

HiCrowd: Hierarchical Crowd Flow Alignment for Dense Human Environments

Yufei Zhu¹, Shih-Min Yang¹, Martin Magnusson¹, Allan Wang²

Abstract—Navigating through dense human crowds remains a significant challenge for mobile robots. A key issue is the freezing robot problem, where the robot struggles to find safe motions and becomes stuck within the crowd. To address this, we propose HiCrowd, a hierarchical framework that integrates reinforcement learning (RL) with model predictive control (MPC). HiCrowd leverages surrounding pedestrian motion as guidance, enabling the robot to align with compatible crowd flows. A high-level RL policy generates a follow point to align the robot with a suitable pedestrian group, while a low-level MPC safely tracks this guidance with short horizon planning. The method combines long-term crowd aware decision making with safe short-term execution. We evaluate HiCrowd against reactive and learning-based baselines in offline setting (replaying recorded human trajectories) and online setting (human trajectories are updated to react to the robot in simulation). Experiments on a real-world dataset and a synthetic crowd dataset show that our method outperforms in navigation efficiency and safety, while reducing freezing behaviors. We further validate through real-world deployment in a public museum and Expo 2025 Osaka, where it navigates dense pedestrian flows without retraining, demonstrating robust and socially aware behavior. Our results suggest that leveraging human motion as guidance, rather than treating humans solely as dynamic obstacles, provides a powerful principle for safe and efficient robot navigation in crowds. Project code and demos are available at <https://github.com/test-bai-cpu/HiCrowd>.

I. INTRODUCTION

With rising cases of robots navigating through dense human environments, mobile robots must balance efficiency, safety, and social compliance [1]. A key challenge in this context is the *freezing robot problem* (FRP) [2], where the robot struggles to find safe motions and becomes stuck within the crowd. Prior research has recognized that modeling human-human and human-robot interactions is a critical aspect in addressing the FRP [3, 4].

Early model-based reactive planners [5, 6] attempted reciprocal collision avoidance or force-based models, but failed to capture intricate interactions in dense human environments, resulting in FRP for the robot. Recent reinforcement learning-based (RL) approaches [7, 8] learn socially compliant behaviors by capturing interactions implicitly from simulation data, but mapping crowd states to low-level actions can be difficult to regulate safety [9].

We mitigate this issue by leveraging *crowd flow*. When moving in dense crowds, humans can make progress by fol-



Fig. 1: Example of HiCrowd navigating in a real-world dense crowd environment. The mobile robot platform is marked with an orange circle. Pedestrians moving in a similar direction are shown in cyan, while those walking in the opposite direction are shown in magenta. At $t = 4$, the robot aligns with the pedestrian flow in front and follows it to the right, without yet observing the opposing group that appears later. When the oncoming pedestrians (magenta) appear at $t = 7$, the robot, already aligned with the group ahead, smoothly follows the flow to the right and avoids collision. By $t = 12$ and $t = 17$, the robot maintains alignment with the surrounding flow and avoids another group, achieving safe and efficient progress.

lowing others moving in compatible directions [10] instead of being stalled by repeated individual avoidance maneuvers. This observation motivates our hypothesis that in dense human environments, following a moving crowd can be an effective navigation strategy, enabling robots to maintain consistent progress toward their goals in dense crowds.

In this work, we propose HiCrowd, a hierarchical architecture that integrates reinforcement learning (RL) with model predictive control (MPC) to enable efficient and safe navigation in dense human environments, as shown in Fig. 1. HiCrowd’s high-level RL policy predicts a socially aware *follow point* that guides the robot to align with a suitable pedestrian group (i.e., the local crowd flow). The low-level MPC then optimizes a short-horizon trajectory to track this guidance while ensuring dynamic feasibility and safety with respect to nearby humans. This hierarchical design separates long-term decision making (group following and goal progress) from short-term motion control (collision avoidance), enabling consistent progress toward the goal in dense crowds while ensuring reliable safety in highly dynamic environments.

We evaluate HiCrowd against both classical and learning-based baselines, including the reactive planner ORCA [6],

¹Robot Navigation and Perception Lab, AASS Research Center, Örebro University, Sweden yufei.zhu@oru.se

²Miraikan – The National Museum of Emerging Science and Innovation, Japan. This work received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 101070596 (euRobin).

and the RL-based methods SARL [8] and CrowdAttn [11]. We include an MPC-only ablation to isolate the contribution of the crowd-following module. Experiments are conducted in two settings: (1) an *offline* setting, where humans replay recorded motions without reacting to the robot, and (2) an *online* setting, where humans (controlled via ORCA) react to the robot in simulation. We use the ETH-UCY dataset [12, 13] and a synthetic crowd dataset to test navigation performance in dense human environments.

Experimental results show that HiCrowd outperforms baseline methods, achieving higher success rates and more efficient navigation. It also maintains safe social distances and significantly reduces freezing in dense crowds. The qualitative analysis highlights that HiCrowd consistently aligns with crowd flows and reaches goals with minimal delay. In contrast, baseline methods tend to become stuck in dense opposing flows, leading to abrupt reactive actions and frequent freezing. Finally, we deployed HiCrowd on a mobile robot in two real-world crowded environments: a public museum and Expo 2025 Osaka. Without any retraining, the robot successfully navigated dense pedestrian flows, robustly demonstrating safe and socially aware behavior.

In summary, we make the following contributions:

- Demonstrating that leveraging human motion as guidance, rather than solely as obstacles, provides a powerful principle for safe and efficient crowd navigation.
- Introducing a hierarchical RL-MPC framework where the RL agent is trained with a crowd-following reward that encourages alignment with compatible crowd flow.
- Showing that the crowd-following reward accelerates learning, improves performance, and can be effectively deployed in the real world.

II. RELATED WORK

Robot navigation in dynamic crowds has been studied for decades. Classical rule-based approaches such as social force [5] and Optimal Reciprocal Collision Avoidance (ORCA) [6] model interactions through attractive and repulsive forces or reciprocal velocity obstacles. While effective in structured scenarios, their performance degrades in complex crowd settings, especially when pedestrians behave unpredictably, resulting in the freezing robot problem [2].

