

# GSON: A Group-based Social Navigation Framework with Large Multimodal Model

Shangyi Luo<sup>1</sup>, Peng Sun<sup>1</sup>, Ji Zhu<sup>1</sup>, Yuhong Deng<sup>2</sup>, Cunjun Yu<sup>2</sup>, Anxing Xiao<sup>2</sup>, Xueqian Wang<sup>1</sup>

**Abstract**—With the increasing presence of service robots and autonomous vehicles in human environments, navigation systems need to evolve beyond simple destination reach to incorporate social awareness. This paper introduces GSON, a novel group-based social navigation framework that leverages Large Multimodal Models (LMMs) to enhance robots’ social perception capabilities. Our approach uses visual prompting to enable zero-shot extraction of social relationships among pedestrians and integrates these results with robust pedestrian detection and tracking pipelines to overcome the inherent inference speed limitations of LMMs. The planning system incorporates a mid-level planner that sits between global path planning and local motion planning, effectively preserving both global context and reactive responsiveness while avoiding disruption of the predicted social group. We validate GSON through extensive real-world mobile robot navigation experiments involving complex social scenarios such as queuing, conversations, and photo sessions. Comparative results show that our system significantly outperforms existing navigation approaches in minimizing social perturbations while maintaining comparable performance on traditional navigation metrics.

**Index Terms**—Semantic Scene Understanding, AI-Enabled Robotics, Human-Centered Robotics, Vision-Based Navigation

## I. INTRODUCTION

THE growth of service robots has driven significant research on autonomous systems capable of navigating human-centered environments [1]–[3]. However, a critical gap exists in current navigation systems: while they excel at trajectory prediction and obstacle avoidance [4]–[8], they often fail to recognize and respect complex social contexts within crowds, such as photography sessions or queuing behaviors, as illustrated in Fig. 1. In the broader context of social robot navigation [9], [10], the goal is not only for the robot to reach its destination but also to interact appropriately with humans without degrading their experience. In this work, we focus on *developing a robot navigation system that can understand and respect contextual social interactions*, which we represent as social groups. These social groups represent meaningful



Fig. 1. Example of a social navigation scenario: Rather than walking between two individuals in a queue, the robot correctly identifies the social context and navigates around the queue to avoid disturbance.

patterns of semantic interaction among individuals within the crowd.

The core issue in building such a socially aware navigation system is how to accurately *identify* contextual social interactions in dynamic and unpredictable crowds, and *exploit* these social interactions to guide the planning system. For perception, accurately identifying contextual social interactions in dynamic, unpredictable crowds requires semantic understanding beyond simple proximity or velocity patterns. Traditional methods based on predefined rules or in-domain learning [4], [5], [11]–[16] struggle in open-world environments due to data scarcity and high visual variability. Regarding planning, generating socially appropriate motions that respect these interactions demands planning at the right abstraction level—traditional local planners often fall into local minima due to short planning horizons [17]–[20], while global path planners cannot adapt quickly to dynamic social interactions. In contrast to modular systems, end-to-end solutions aim to bypass the traditional perception–planning decomposition [21]–[24]. However, these approaches often overfit demonstrations and struggle to interpret complex contextual information. More recently, some works have explored the use of large multimodal models (LMMs) to select sampled trajectories [25] or high-level directions [26]–[28] to improve social awareness without explicit social interactions modeling. However, these methods are inefficient as they require querying the LMM at every planning step, and the resulting social behaviors are often not interpretable.

To systematically address these challenges, we propose GSON that combines the strengths of Large Multimodal Models (LMMs) with a robust navigation system. Our key insight is to utilize the visual understanding ability of the Large Multimodal Model (LMM) to enable zero-shot reasoning of

Manuscript received: April, 7, 2025; Revised June, 26, 2025; Accepted July, 21, 2025.

This paper was recommended for publication by Editor Markus Vincze upon evaluation of the Associate Editor and Reviewers’ comments. This work was supported by the National Key R&D Program of China (2022YFB4701400/4701402), and National Natural Science Foundation of China under Grant 62293545. (Shangyi Luo, Peng Sun, and Ji Zhu are co-first authors.) (Corresponding authors: Anxing Xiao; Xueqian Wang.)

<sup>1</sup>Center for Artificial Intelligence and Robotics, Tsinghua Shenzhen International Graduate School, Shenzhen, China. Correspondence to: wang.xq@sz.tsinghua.edu.cn

<sup>2</sup>School of Computing, National University of Singapore, Singapore. Correspondence to: anxingx@comp.nus.edu.sg

Digital Object Identifier (DOI): see top of this page.

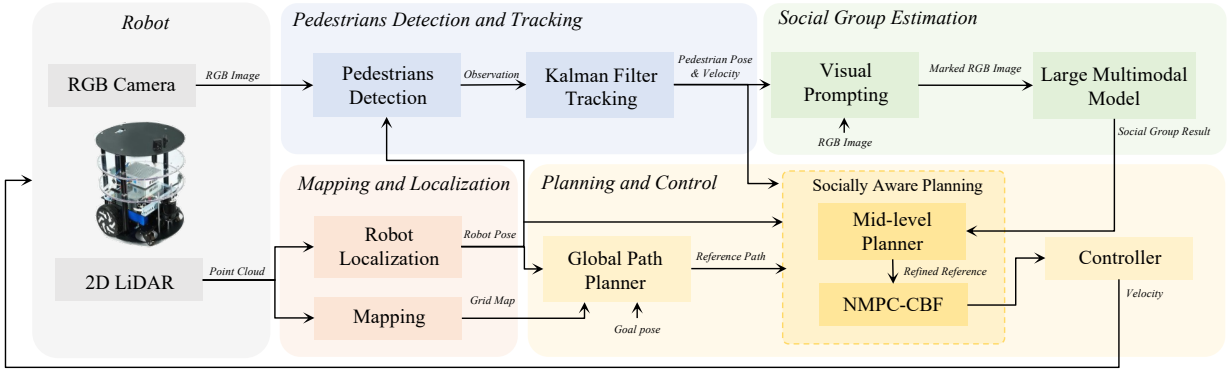


Fig. 2. Overall Framework for GSON: GSON integrates RGB cameras and 2D LiDAR data for dynamic crowd perception. The large multimodal model allows group estimation and enables the planning system to perform socially-aware planning for navigation.

