y. Cao, Y. Kang, C. Wang, and L. Sun, "Instruction mining: Instruction data selection for tuning large language models," *arXiv preprint arXiv:2307.06290*, 2023.
- [25] A. Zhang, C. Eranki, C. Zhang, J.-H. Park, R. Hong, P. Kalyani, L. Kalyanaraman, A. Gamare, A. Bagad, M. Esteva, and J. Biswas, "Towards robust robot 3d perception in urban environments: The ut campus object dataset," 2023.
- [26] H. Liu, C. Li, Y. Li, B. Li, Y. Zhang, S. Shen, and Y. J. Lee, "Llava-next: Improved reasoning, ocr, and world knowledge," January 2024. [Online]. Available: <https://llava-vl.github.io/blog/2024-01-30-llava-next/>