Instead of explicitly modeling agent behaviors, recent approaches have used reinforcement learning (RL) to implicitly capture complex inter-agent interactions. Deep RL methods [7] have been introduced for decentralized multi-agent collision avoidance, where each agent independently learns safe behaviors without explicit communication. Building on this line of work, SARL [8], an attention-based Deep RL framework, has been proposed to enable robots to reason about interactions with surrounding pedestrians in a socially compliant manner. More recently, Liu et al. [11, 14] extended this direction by using a spatio-temporal attention graph to encode interactions, while Xie and Dames [15] combined Deep RL with velocity obstacles to navigate crowded dynamic scenes. These works demonstrate the potential of Deep RL to learn socially aware navigation strategies directly from

simulated data, though the ability to generalize to dense and diverse crowd environments remains an open challenge.

Beyond RL policies, hierarchical reinforcement learning (HRL) has been widely explored in robotics as a way to improve learning efficiency. Nachum et al. [16] propose a HRL method that generates subgoals to enable efficient robot control for navigation tasks. Nasiriany et al. [17] replace the low-level agent in HRL with predefined parametric primitives to separate primitives from parameter estimation. Yang et al. [18] apply HRL to combine a sequence of learned parameterized primitives for robotic manipulation. These works demonstrate how hierarchical structures can simplify decision making by decomposing long-horizon problems into subgoals or primitives, a principle that is also relevant for navigation in dynamic crowds.

Early work demonstrated the feasibility of integrating RL with MPC for control tasks [9]. Wang and Ba [19] proposed using policy networks to guide sampling-based MPC by generating candidate actions for high-dimensional locomotion control tasks. Brito et al. [20] introduced an RL subgoal planner to guide an MPC local planner, rewarding goal-reaching while penalizing collisions. Most recently, Han et al. [21] proposed DR-MPC, integrating MPC path tracking with model-free Deep RL to address the challenges of large data requirements and unsafe initial behaviors.

Moreover, human flow is important for robots to conform to surrounding motion patterns [22]. Previous approaches have leveraged pedestrian group formations, a simple form of crowd flow. Wang et al. [23] and Katyal et al. [24] considered groups in their cost terms in MPC- and RL-based frameworks respectively. However, these methods optimized goal reaching while avoiding group intrusions, without explicitly considering the integration of underlying crowd flow. Swaminathan et al. [25] incorporated learned site-specific motion patterns into an RRT* planner, but these patterns are computed offline, making real-time adaptation difficult. [26].

Recent work leverages online observations of pedestrians to guide navigation. Liao et al. [10] propose a people-as-planners framework that selects a human leader via rule-based evaluation. But relying on a single leader and rule-based selection limits generalization to diverse real-world scenarios. In this work, we bridge this gap by focusing on *crowd flow alignment* and learning a socially aware follow point with a high-level RL policy, enabling safe and efficient navigation in dense human environments.

III. METHOD

We propose a hierarchical crowd navigation approach that integrates reinforcement learning (RL) with model predictive control (MPC) for robot navigation in dense human environments. An overview of the framework is shown in Fig. 2. At the high level, the RL policy produces a *follow point* to guide the robot to align with a suitable group in the current crowd. At the low level, the MPC controller uses this follow point as a reference while optimizing a short-horizon trajectory that ensures dynamic feasibility and safety with respect to nearby humans. The framework separates long-term decision




Fig. 2: Method overview. The robot observes nearby humans within its sensing radius r_{obs} along with its own state s and goal g . A high-level RL policy outputs a follow point (f_x, f_y) that guides the robot to align with a suitable crowd flow. The low-level MPC controller then generates a control action u to track this follow point while ensuring collision avoidance and dynamic feasibility. The RL policy is updated based on rewards that combine goal reaching, progress, and crowd following.

making from short-term motion control, enabling the robot to progress consistently toward the goal even in dense crowds, mitigating the freezing robot problem.

A. Problem Definition

We consider the problem of robot navigation in dense human environments. The robot navigates in a 2D workspace $\mathcal{W} \subset \mathbb{R}^2$ among human agents.

At time t , the robot state is represented by its position and orientation, $s_t = (p_{x,t}^r, p_{y,t}^r, \theta_t)$, and the robot is tasked with reaching a goal $g = (g_x, g_y)$. The robot does not know the humans' goals, but can observe them within a sensing range r_{obs} . Within r_{obs} , the robot observes n humans, each represented by position and velocity, $h_t^i = (p_{x,t}^i, p_{y,t}^i, v_{x,t}^i, v_{y,t}^i)$, $i = 1, \dots, n$ and $H_t = (h_t^1, h_t^2, \dots, h_t^n)$ denotes the set of human states at time t . Humans may form groups when spatially close and moving in similar directions.

At each step, the robot executes a control input $u_t = (v_t, \omega_t)$ where v_t is the linear velocity and ω_t is the angular velocity. The robot state is updated according to a unicycle model with time step Δt :

$$\begin{aligned} x_{t+1} &= x_t + v_t \cos(\theta_t) \Delta t, \\ y_{t+1} &= y_t + v_t \sin(\theta_t) \Delta t, \\ \theta_{t+1} &= \theta_t + \omega_t \Delta t. \end{aligned} \quad (1)$$

For $\omega_t \neq 0$, the exact displacement can be expressed as an arc of radius v_t/ω_t .

B. Reinforcement Learning Module

The reinforcement learning (RL) policy acts as a high-level decision maker, responsible for generating a *follow point* that guides the robot through the crowd. Unlike RL-based motion planners that directly output control action [8, 11], our RL agent predicts a point in robot-centric coordinates. This follow point serves as a target reference for the low-level MPC controller, which will be described later in Sec. III-C.