social groups. The LMM serves as an auxiliary to our perception pipeline, which supports reliable detection and tracking, to enhance the estimation of social groups. Our design utilizes LMMs only for low-frequency semantic understanding, while high-frequency tracking and planning are handled without querying large models, leading to a more practical sense. In addition, we adapted a mid-level planner as a bridge between global path planning and local motion planning, enabling a balance between global awareness of social interactions and reactive responsiveness.

Overall, we develop a group-based social navigation framework GSON that can identify and exploit contextual social interactions. We conduct extensive experiments with several daily social interaction scenarios in the real world. Comparative experiments indicate that GSON outperforms all baseline methods in terms of minimizing perturbations to the social group.

## II. RELATED WORK

### A. Social Group Detection in the Crowd.

Significant research has focused on social group detection in crowded pedestrian environments, which can be approached from both exo-centric and ego-centric perspectives. The exo-centric approach [29] employs external sensors to provide a comprehensive top-down view for crowd analysis. Various methods have been proposed for exo-centric group detection to support motion tracking [30], [31], trajectory prediction [12], [32], group activity recognition [33], [34], and crowd analysis [35]–[37]. However, this approach has two key limitations: it relies on external sensors and offline processing, and it often overlooks the rich semantics of crowd interactions by focusing mainly on motion patterns or domain-specific overfitting. For ego-centric group detection, prior approaches to crowd grouping have employed clustering based on motion patterns [13], [14], [38] and distance [14], [39], [40], or learning group information in a specific dataset using graph neural network [15], [16], predicting cost function [41], or leveraging F-formation theories [42]. While these methods excel at robot-centric perception, they still face limitations in capturing semantic context and tend to overfit specific scenarios. In GSON, we leverage the zero-shot visual reasoning capabilities of the Large Multimodal Model (LMM) to enable semantic reasoning of social groups given ego-centric observation.

### B. Social Robot Navigation

Previous studies have proposed various methods to improve collision avoidance in robot navigation [5], [6], [8], [18]. With the increasing number of mobile robots in human environments, there’s a need beyond collision avoidance to improve social awareness. Traditionally, researchers have applied physics-inspired approaches such as the Social Force Model [43] to robot navigation [4], [44]. However, SFM-based methods only implicitly represent social awareness as pedestrian motion and suffer from limited generalization due to parameter selection. Group-based navigation strategies integrate dynamic human grouping into social navigation. Methods that combine group information with model-based planners [14], [38], [40] or reinforcement learning controllers [16], [45] show that representing human groups improves robot safety and social awareness. However, most group-based methods define groups only in terms of position and velocity, ignoring the semantic social interaction between pedestrians. Another line of research uses imitation learning [21]–[24] to implicitly learn social awareness from demonstrations. However, due to the data scarcity problem, the learned policies capture only limited behaviors in the dataset. Recent advances in foundation models have integrated large multimodal models to generate high-level actions [27], [28], score trajectories [25] or directions [26], and distill cost into imitation learning [46]. These methods are inefficient, requiring LMM queries at every planning step, with often uninterpretable social behaviors. Our approach uses LMMs only for low-frequency semantic social group detection when new pedestrians appear in observation, handling high-frequency tracking and planning without querying large models, creating a more practical system.

## III. PROBLEM FORMULATION

We consider the problem of social robot navigation in a 2D Euclidean space  $\mathbb{R}^2$ . This space is populated by  $n$  dynamic pedestrians  $H_i$  with a specific social structure. We consider the social structure as  $m$  social groups  $G_j$ . The robot is initialized with a 2D occupancy map  $M$  at the start position  $S_s$  and a target position  $S_g$  in the map. Given the 2D lidar point cloud  $P^t$  and the RGB image observation  $I^t$ , the goal is to enable the robot to safely navigate to the goal position while minimizing the time it spends trespassing on any social group. The robot

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 3. The framework of pedestrians detection and tracking: It combines 2D LiDAR-based foot detection with human detection from RGB images

must infer the social groups of pedestrians from the limited observation and generate an appropriate trajectory  $T$ . We focus on designing the method to estimate the social group with the aid of the Large Multimodal Model and integrate the estimated result into the navigation system.

## IV. METHOD

### A. Overview

The overall architecture of GSON is illustrated in Fig. 2. The robot is equipped with two RGB cameras, a 2D lidar sensor, and a mobile base with a differential drive. The system has two critical modules: the social group estimation module and the socially aware planning module. The social group estimation module takes as input the current observation RGB image  $I^t$  (formed by merging images from multiple cameras), 2D point clouds  $P^t$ , and the current robot pose  $S^t$  from the SLAM module. It generates the estimated social group score  $\hat{G}$ , quantifying the likelihood that individuals belong to the same social group. The estimated social group score is fed into the socially aware planning modules, which consist of a global path planner, a mid-level planner, and a local motion planner. The global path planner computes a collision-free path from the robot’s starting position to the target position, taking into account the environment information and static obstacles such as walls, furniture, and other permanent structures. The mid-level planner utilizes the estimated social group information and generates intermediate references that guide the local planner in navigating through the crowd while minimizing the risk of intruding upon the social groups. The local motion planner then ensures precise and safe navigation based on the intermediate references provided by the mid-level planner.