The RL policy observes the robot state s_t , the robot goal g , and the set of nearby human states H_t within the observation range r_{obs} . The action of the RL policy is a 2D follow point $f_t = (f_{x,t}, f_{y,t})$ in the robot's local frame. Formally as, $f_t \sim \pi(\cdot | s_t, H_t, g)$. The optimal policy π^* maximizes the expected return: $\pi^* = \arg \max_{\pi} \mathbb{E}_{\pi} [\sum_{k=0}^{\infty} \gamma^k r_{t+k}]$, where r_{t+k} is the reward at time $t+k$, and $\gamma \in [0, 1)$ is the

discount factor. At each time step t , the reward is defined as the weighted sum of three components:

$$r_t = r_t^{\text{goal}} + \lambda_d r_t^{\text{dist}} + \lambda_f r_t^{\text{follow}}, \quad (2)$$

Goal reaching sparse reward (r_t^{goal}). The agent receives positive reward when it reaches goal location g .

Goal reaching dense reward (r_t^{dist}). The agent is encouraged to be closer to the goal, defined as $r_t^{\text{dist}} = \|\mathbf{p}_{t-1} - \mathbf{g}\| - \|\mathbf{p}_t - \mathbf{g}\|$, where $\mathbf{p}_t = (p_{x,t}^r, p_{y,t}^r)$, i.e. the robot position at time t . This term is weighted by λ_d .

Crowd following reward (r_t^{follow}). To compute r_t^{follow} , we first cluster the n observed humans into groups based on spatial proximity and motion similarity. We apply the DBSCAN algorithm [27] to human states $\{h_t^i\}_{i=1}^n$ with thresholds $(\varepsilon_p, \varepsilon_{\theta}, \varepsilon_v)$ on position, orientation, and speed, respectively. DBSCAN assigns each human i a group label ℓ_t^i . Each group is then defined as $G_j = \{i | \ell_t^i = j\}$, where a group may include one or more people. For each group G_j , we compute scores that measure how well the robot's current motion matches the group's average motion: $s_j^{\text{speed}} = \exp(-(\frac{\Delta v_j}{\sigma_v})^2)$, $s_j^{\text{angle}} = \exp(-(\frac{\Delta \theta_j}{\sigma_{\theta}})^2)$, $s_j^{\text{dist}} = \exp(-d_j^2)$, where Δv_j and $\Delta \theta_j$ are the differences between the robot's speed and heading and the average speed and heading of group G_j , and d_j is the minimum distance from the robot to any member of group G_j . The final crowd following reward is the maximum score across all groups:

$$r_t^{\text{follow}} = \max_j (s_j^{\text{speed}} \cdot s_j^{\text{angle}} \cdot s_j^{\text{dist}}). \quad (3)$$

Using the maximum score across groups, the robot aligns with the most suitable group, typically the one moving in a similar direction and toward a compatible destination, while retaining flexibility to switch groups as conditions change. This term is weighted by λ_f .

During training, the RL agent generates a follow point, which is passed to the MPC-based low-level controller. The MPC computes a safe robot control input u_t to execute, and the environment transitions to a new state as in Eq. (1). The reward is computed based on the outcome of the MPC, and the RL policy is updated accordingly. This hierarchical design separates strategic decision-making (selecting a suitable group to follow) from low-level motion control (ensuring dynamic feasibility and collision avoidance).

C. Model Predictive Control Module

At time t , the MPC module takes as input the robot state s_t , the human states H_t , and the follow point f_t . It then optimizes a short-horizon trajectory by minimizing a cost function that balances progress toward the RL-provided follow point and maintaining safety with respect to nearby humans. The total cost is given by $J = J_{\text{follow}} + J_{\text{safe}}$.

Follow cost (J_{follow}). The follow cost encourages the robot to move towards the follow point, defined as the minimum Euclidean distance between the predicted trajectory and the RL-provided follow point f_t :

$$J_{\text{follow}} = \min_{\tau \in [0, H-1]} \|\mathbf{p}_{t+\tau}^r - \mathbf{f}_t\|, \quad (4)$$

where $\mathbf{p}_{t+\tau}^r$ is the predicted robot position at step $t + \tau$ over a planning horizon H .

Safety cost (J_{safe}). The safety cost penalizes proximity to humans over the planning horizon. For each human i , we predict the future position using a constant velocity model (CVM), and compute the distance $d_{t+\tau}^i$ between the predicted robot position and the predicted human position at horizon step τ . The safety cost is defined as

$$J_{\text{safe}} = \sum_{\tau=0}^{H-1} \beta^\tau \sum_{i=1}^n \exp(d_c - d_{t+\tau}^i), \quad (5)$$

where d_c is a collision radius, $\beta \in (0, 1)$ is a discount factor that decrease the weights of future collision risks.

The MPC selects the control sequence $\mathbf{u}_{t:t+H-1}$ that minimizes J , and applies the first control input \mathbf{u}_t to update the robot dynamics as in Eq. (1).

IV. EXPERIMENTS

A. Datasets

The evaluation is carried out in real-world datasets and a synthetic dataset.

ETH-UCY: The ETH-UCY dataset [12, 13] contains 5 subsets: ETH, HOTEL, UNIV, ZARA1 and ZARA2, with various motion patterns. Similar to [23], from these recordings, we construct navigation episodes by segmenting the trajectories into blocks that involve dense human interactions. A segment contains at least five pedestrians present in the scene. Across all scenarios, the number of pedestrians ranges from 5 to 22, with an average of 7.3 pedestrians per scene.

Synthetic: To evaluate navigation performance under even denser conditions, we construct a synthetic dataset simulating pedestrian flows. Each episode consists of two opposing flows of pedestrians. Groups are spawned at randomized intervals around 5s from opposite sides of the scene, each containing 2 to 5 pedestrians. Within a group, each pedestrian is assigned an initial velocity sampled around a mean speed of 0.9 m s^{-1} , with a random noise of $\pm 0.05 \text{ m s}^{-1}$, and moves with constant velocity.

We construct navigation episodes by applying the same segmentation procedure as used for ETH-UCY. Across all episodes, the number of pedestrians ranges from 6 to 63, with an average of 19.8 pedestrians per scene.

For both datasets, we generate multiple episodes per segment by uniformly sampling robot start and goal positions from start and goal regions placed on opposite sides of the dominant pedestrian flow to encourage interactions. We sample 5,000 training episodes from ETH-UCY and 3,000 from the synthetic dataset, in total across all segments. For evaluation, we generate 500 test episodes per dataset, and use the same test episodes for evaluating all methods.

B. Experimental Setup

We evaluate in two settings that differ in whether pedestrians react to the robot.

Offline: Pedestrians replay recorded trajectories in the dataset and do not react to the robot. This setting evaluates whether a crowd navigation planner can reach the goal without causing deviations in pedestrian behavior. It also reflects real deployments with small robots that nearby pedestrians may overlook and fail to react in time.

Online: Pedestrian are initialized with the same start and goal locations as in the corresponding dataset segment. Their subsequent motion is simulated by ORCA [6], and pedestrians react to the robot and to each other, with maximum speed set to 1.75 m s^{-1} .

C. Baselines

We compare with the following baseline methods.

MPC: The MPC baseline uses the same setup as the MPC component (Sec. III-C) of HiCrowd, but instead of taking an RL-provided follow point as the reference, it directly uses the robot’s goal location. This baseline evaluates the effect of the RL-guided crowd-following module in HiCrowd.

ORCA: Optimal Reciprocal Collision Avoidance (ORCA) [6] serves as a reactive-based baseline, generating collision-free movements through reciprocal collision avoidance. ORCA is evaluated as a baseline in both offline and online settings, and in online settings it also serves as the controller for pedestrians.

SARL: The Socially Attentive Reinforcement Learning (SARL) [8] learns a policy that maps crowd states directly to robot actions. This baseline evaluates the performance of the hierarchical RL-MPC design and the crowd-following strategy over a flat RL approach.

CrowdAttn: Crowd Navigation with Attention-Based Interaction Graph method [11] augments RL-based navigation with a spatio-temporal attention graph. CrowdAttn serves as a recent state-of-the-art baseline.

D. Evaluation Metrics

The evaluation metrics include the success rate (**SR**), with failures arising from either collisions (**CR**) or timeouts (**TR**). In the case of a collision, the episode is immediately terminated as a failure. A timeout occurs when the robot fails to reach the goal within a time limit set to three times the duration of reaching the goal in a straight line in the dataset.

We also report the average navigation time in seconds (**NT**), the average path length for non-collision episodes in meters (**PL**), and the average minimum pedestrian distance in meters (**MP**). In addition, we measure the freezing frequency (**FF**), defined as the percentage of time steps in which the robot’s velocity reaches zero (following [4]), indicating how often the robot becomes stuck among human agents.

E. Implementation Details

The RL agent is trained for 100K transitions using the Soft Actor-Critic (SAC) algorithm [28]. We use a discount factor of $\gamma = 0.99$ and a batch size of 256. The actor and critic are trained with learning rates of 3×10^{-4} and 1×10^{-3} , respectively. The entropy coefficient is trained with a learning rate of 1×10^{-3} . We use Adam for all optimizations. Observation radius r_{obs} is set to 5 m. The MPC module uses

Method	SR \uparrow	CR \downarrow	TR \downarrow	NT \downarrow	PL \downarrow	MP \uparrow	FF \downarrow
HiCrowd	1.00	0.00	0.00	19.10	19.02	1.50	0.001
MPC	0.98	0.00	0.02	28.29	27.78	2.01	0.010
CrowdAttn	1.00	0.00	0.00	26.49	25.91	1.12	0.008
ORCA	0.77	0.00	0.23	43.43	39.97	1.36	0.013
SARL	0.78	0.00	0.22	33.25	25.90	1.19	0.051

TABLE I: Navigation results in the **online** setting on the ETH-UCY dataset, where pedestrians react to the robot using ORCA simulation. Metrics include success rate (SR \uparrow), collision rate (CR \downarrow), timeout rate (TR \downarrow), navigation time (NT \downarrow), path length (PL \downarrow), minimum pedestrian distance (MP \uparrow), and freezing frequency (FF \downarrow).

Method	SR \uparrow	CR \downarrow	TR \downarrow	NT \downarrow	PL \downarrow	MP \uparrow	FF \downarrow
HiCrowd	1.00	0.00	0.00	30.88	30.85	1.35	0.000
MPC	1.00	0.00	0.00	36.83	36.35	1.60	0.008
CrowdAttn	0.96	0.00	0.04	45.45	43.22	1.45	0.017
ORCA	0.97	0.00	0.03	40.89	39.85	1.61	0.002
SARL	0.85	0.00	0.15	58.35	42.59	1.29	0.003

TABLE II: Navigation results in the **online** setting on the Synthetic dataset.

a planning horizon $H = 10$ and discount factor $\beta = 0.9$. The collision radius d_c is 0.5 m and is used as the threshold for collision detection. The ratio between high-level RL action and low-level MPC action is 10. For DBSCAN grouping, $\varepsilon_\theta = \pi/6$, $\varepsilon_v = 1.0$, $\varepsilon_p = 2$. For computing crowd following reward r_t^{follow} , $\sigma_v = 0.5$ and $\sigma_\theta = \pi/12$. Both goal reaching dense reward weight λ_d and crowd following reward weight λ_f are set to 1. In the ablation, we run experiments with $\lambda_f \in \{0, 1, 5\}$ to evaluate its effect on learning efficiency. The robot’s maximum linear speed is 1 m s^{-1} and angular speed is 3.14 rad s^{-1} .

For SARL training, the imitation learning part follows the same configurations as in [8]. The RL part is trained on both datasets in online and offline setup, for 5×10^5 transitions with a learning rate of 0.001. CrowdAttn is used with its pretrained model and the Gumbel Social Transformer (GST) for motion prediction, trained on ETH-UCY. For CrowdAttn and ORCA, we map holonomic action outputs to differential drive commands. The linear velocity is the velocity vector norm, and the angular velocity is the turning angle between the robot orientation and the holonomic action. This ensures comparability with our method and other baselines ¹.