### B. Social Group Estimation with Large Multimodal Models

The core idea of GSON is to model contextual social interactions in dynamic and unpredictable crowds using social groups, which then guide the planning system. In this context, a social group refers to two or more individuals who are interacting, conversing, or moving together with shared intentions or social connections. To identify these social groups, we first build a robust pipeline for pedestrian detection and tracking. After obtaining detection and tracking results of pedestrians, identifying social groups requires commonsense knowledge about social interaction. Therefore, we leverage the multimodal reasoning capabilities of the LMM and employ visual prompting to identify social groups.

**Pedestrian Detection and Tracking:** We design a robust pipeline for pedestrian detection and tracking. Given the point cloud data  $P^t$  of the environment, we first use an off-the-shelf

model proposed in [47] for preliminary pedestrian detection. This preliminary detection provides potential pedestrian locations in the 3D point cloud but may contain false positives.

To improve accuracy and remove noise, we combine information from RGB images using YOLO-V5 [48]. Specifically, the preliminary detected pedestrian coordinates  $p_i^{3D} = [x'_i, y'_i, z'_i]$  from the point cloud are projected into the camera coordinate system. We then compare these projections with the bounding box results from YOLO-V5, retaining only those detections that align with YOLO-V5’s pedestrian identifications. For each YOLO-V5 detection box  $B_j$  (where  $j \in \{1, \dots, N\}$ ), we define the final selected coordinate  $h_j^*$  by:

$$h_j^* = \arg \min_{h_i \in \mathcal{H}_j^{proj}} \|\pi(h_i) - \mathbf{c}_j\|_2, \quad (1)$$

where  $\mathcal{H}_j^{proj} = \{p_i^{3D} \mid \pi(p_i^{3D}) \in B_j\}$  is the set of point cloud coordinates whose projections fall within the bounding box  $B_j$ ,  $\pi: \mathbb{R}^3 \rightarrow \mathbb{R}^2$  denotes the projection from 3D point cloud coordinates to 2D image coordinates,  $\mathbf{c}_j = (c_j^x, c_j^y)$  is the center coordinates of bounding box  $B_j$ ,  $\|\cdot\|_2$  represents the Euclidean distance.

After this filtering process, we obtain refined pedestrian positions in 2D coordinates  $\mathbf{p}_i = [x_i, y_i]$ , where the  $z$  dimension is omitted as we focus on the ground plane for navigation purposes. Finally, we apply the Hungarian algorithm for data association and Kalman filtering for estimating and predicting the position of each pedestrian with a unique tracking ID, enabling robust tracking for subsequent social group detection. Each pedestrian is modeled using a constant velocity motion model, represented by a state vector  $\mathbf{s}_k = [x_k, y_k, \dot{x}_k, \dot{y}_k]^\top$ , which encodes both position and velocity in 2D space. At each timestep, Kalman filters first predict target positions using linear constant-velocity dynamics, then the Hungarian algorithm optimally matches these predictions with new detections by minimizing the total Euclidean distance between corresponding points. Unmatched detections initialize new tracks with zero initial velocity, while tracks failing to match consecutively for several frames are terminated to prevent ghost tracking.

This pipeline enables robust pedestrian detection and tracking, providing bounding boxes and tracking IDs necessary for subsequent social group detection. In addition to supporting group detection, the tracking results—specifically each pedestrian’s estimated position and velocity—serve as crucial input to the planning module. These motion states help the planner reason about group formations and predict pedestrian movement, enabling proactive, socially aware trajectory generation that avoids disrupting group interactions.

**Social Group Detection with Visual Prompting:** We leverage a LMM<sup>1</sup> to identify social groups within crowds by analyzing human behaviors and interactions. Our approach consists of three main steps: image annotation, LMM querying, and social group relationship extraction as shown in Fig. 4. First, given bounding boxes and tracking IDs of pedestrians, We use the mask that results from YOLO-V5, creating an annotated image where each individual is marked with their bounding box,

<sup>1</sup>specifically gpt-4o-2024-08-06

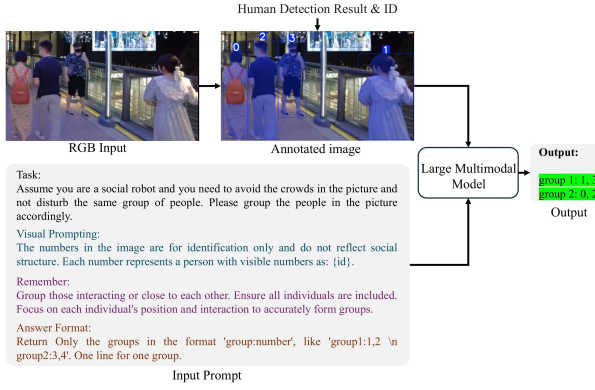


Fig. 4. With visual prompting, the LMM identifies social interactions among people and infers social groupings for downstream planning modules. segmentation mask, and unique tracking ID  $h_i$ . These visual marks enhance the LMM's reasoning. Second, we query the LMM with a structured input consisting of the annotated image and a text prompt. In the text prompt, we provide the LMM with contextual information, including the expected behavior of the LMM, the format of queries and responses, and cues to assist the LMM in accurately identifying social groups. Finally, we process the LMM's output using a parsing script that extracts the predicted social group relationships. These relationships are then integrated into our tracking system's data structure, maintaining an up-to-date record of social groups over time. Formally, we can express this process as:

$$G = \text{LMM}(A(I, \mathcal{H}^{\text{track}}), T), \quad (2)$$

where  $I$  denotes the raw input image frame,  $\mathcal{H}^{\text{track}} = \{p_i\}_{i=1}^N$  represents the filtered tracking states, and  $A$  is our annotation function that applies masks, bounding boxes, and IDs to the image,  $T$  is the text prompt, and  $G$  is the resulting social group assignments. Our approach enables the system to identify social groups by analyzing visual cues and spatial relationships between individuals.