V. RESULTS

A. Quantitative Evaluation

We evaluate our method HiCrowd against MPC, ORCA, SARL, and CrowdAttn in both *online* (pedestrians react via ORCA) and *offline* (pedestrians replay and do not react) settings on the ETH-UCY and Synthetic dataset.

Evaluation results in the online setting are shown for ETH-UCY in Table I and for Synthetic dataset in Table II. HiCrowd achieves the best overall performance, with full success and the lowest navigation time and path length. In the online setting, baseline methods achieve high success rates but require longer navigation times, exhibiting more frequent reactive adjustments in dense crowd environments.

¹We also trained CrowdAttn under a differential drive setting, but the results were suboptimal compared to those with the converted pre-trained holonomic model, so we present pre-trained model results.

Method	SR \uparrow	CR \downarrow	TR \downarrow	NT \downarrow	PL \downarrow	MP \uparrow	FF \downarrow
HiCrowd	0.88	0.12	0.00	21.82	21.71	1.55	0.001
MPC	0.76	0.09	0.15	40.65	36.53	2.11	0.056
CrowdAttn	0.55	0.44	0.01	26.82	26.25	0.94	0.005
ORCA	0.61	0.26	0.13	40.78	36.05	1.20	0.044
SARL	0.24	0.71	0.05	29.73	25.67	0.76	0.014

TABLE III: Navigation results in the **offline** setting on the ETH-UCY dataset, where pedestrians replay recorded trajectories and do not react to the robot.

Method	SR \uparrow	CR \downarrow	TR \downarrow	NT \downarrow	PL \downarrow	MP \uparrow	FF \downarrow
HiCrowd	0.99	0.01	0.00	35.21	35.18	2.30	0.000
MPC	0.91	0.09	0.00	36.08	34.93	1.92	0.023
CrowdAttn	0.85	0.12	0.03	46.40	42.96	1.62	0.019
ORCA	0.75	0.23	0.02	37.53	35.69	1.60	0.002
SARL	0.45	0.50	0.05	57.57	42.68	1.33	0.040

TABLE IV: Navigation results in the **offline** setting on the Synthetic dataset.

Evaluation results in the offline setting are shown for ETH-UCY in Table III and for Synthetic dataset in Table IV. When pedestrians do not react to the robot, RL-based baselines, including SARL and CrowdAttn, experience more collisions. The pure MPC baseline and ORCA require longer navigation times and make more frequent reactive adjustments. HiCrowd maintains a higher success rate, which shows the benefit of aligning with crowd flows rather than relying only on late reactive avoidance.

Notably, in the offline setting HiCrowd sometimes gets longer path lengths compared to baselines. These detours arise from following the pedestrian flow, which helps avoid potential future conflicts and collisions, leading to shorter navigation times and less freezing frequency. This illustrates the advantage of proactively aligning with social dynamics, where small path deviations ensure continuous and efficient progress toward the goal. Our method consistently achieves the lowest freezing frequency (FF) across all datasets and settings, showing that aligning with pedestrian flows enables continuous progress without getting stuck. In contrast, baselines can cause freezing situations, forcing the robot to stop.

On the other hand, although our method does not maximize the minimum pedestrian distance (MP) compared to baselines, it consistently maintains a social distance above 1.2 m [29] across all settings and datasets.

B. Qualitative Evaluation

Fig. 3 and Fig. 4 show example trajectories for our method and all baselines in the ETH-UCY dataset (offline setting) and the Synthetic dataset (online setting), respectively. To visualize and compare complete trajectories, we allow all methods to continue until the end, even after collisions. In Fig. 3, which shows a dense crowd episode, baselines take longer paths and more reactive actions, and collisions happen with ORCA, SARL, and CrowdAttn. Our method, HiCrowd, reaches the goal efficiently without any collision. In Fig. 4, two opposing pedestrian flows form a trap. Although the straight path to the goal is initially clear, it leads directly into the oncoming opposing flow, which lies beyond observation radius. In this situation, MPC and ORCA push into the opposing flow and freeze. CrowdAttn also takes the straight

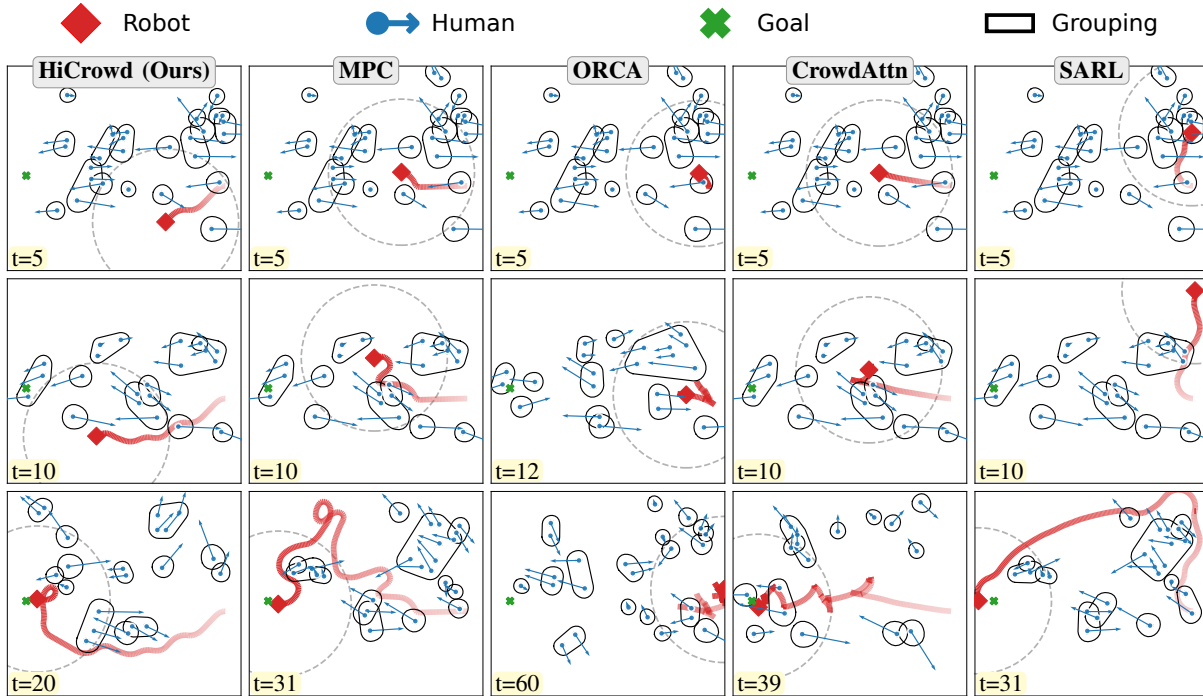


Fig. 3: Example episode from the ETH-UCY dataset in the **offline** setting (t in seconds). This case involves a dense crowd with both moving and static groups. Groups are visualized as convex hulls. Robot observation range is marked as a grey circle. Baseline methods take longer paths and more time to reach the goal, with multiple reactive avoidance actions. ORCA gets stuck near dense crowds. CrowdAttn freezes several times. And ORCA, SARL and CrowdAttn result in collisions. In contrast, HiCrowd successfully reaches the goal with a shorter navigation time.

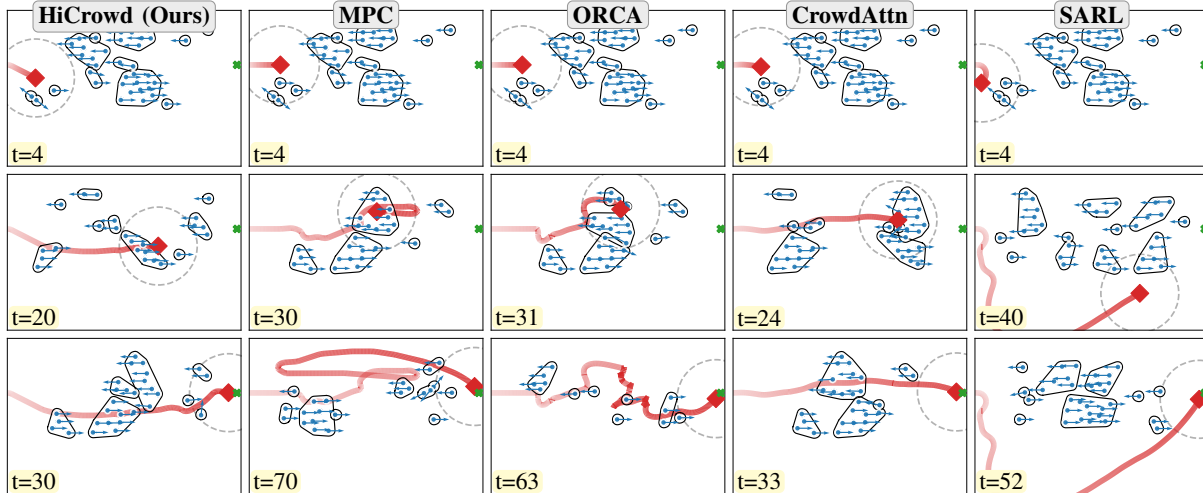


Fig. 4: Example episode from the Synthetic dataset in the **online** setting (t in seconds). The scenario contains two opposing pedestrian flows: one moving left to right and the other right to left. The robot starts on the left and aims for a goal on the right, located in the upper region where pedestrians move in the opposite direction. Taking the shortest direct path leads the robot into the dense opposing flow, which is initially outside the robot's observation radius. MPC and ORCA exhibit this behavior, becoming stuck in the oncoming flow and requiring much longer time to reach the goal. SARL avoids the dense crowd entirely, taking a long detour. CrowdAttn follows a straight path and later runs into the opposing flow, forcing the oncoming social groups to split apart for avoidance, resulting in multiple reactive actions and periods of freezing. In contrast, HiCrowd moves align with the pedestrian flow in the lower part of the scene, taking a detour but moving with the crowd, and reaches the goal in the shortest time.

path and later runs into the opposing flow, forcing the oncoming social groups to split apart for avoidance, resulting in multiple reactive actions and periods of freezing. In contrast, our method detours to align with the left-to-right crowd flow and arrives in the shortest navigation time without collision.

C. Ablation Study

To study the role of the crowd following reward r_t^{follow} , we perform an ablation with $\lambda_f \in \{0, 1, 5\}$ in Fig. 6.

When $\lambda_f > 0$, the agent receives explicit guidance to follow the crowd, which accelerates learning and achieves higher performance with fewer transition steps. The faster convergence observed with higher λ_f shows the usefulness of explicitly rewarding crowd following behavior, indicating that in dense environments, aligning with a moving crowd is an optimal navigation strategy. When $\lambda_f = 0$, the agent relies only on goal-reaching rewards. In this case, learning




Fig. 5: Real-world experiment setup and examples. (a) Our differential-drive robot and its sensor configuration. (b) Illustration of the footprint of our robot. It is a circle of 0.5 m radius that encompasses the robot and the predefined nominal human offset. (c) Qualitative examples. Robot movement is shown in yellow arrows. Different colors represent different groups of surrounding pedestrians. **Top row:** With the flow scenario. The robot merged into the pedestrian flow formed by the cyan and blue groups and naturally avoided the magenta group. **Middle row:** Against the flow scenario. The robot navigated successfully against the pedestrian flow, which consists of at least 3 groups. **Bottom row:** Crossing the flow scenario. The robot navigated through a crossing flow and reached to the goal successfully without collision.

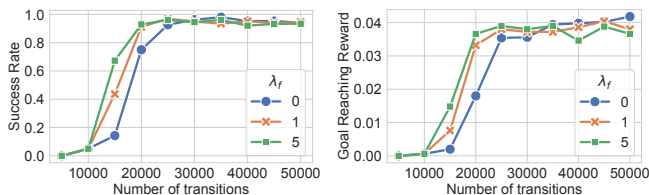


Fig. 6: Ablation on crowd following reward weight λ_f . Explicitly rewarding crowd following ($\lambda_f \in \{1, 5\}$) accelerates learning and improves early performance, showing that following a moving crowd is an effective navigation strategy in dense environment.

is slower but still converges eventually, as crowd following behavior implicitly emerges as part of the optimal strategy in dense environments.