Due to the computational constraints of LMM, we implement a keyframe-based approach to reduce the frequency of LMM queries while maintaining accurate social group tracking. Instead of analyzing every frame, we maintain a keyframe that captures the most informative scene within a 1-second window - specifically, the frame containing the highest number of visible people. Each keyframe stores both the image and corresponding metadata (position and velocity) for all detected individuals. The LMM is only queried when this keyframe is updated, which occurs either when a more representative frame is found or when the current keyframe expires after 1 second.

### C. Socially Aware Planning Module

To leverage social group identification results for guiding robot planning, we propose a socially aware planning module, which consists of three levels: global path planning, mid-level planning, and local motion planning.

**Global Path Planning:** The first step of our planning module is to compute a collision-free path from the robot's current position to the target location. We begin by inflating the occupancy map  $M$  to account for the robot's size and safety buffer,




Fig. 5. The framework of socially aware planning modules.

resulting in the grid-based cost map  $C$ . This inflated cost map provides a conservative estimate of obstacle boundaries. We then implement the A\* algorithm on  $C$ , further optimized with Jump Point Search (JPS) [49]. JPS improves the efficiency of path planning in large-scale grid environments by reducing unnecessary node explorations. The generated global path  $R$  serves as the reference for the mid-level and local planners.

**Mid-level planning:** Given the global reference path  $R$ , we introduce a mid-level planning algorithm that leverages social groups to provide socially aware guidance in the planning module.

Once the LMM module has produced the social group assignments  $G$ , we proceed to estimate the spatial extent of each group. Specifically, we compute the convex hull of each group, denoted as  $G'$ , to represent the minimal boundary enclosing all its members. To further approximate the shape and spatial influence of each group, we fit an ellipse to the boundary defined by  $G'$ . The resulting enclosed two-dimensional regions are defined as social spaces  $S$ , as shown in Fig. 5. Once the global reference path  $R = \{p_0, \dots, p_N\}$  intersects a social space  $S \subset C^+$ , the robot's current state is set as a new start  $S'_s$ . Let  $p_k$  be the last waypoint within  $S$ . From  $p_k$ , we move a fixed arc-length  $h$  along  $R$  to define a sub-goal  $S'_g$ . If the remaining path is shorter than  $h$ , set  $S'_g = p_N$ ; otherwise, find the smallest  $m$  such that the arc length from  $p_k$  to  $p_m$  is at least  $h$ , and interpolate between  $p_{m-1}$  and  $p_m$  to locate  $S'_g$ . This ensures  $S'_g$  lies on  $R$  but outside  $S$ , so only the segment  $S'_s S'_g$  requires re-planning.

To compute a socially compliant trajectory under the updated cost constraints, we employ the Batch Informed Trees (BIT\*) algorithm [50]. This process yields a revised path, referred to as the *CrowdFreePath*  $F$ , which effectively avoids traversing through socially sensitive regions and minimizes potential interactions with social groups. The socially-aware mid-level planning incorporates social group information, enabling dynamic updates to the cost map. This ensures the robot navigates safely while avoiding significant impacts on social groups.

**Local Motion Planning:** Given the socially aware path  $F$  from the mid-level planner and current pedestrian states (positions and velocities), the local planner computes dynamically feasible trajectories that track  $F$  and ensure real-time, collision-free navigation by extrapolating pedestrian motion using a constant-velocity model. To this end, we develop a local planning framework based on Nonlinear Model Predictive Control with Control Barrier Functions (NMPC-CBF) [51], which explicitly accounts for both the positions and velocities of nearby pedestrians when computing control actions. We

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.



Fig. 6. Sampled scenario visualizations from the dataset.

define a Control Barrier Function (CBF) as follows:

$$h^i(x) = (x - x_p^i)^2 + (y - y_p^i)^2 - d_{\text{safe}}^2, \quad (3)$$

where  $\mathbf{x}_p = [x_p, y_p]^T$  denotes the predicted pedestrian position based on their current velocity, and  $d_{\text{safe}}$  represents a safety radius. The local planning problem is then formulated as a constrained NMPC optimization similar to [52]:

$$\min_{\{\mathbf{x}_k, \mathbf{u}_k\}} \|\mathbf{x}_N - \mathbf{x}_{\text{goal}}\|_{P_f}^2 + \sum_{k=0}^{N-1} \|\mathbf{u}_k\|_{Q_u}^2 \quad (4)$$

$$\text{s.t. } \mathbf{x}_{k+1} = f(\mathbf{x}_k, \mathbf{u}_k) \quad (5)$$

$$\mathbf{x}_0 = \mathbf{x}_{\text{init}} \quad (6)$$

$$\mathbf{x}_k \in \mathcal{X}, \mathbf{u}_k \in \mathcal{U} \quad (7)$$

$$\Delta h_{\text{ob}}^i(\mathbf{x}_k, \mathbf{u}_k) + \lambda_k h_{\text{ob}}^i(\mathbf{x}_k) \geq 0. \quad (8)$$

where the last constraint ensures forward invariance of the safe set defined by each pedestrian, which evolves dynamically with their predicted trajectories.

## V. EXPERIMENTS

In this section, we aim to investigate the following research questions through a series of experiments: (1) *Can Large Multimodal Models effectively identify social groups among pedestrians?* (2) *How well does GSON perform in real-world scenarios that require social awareness?*

### A. Evaluation of LMM Social Group Prediction

To evaluate LMMs in predicting social groups from RGB images, we selected 100 scenarios from datasets [53], [54] and city walk videos, covering interactions such as queuing, walking, talking, and photographing. Human annotators provided group labels based on marked images, with the most agreed-upon annotations as ground truth. Example scenarios are shown in Fig. 6. On average, each scenario contains 6.0 pedestrians and 3.5 distinct social groups. We evaluated six LMMs for this task: four commercial models: Gemini 2.0 Flash, Gemini 2.5 Pro, Claude 3.7 Sonnet, GPT-4o, and two open-source models: Qwen2.5-VL (32B and 7B variants). Each scenario was tested 10 times with each model. The results were classified into four categories based on the predicted groupings: (1) *Accurate*: Group matches the ground truth (preferred), (2) *Miss*: Omits some individuals from the ground-truth group (acceptable), (3) *Extra*: Includes individuals not in




Fig. 7. Grouping results of LMMs and traditional clustering in the dataset the ground-truth group (not preferred), (4) *Error*: Both misses and includes incorrect individuals (not acceptable).