In addition, the MPC-only baseline serves as an ablation that removes the RL module and plans directly toward the goal. Compared to HiCrowd, pure MPC provides safe short-horizon control but lacks long-term crowd reasoning, and can commit to locally clear paths that lead into opposing flows. The RL module complements MPC by generating a follow point that guides the robot to align with suitable crowd flows, effectively abstracting long-term guidance beyond the immediate sensing horizon. This combination improves efficiency and reduces freezing, enabling efficient and safe navigation in dense crowds.

D. Real world experiment

To validate the effectiveness of HiCrowd in the real world, we deployed our model on the robot shown in Fig. 5-a. The robot is a custom platform built from scratch, referencing the open-source platform [30]. We mounted a front-facing Intel RealSense D435 stereo camera and used the MMDetection toolbox [31] combined with depth images to obtain trajectories of the pedestrians in world coordinates. We also mounted a Velodyne VLP-16 light detection and ranging (LiDAR) sensor to localize the robot and safeguard against collisions. Due to the camera placement, the robot could not see the pedestrians to the sides or the rear. The robot’s computer has an Intel Core i7-9750H CPU and a NVIDIA GeForce RTX 2070 GPU. The robot’s speed is limited to 1 m s^{-1} linear and $\frac{\pi}{4} \text{ rad s}^{-1}$.

We deployed the robot in a crowded real-world science museum during a holiday (Fig. 5-c). In line with our open-space assumption, we selected a wide area with few static obstacles. Our deployment was approved by our institution’s ethics committee. To ensure the safety of the unsuspecting visitors, we attached a touch sensor to the top handle part of the robot, and the robot would only navigate autonomously when the handle was held². We directly deployed HiCrowd trained in the simulator, and integrated it as a controller and planner plugin to the robot’s ROS2-based software stack. The robot’s footprint is a circle that encompasses both the robot and the nominal location of the human supervisor (Fig. 5-b).

We conducted 20 trials and tasked the robot to cross the selected museum area repeatedly. During the trials, the robot often needed to navigate around 10 people, and we did not observe any significant issues such as collisions or the freezing robot problem. In terms of computation, the RL policy estimated a follow point in 0.0047 ± 0.0053 seconds, and the time to produce an MPC action was 0.15 ± 0.073 seconds. Example robot behaviors are shown in Fig. 5-c. In the first example, the robot merged into the pedestrian flow formed by the cyan and blue groups and naturally avoided the magenta group. In the second example, the robot navigated successfully against the pedestrian flow characterized by at least 3 groups. In the third example, the robot navigated through a crossing flow defined by multiple groups. Please refer to our videos for more details.

VI. CONCLUSION

We introduced HiCrowd, a hierarchical navigation framework that integrates RL and MPC to address the challenges of robot navigation in dense human environments. Rather than treating humans as dynamic obstacles [15], HiCrowd leverages surrounding pedestrian motion as guidance. A high-level RL policy generates a socially aware follow point that steers the robot toward a compatible crowd flow, while a low-level MPC controller tracks this guidance with safe, short-horizon planning. By encouraging alignment with

²Experiments would not be permitted without this feature.

nearby groups through a crowd-following reward, the framework improves both learning efficiency and navigation performance. Extensive experiments across synthetic and real-world datasets demonstrate that HiCrowd outperforms classical and learning-based baselines, achieving higher success rates and more efficient navigation, also reducing freezing. The benefits are most pronounced in dense scenarios with structured human flows, where baseline methods can freeze or collide when encountering opposing crowds.

Finally, we deployed HiCrowd on a mobile robot and evaluated in two crowded real-world environments, a public museum and Expo 2025 Osaka, without any retraining. The robot consistently demonstrated socially compliant behaviors in scenarios of aligning with and crossing pedestrian flows, showing the real-world applicability of our approach.

The proposed method has limitations. First, the ratio between high-level and low-level actions is fixed, which is, the robot executes a number of MPC steps per follow point. This may lead to delayed reactions when crowd conditions change rapidly, or unstable behavior when follow points change too frequently. Future work will explore dynamically adapting the ratio based on flow variations or environmental uncertainty to improve responsiveness and stability. Static obstacles such as walls will also be explored. Lastly, the current framework treats groups primarily through their motion features without modeling their spatial extent. Incorporating group geometry could improve decision-making, particularly when navigating around large or irregularly shaped groups.