Fig. 7 shows the results for each model. All four LMMs achieved acceptable grouping rates above 67%, with the lowest from Qwen2.5-VL-7B (53.6% *Accurate* + 14.8% *Miss*) and the highest from GPT-4o at 86.6% (77.8% *Accurate* + 8.8% *Miss*). Given the complexity and diversity of the dataset, these results highlight the capability of Large Multimodal Models to identify social groups among pedestrians. Inference time varies across models: Gemini 2.5 Pro and Claude 3.7 Sonnet take around 8s, GPT-4o and Gemini 2.0 Flash about 4s, Qwen2.5-VL-32B around 2s, and Qwen2.5-VL-7B is the fastest, around 1s. In our pipeline, group detection doesn't require high frequency, making GPT-4o and Qwen2.5-VL-32B well-suited for our system. For comparison, we evaluated two classical clustering baselines using DBSCAN on 2D and 3D coordinates. Both yielded significantly lower accuracy (10.8–14.8%) and higher extra grouping errors (41.9–46.5%), underscoring the limitations of proximity-based methods and the benefits of semantic reasoning in LMMs.

### B. Real-world Evaluation of GSON

1) *Experimental Setup*: We tested GSON in real-world settings using a two-wheeled nonholonomic mobile robot equipped with a 2D LiDAR and two RGB cameras. The experiments focused on short-range navigation around people engaged in social interactions like queuing, photographing, talking, and walking, as shown in Fig. 9. Each scenario was run three times, with variations in participants' initial positions, paths, and speeds to capture real-world complexity.

2) *Baseline and Evaluation Metrics*: We compared GSON's performance against several baselines, including classical planners like TEB and TEB enhanced with our group estimation, as well as recent open-source approaches: SACSON [23], an imitation learning-based method; GMPC [14], a group-based motion prediction controller; and RLCA [55], a distributed deep reinforcement learning method. Note that we excluded several LMM-based methods [25]–[28] from the comparison because they are not open source and did not perform reliably in our pilot experiments with our own implementations. For evaluation, traditional navigation metrics include roughness, curvature, jerk, and runtime [9], [56]. To assess social navigation, we introduced metrics to capture the robot's impact on individuals and groups. Following prior guidelines [3], an individual was considered disturbed if the

## IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.




Fig. 8. **Real-world experiment results.** GSON performs comparably on traditional navigation metrics—exhibiting smooth curvature, low roughness, and slightly higher jerk—while excelling in key social navigation metrics. These results highlight its effectiveness in navigating complex social interactions.

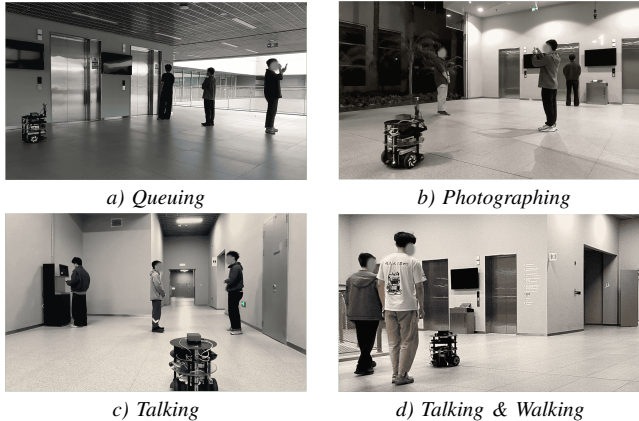


Fig. 9. Test scenarios in real-world experiments.

robot approached within 1.2 meters directly in front. For groups, a disturbance was defined by the robot entering the convex hull formed by group members. We measured both the duration of group disturbance and the “comfort distance”—the average minimum distance between the robot and social groups.

3) *Result and Analysis:* The results are shown in Fig. 8. GSON performs comparably to other methods in traditional navigation metrics like curvature, jerk, and roughness, achieving smooth and efficient motion with only a slight increase in jerk and runtime. Our speed is slightly lower than the comparison methods, mainly due to two reasons. First, calling the LMM adds at least 1 second of response time. Second, our method actively avoids groups, even if it means taking longer detours. However, the time is not significantly longer since the mid-level planning chooses the optimal path that does not interrupt the group. In terms of social navigation, GSON clearly outperforms baselines. It records the lowest Time of Disturbing Individuals and Groups, showing minimal disruption to pedestrians and social groups. Notably, even compared to TEB+GROUP ESTIMATE, GSON maintains better group awareness. This is due to the short planning horizon of the TEB planner—when group members are spaced far apart, as in our scenarios, *TEB+GROUP ESTIMATE* can become stuck in local minima. GSON also achieves the highest Comfort Distance, maintaining more personal space from social groups—crucial in semantically rich environments. This demonstrates its strong respect for social norms during navigation. Overall, GSON balances traditional performance with superior social compliance, making it highly effective for real-world deployment in human-centric environments.

### C. Long-Range Social Navigation Demonstration

We conducted a long-range demonstration in a structured, large-scale building comprising a hall, corridors, and an open

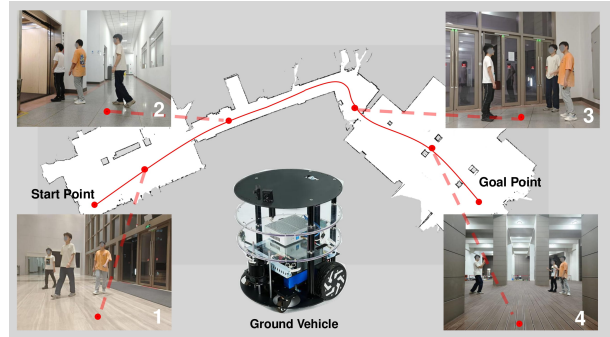


Fig. 10. The setting of the long-range social navigation task.

square. This extended test integrated four smaller scenarios into a complex setting, where the robot was tasked with identifying social interactions and navigating through these environments (see Fig. 10). The start and goal points were defined on opposite sides of the map.

Fig. 11 shows keyframes from the navigation process. At approximately 10 seconds, the robot encounters a small group of individuals walking ahead. Leveraging its perception and planning modules, it identifies the group and adjusts its trajectory, successfully avoiding the two people. At around 60 seconds, the robot approaches the entrance of an elevator where a queue has formed. It recognizes the queue and reroutes to avoid cutting through it, ensuring it does not disturb the individuals waiting. By 146 seconds, while inside a narrow corridor, the robot detects a group engaged in conversation. Based on the group estimation, it computes a seclusive path to bypass the group and maintain smooth navigation through the confined space. Finally, at around 191 seconds, upon entering an outdoor square, the robot encounters individuals engaged in photography. It identifies their positions and reroutes its path to go around them, ensuring progress while respecting their activity. In summary, this long-range demonstration validated the robot’s ability to navigate diverse, real-world social environments while maintaining seamless and respectful interactions. Its consistent performance across various challenging scenarios—from avoiding moving individuals to navigating around groups in tight or open areas—highlights its advanced reasoning about social interactions.

### D. Group-Aware Social Navigation in Dense Crowds: Limitations and Future Directions

To further investigate the ability of our system in high-density social environments and provide insights, we conducted additional experiments in three real-world scenarios: a university canteen during peak lunch hours, the entrance of a teaching building near an elevator, and the main campus gate. Each involved an average of over 13 pedestrians exhibiting

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

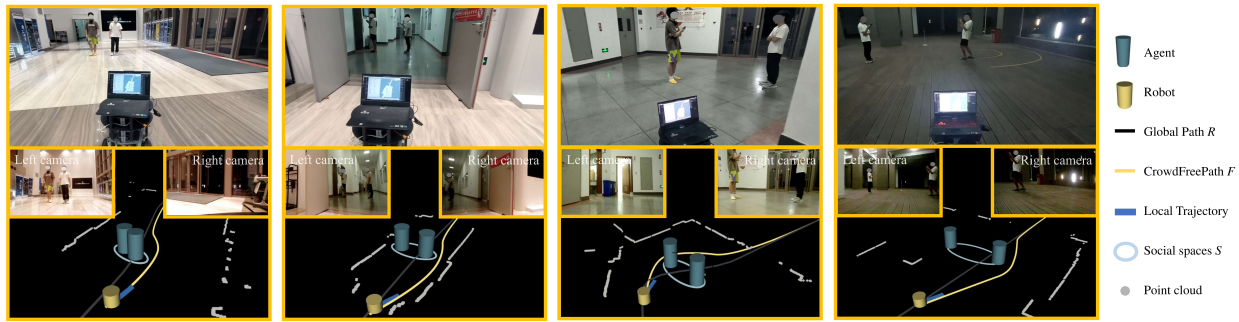


Fig. 11. **Demonstration of Real-World Long-Range Experiment.** This figure illustrates the main segment of the long-range real-world experiment, which integrates four smaller scenarios over a nearly 100-meter route from start to finish. The path includes a narrow corridor and an outdoor area with uneven terrain. The robot successfully reached the endpoint while consistently avoiding disturbance to social groups.

dynamic, diverse social behaviors within the robot’s immediate workspace. While the robot successfully completed all navigation tasks, several limitations emerged under dense crowd conditions. We analyze these challenges and offer insights to guide future improvements.

**LMM inference latency** Latency from LMM-based group estimation—typically 1.5 to 2 seconds using Qwen2.5-VL-32B—sometimes reduced planner responsiveness to rapidly changing social contexts. To balance reasoning capability and speed, we plan to distill this ability into a smaller model, Qwen2.5-VL-3B, via fine-tuning.

**Occlusions and partial observability** Egocentric perception suffered from partial occlusions, where individuals blocked one another from the camera’s view, leading to incomplete or inaccurate group detection. In future work, we aim to incorporate predicted bird’s-eye views to infer social relationships beyond egocentric observations.

**Complex and adaptive group behavior** Currently, our framework outputs binary group labels (0 or 1), and the planner conservatively avoids all identified social groups. While sufficient for constructing conservative social spaces, this approach lacks flexibility to adapt to nuanced social contexts (e.g., traversing a queue may be more acceptable than interrupting a group chat). Incorporating finer-grained social understanding is a key next step. One promising direction is to integrate LMM-generated costs—based on different social interaction types—into the planning process.

## VI. CONCLUSION

In this work, we present a novel approach that integrates the visual reasoning capabilities of LMMs with a social structure-aware perception and planning system for mobile robots operating in human-centered environments. Our method utilizes the common sense reasoning ability of LMM by applying the visual prompting to predict the social relationships among pedestrians, which allows our planning system to generate socially aware behavior. Through extensive real-world navigation experiments, our approach demonstrated superior performance and robustness compared to baseline methods, illustrating its potential for improving social awareness and interaction in autonomous robot navigation. Future work could explore extending this framework to more dense social interactions, combining compact state representation [57], and also distilling knowledge into smaller models for faster inference.