REFERENCES

- [1] R. Triebel et al. "Spencer: A socially aware service robot for passenger guidance and help in busy airports". In: *Field and service robotics*. Springer. 2016, pp. 607–622.
- [2] P. Trautman and A. Krause. "Unfreezing the robot: Navigation in dense, interacting crowds". In: *Proc. of the Int. Conf. on Intell. Robots and Syst. (IROS)*. IEEE. 2010, pp. 797–803.
- [3] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh. "Core challenges of social robot navigation: A survey". In: *ACM Transactions on Human-Robot Interaction* 12.3 (2023), pp. 1–39.
- [4] S. Samavi, J. R. Han, F. Shkurti, and A. P. Schoellig. "SICNav: Safe and Interactive Crowd Navigation using Model Predictive Control and Bilevel Optimization". In: *Trans. on Robotics and Automation (TRO)* 41 (2024), pp. 80–818.
- [5] D. Helbing and P. Molnar. "Social force model for pedestrian dynamics". In: *Physical review E* 51.5 (1995), p. 4282.
- [6] J. van den Berg, S. J. Guy, M. Lin, and D. Manocha. "Reciprocal n-Body Collision Avoidance". In: *Robotics Research*. 2011, pp. 3–19.
- [7] Y. F. Chen, M. Liu, M. Everett, and J. P. How. "Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning". In: *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*. IEEE. 2017, pp. 285–292.
- [8] C. Chen, Y. Liu, S. Kreiss, and A. Alahi. "Crowd-Robot Interaction: Crowd-Aware Robot Navigation With Attention-Based Deep Reinforcement Learning". In: *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*. 2019, pp. 6015–6022.
- [9] D. Ernst, M. Glavic, F. Capitanescu, and L. Wehenkel. "Reinforcement Learning Versus Model Predictive Control: A Comparison on a Power System Problem". In: *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)* 39.2 (2009), pp. 517–529.
- [10] Y. Liao, X. Xu, R. Bai, Y. Yang, M. Cao, S. Yuan, and L. Xie. "Following is All You Need: Robot Crowd Navigation Using People as Planners". In: *Robotics and Automation Letters* 10.10 (2025), pp. 9814–9821.
- [11] S. Liu, P. Chang, Z. Huang, N. Chakraborty, K. Hong, W. Liang, D. L. McPherson, J. Geng, and K. Driggs-Campbell. "Intention Aware Robot Crowd Navigation with Attention-Based Interaction Graph". In: *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*. 2023, pp. 12015–12021.
- [12] A. Lerner, Y. Chrysanthou, and D. Lischinski. "Crowds by example". In: *Computer Graphics Forum*. Vol. 26. 3. Wiley Online Library. 2007, pp. 655–664.
- [13] S. Pellegrini, A. Ess, and L. van Gool. "Improving data association by joint modeling of pedestrian trajectories and groupings". In: *Proc. of the Europ. Conf. on Comp. Vision (ECCV)*. Springer. 2010, pp. 452–465.
- [14] S. Liu, H. Xia, F. C. Pouria, K. Hong, N. Chakraborty, Z. Hu, J. Biswas, and K. Driggs-Campbell. "HEIGHT: Heterogeneous Interaction Graph Transformer for Robot Navigation in Crowded and Constrained Environments". In: *Trans. on Automation Science and Engineering* 23 (2026), pp. 1211–1230.
- [15] Z. Xie and P. Dames. "DRL-VO: Learning to Navigate Through Crowded Dynamic Scenes Using Velocity Obstacles". In: *Trans. on Robotics and Automation (TRO)* 39.4 (2023), pp. 2700–2719.
- [16] O. Nachum, S. S. Gu, H. Lee, and S. Levine. "Data-efficient hierarchical reinforcement learning". In: *Advances in Neural Inf. Proc. Syst. (NeurIPS)* 31 (2018).
- [17] S. Nasiriany, H. Liu, and Y. Zhu. "Augmenting reinforcement learning with behavior primitives for diverse manipulation tasks". In: *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*. IEEE. 2022, pp. 7477–7484.
- [18] S.-M. Yang, M. Magnusson, J. A. Stork, and T. Stoyanov. "Learning extrinsic dexterity with parameterized manipulation primitives". In: *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*. IEEE. 2024, pp. 5404–5410.
- [19] T. Wang and J. Ba. "Exploring Model-Based Planning with Policy Networks". In: *Int. Conf. on Learning Representations (ICLR)*. 2020.
- [20] B. Brito, M. Everett, J. P. How, and J. Alonso-Mora. "Where to go Next: Learning a Subgoal Recommendation Policy for Navigation in Dynamic Environments". In: *Robotics and Automation Letters* 6.3 (2021), pp. 4616–4623.
- [21] J. R. Han, H. Thomas, J. Zhang, N. Rhinehart, and T. D. Barfoot. "DR-MPC: Deep Residual Model Predictive Control for Real-World Social Navigation". In: *Robotics and Automation Letters* 10.4 (2025), pp. 4029–4036.
- [22] Y. Zhu, S.-M. Yang, A. Rudenko, T. P. Kucner, A. J. Lilienthal, and M. Magnusson. "NeMo-map: Neural Implicit Flow Fields for Spatio-Temporal Motion Mapping". In: *Int. Conf. on Learning Representations (ICLR)*. 2026.
- [23] A. Wang, C. Mavrogiannis, and A. Steinfeld. "Group-based motion prediction for navigation in crowded environments". In: *Conf. on Robot Learning*. PMLR. 2022, pp. 871–882.
- [24] K. Katyal, Y. Gao, J. Markowitz, S. Pohlman, C. Rivera, I.-J. Wang, and C.-M. Huang. "Learning a group-aware policy for robot navigation". In: *Proc. of the Int. Conf. on Intell. Robots and Syst. (IROS)*. IEEE. 2022, pp. 11328–11335.
- [25] C. S. Swaminathan, T. P. Kucner, M. Magnusson, L. Palmieri, and A. J. Lilienthal. "Down The CLiFF: Flow-aware Trajectory Planning under Motion Pattern Uncertainty". In: *Proc. of the Int. Conf. on Intell. Robots and Syst. (IROS)*. IEEE. 2018, pp. 7403–7409.
- [26] Y. Zhu, A. Rudenko, L. Palmieri, L. Heuer, A. J. Lilienthal, and M. Magnusson. "Fast online learning of cliff-maps in changing environments". In: *Proc. of the Int. Conf. on Robotics and Automation (ICRA)*. IEEE. 2025, pp. 10424–10431.
- [27] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. "A density-based algorithm for discovering clusters in large spatial databases with noise". In: *Proc. of the Int. Conf. on Knowledge Discovery and Data Mining*. 1996, pp. 226–231.
- [28] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor". In: *Int. Conf. on Machine Learning (ICML)*. 2018, pp. 1861–1870.
- [29] E. T. Hall. *The Hidden Dimension*. Anchor Books, 1966.
- [30] J. Guerreiro, D. Sato, S. Asakawa, H. Dong, K. M. Kitani, and C. Asakawa. "Cabot: Designing and evaluating an autonomous navigation robot for blind people". In: *Proc. of the ACM/SIGACCESS Int. Conf. on Comput. and Accessibility (CHI)*. 2019, pp. 68–82.
- [31] K. Chen et al. "MMDetection: Open MMLab Detection Toolbox and Benchmark". In: *arXiv preprint arXiv:1906.07155* (2019).