## REFERENCES

- [1] S. Thrun, M. Beetz, M. Bennewitz, W. Burgard, A. B. Cremers, F. Dellaert, D. Fox, D. Haehnel, C. Rosenberg, N. Roy *et al.*, “Probabilistic algorithms and the interactive museum tour-guide robot *minerva*,” *The international journal of robotics research*, vol. 19, no. 11, pp. 972–999, 2000.
- [2] B. Xia, H. Luan, Z. Zhao, X. Gao, P. Xie, A. Xiao, J. Wang, and M. Q.-H. Meng, “Collaborative trolley transportation system with autonomous nonholonomic robots,” in *2023 IEEE/RISJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 8046–8053.
- [3] A. Francis, C. Pérez-D’Arpino, C. Li, F. Xia, A. Alahi, R. Alami, A. Bera, A. Biswas, J. Biswas, R. Chandra, H.-T. L. Chiang, M. Everett, S. Ha, J. Hart, J. P. How, H. Karnan, T.-W. E. Lee, L. J. Manso, R. Mirksy, S. Pirk, P. T. Singamaneni, P. Stone, A. V. Taylor, P. Trautman, N. Tsoi, M. Vázquez, X. Xiao, P. Xu, N. Yokoyama, A. Toshev, and R. Martín-Martín. Principles and Guidelines for Evaluating Social Robot Navigation Algorithms.
- [4] G. Ferrer and A. Sanfeliu, “Proactive kinodynamic planning using the extended social force model and human motion prediction in urban environments,” in *2014 IEEE/RISJ International Conference on Intelligent Robots and Systems*, 2014, pp. 1730–1735.
- [5] Y. F. Chen, M. Everett, M. Liu, and J. P. How, “Socially aware motion planning with deep reinforcement learning,” in *2017 IEEE/RISJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 1343–1350.
- [6] L. Huber, J.-J. Slotine, and A. Billard, “Fast obstacle avoidance based on real-time sensing,” *IEEE Robotics and Automation Letters*, vol. 8, no. 3, pp. 1375–1382, 2022.
- [7] Y. Chen, Z. Xu, Z. Jian, G. Tang, L. Yang, A. Xiao, X. Wang, and B. Liang, “Quadruped guidance robot for the visually impaired: A comfort-based approach,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 12078–12084.
- [8] L. Huber, J.-J. Slotine, and A. Billard, “Avoidance of concave obstacles through rotation of nonlinear dynamics,” *IEEE Transactions on Robotics*, vol. 40, pp. 1983–2002, 2023.
- [9] 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.
- [10] P. T. Singamaneni, P. Bachiller-Burgos, L. J. Manso, A. Garrell, A. Sanfeliu, A. Spalanzani, and R. Alami, “A survey on socially aware robot navigation: Taxonomy and future challenges,” *The International Journal of Robotics Research*, vol. 43, no. 10, pp. 1533–1572, 2024.
- [11] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, “Social gan: Socially acceptable trajectories with generative adversarial networks,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2255–2264.
- [12] A. Wang and A. Steinfeld, “Group split and merge prediction with 3d convolutional networks,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1923–1930, 2020.
- [13] A. Taylor, D. M. Chan, and L. D. Riek, “Robot-centric perception of human groups,” *ACM Transactions on Human-Robot Interaction (THRI)*, vol. 9, no. 3, pp. 1–21, 2020.
- [14] A. Wang, C. Mavrogiannis, and A. Steinfeld, “Group-based motion prediction for navigation in crowded environments,” in *Conference on Robot Learning*, 2022, pp. 871–882.

## IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

- [15] S. Jahangard, M. Hayat, and H. Rezaatfoghi, "Real-time trajectory-based social group detection," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 1901–1908.
- [16] X. Lu, A. Faragasso, Y. Wang, A. Yamashita, and H. Asama, "Group-aware robot navigation in crowds using spatio-temporal graph attention network with deep reinforcement learning," *IEEE Robotics and Automation Letters*, 2025.
- [17] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.
- [18] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics Research: The 14th International Symposium ISRR*, 2011, pp. 3–19.
- [19] D. Panagou, "Motion planning and collision avoidance using navigation vector fields," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 2513–2518.
- [20] C. Rösmann, F. Hoffmann, and T. Bertram, "Timed-elastic-bands for time-optimal point-to-point nonlinear model predictive control," in *2015 European conference (ECC)*, 2015, pp. 3352–3357.
- [21] L. Tai, J. Zhang, M. Liu, and W. Burgard, "Socially compliant navigation through raw depth inputs with generative adversarial imitation learning," in *2018 IEEE international conference on robotics and automation (ICRA)*, 2018, pp. 1111–1117.
- [22] Z. Xie, P. Xin, and P. Dames, "Towards safe navigation through crowded dynamic environments," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 4934–4940.
- [23] N. Hirose, D. Shah, A. Sridhar, and S. Levine, "Sacson: Scalable autonomous control for social navigation," *IEEE Robotics and Automation Letters*, 2023.
- [24] D. M. Nguyen, M. Nazeri, A. Payandeh, A. Datar, and X. Xiao, "Toward human-like social robot navigation: A large-scale, multi-modal, social human navigation dataset," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 7442–7447.
- [25] 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.
- [26] 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.
- [27] 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.
- [28] Y. Kong, D. Song, J. Liang, D. Manocha, Z. Yao, and X. Xiao, "Autospacial: Visual-language reasoning for social robot navigation through efficient spatial reasoning learning," *arXiv preprint arXiv:2503.07557*, 2025.
- [29] P. Trautman, J. Ma, R. M. Murray, and A. Krause, "Robot navigation in dense human crowds: the case for cooperation," in *2013 IEEE International Conference on Robotics and Automation (ICRA)*, 2013, pp. 2153–2160.
- [30] F. Zhu, X. Wang, and N. Yu, "Crowd tracking with dynamic evolution of group structures," in *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part VI 13*, 2014, pp. 139–154.
- [31] Q. Wang, M. Chen, F. Nie, and X. Li, "Detecting coherent groups in crowd scenes by multiview clustering," *IEEE transactions on pattern analysis and machine intelligence*, vol. 42, no. 1, pp. 46–58, 2018.
- [32] Z. Fang, D. Hsu, and G. H. Lee, "Neutralized markov random field for interaction-aware stochastic human trajectory prediction," in *The Thirteenth International Conference on Learning Representations*.
- [33] N. V. Chappa, P. Nguyen, A. H. Nelson, H.-S. Seo, X. Li, P. D. Dobbs, and K. Luu, "Spartan: Self-supervised spatiotemporal transformers approach to group activity recognition," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 5158–5168.
- [34] D. Kim, Y. Song, M. Cho, and S. Kwak, "Towards more practical group activity detection: A new benchmark and model," in *European Conference on Computer Vision*, 2024, pp. 240–258.
- [35] F. Solera, S. Calderara, and R. Cucchiara, "Socially constrained structural learning for groups detection in crowd," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 5, pp. 995–1008, 2015.
- [36] X. Li, M. Chen, F. Nie, and Q. Wang, "A multiview-based parameter free framework for group detection," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 31, no. 1, 2017.
- [37] J. Li, R. Han, H. Yan, Z. Qian, W. Feng, and S. Wang, "Self-supervised social relation representation for human group detection," in *European Conference on Computer Vision*, 2022, pp. 142–159.
- [38] M. Luber and K. O. Arras, "Multi-hypothesis social grouping and tracking for mobile robots," in *Robotics: Science and Systems*, vol. 129, 2013.
- [39] A. Wang, C. Mavrogiannis, and A. Steinfeld, "Group-based motion prediction for navigation in crowded environments," in *Conference on Robot Learning*, 2022, pp. 871–882.
- [40] A. Taylor and L. D. Riek, "Regroup: A robot-centric group detection and tracking system," in *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2022, pp. 412–421.
- [41] A. Eirale, M. Leonetti, and M. Chiaberge, "Learning social cost functions for human-aware path planning," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 5364–5371.
- [42] H. B. Barua, T. H. Mg, P. Pramanick, and C. Sarkar, "Enabling social robots to perceive and join socially interacting groups using f-formation: A comprehensive overview," *ACM Transactions on Human-Robot Interaction*, vol. 13, no. 4, pp. 1–48, 2024.
- [43] D. Helbing and P. Molnar, "Social force model for pedestrian dynamics," *Physical review E*, vol. 51, no. 5, p. 4282, 1995.
- [44] A. Biswas, A. Wang, G. Silvera, A. Steinfeld, and H. Admoni, "Socnavbench: A grounded simulation testing framework for evaluating social navigation," *ACM Transactions on Human-Robot Interaction (THRI)*, vol. 11, no. 3, pp. 1–24, 2022.
- [45] K. Katyal, Y. Gao, J. Markowitz, S. Pohland, C. Rivera, I.-J. Wang, and C.-M. Huang, "Learning a group-aware policy for robot navigation," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 11 328–11 335.
- [46] M. Elnoor, K. Weerakoon, G. Seneviratne, J. Liang, V. Rajagopal, and D. Manocha, "Vi-lad: Vision-language attention distillation for socially-aware robot navigation in dynamic environments," *arXiv preprint arXiv:2503.09820*, 2025.
- [47] D. Jia, A. Hermans, and B. Leibe, "Dr-spaam: A spatial-attention and auto-regressive model for person detection in 2d range data," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 10 270–10 277.
- [48] G. Jocher, "YOLOv5 by Ultralytics," May 2020.
- [49] D. Harabor and A. Grastien, "Online graph pruning for pathfinding on grid maps," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 25, no. 1, 2011, pp. 1114–1119.
- [50] J. D. Gammell, T. D. Barfoot, and S. S. Srinivasa, "Batch informed trees (bit\*): Informed asymptotically optimal anytime search," *The International Journal of Robotics Research*, vol. 39, no. 5, pp. 543–567, 2020.
- [51] J. Zeng, Z. Li, and K. Sreenath, "Enhancing feasibility and safety of nonlinear model predictive control with discrete-time control barrier functions," in *2021 60th IEEE Conference on Decision and Control (CDC)*, 2021.
- [52] A. Xiao, H. Luan, Z. Zhao, Y. Hong, J. Zhao, W. Chen, J. Wang, and M. Q.-H. Meng, "Robotic autonomous trolley collection with progressive perception and nonlinear model predictive control," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022.
- [53] H. Karnan, A. Nair, X. Xiao, G. Warnell, S. Pirk, A. Toshev, J. Hart, J. Biswas, and P. Stone, "Socially compliant navigation dataset (scand): A large-scale dataset of demonstrations for social navigation," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 807–11 814, 2022.
- [54] J. W. Bae, J. Kim, J. Yun, C. Kang, J. Choi, C. Kim, J. Lee, J. Choi, and J. W. Choi, "Sit dataset: socially interactive pedestrian trajectory dataset for social navigation robots," *Advances in neural information processing systems*, vol. 36, pp. 24 552–24 563, 2023.
- [55] P. Long, T. Fan, X. Liao, W. Liu, H. Zhang, and J. Pan, "Towards optimally decentralized multi-robot collision avoidance via deep reinforcement learning," in *2018 IEEE international conference on robotics and automation (ICRA)*, 2018, pp. 6252–6259.
- [56] L. Kästner, R. Carstens, H. Zeng, J. Kmiecik, T. Bhuiyan, N. Khorasandhi, V. Shcherbyna, and J. Lambrecht, "Arena-rosnav 2.0: A development and benchmarking platform for robot navigation in highly dynamic environments," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 11 257–11 264.
- [57] S. Chen, A. Xiao, and D. Hsu, "Llm-state: Expandable state representation for long-horizon task planning in the open world," *arXiv preprint arXiv:2311.17406*, 2023.