

Human-Inspired Multi-Agent Navigation using Knowledge Distillation

Pei Xu and Ioannis Karamouzas

<https://motion-lab.github.io/KDMA>

Abstract—Despite significant advancements in the field of multi-agent navigation, agents still lack the sophistication and intelligence that humans exhibit in multi-agent settings. In this paper, we propose a framework for learning a human-like general collision avoidance policy for agent-agent interactions in fully decentralized, multi-agent environments. Our approach uses knowledge distillation with reinforcement learning to shape the reward function based on expert policies extracted from human trajectory demonstrations through behavior cloning. We show that agents trained with our approach can take human-like trajectories in collision avoidance and goal-directed steering tasks not provided by the demonstrations, outperforming the experts as well as learning-based agents trained without knowledge distillation.

I. INTRODUCTION

The problem of decentralized multi-agent navigation has been extensively studied in a variety of fields including robotics, graphics, and traffic engineering. Existing state-of-the-art planners can provide formal guarantees about the collision-free motion of the agents allowing applicability on physical robots [1–3]. In addition, recent learning-based approaches are capable of end-to-end steering and human-aware robot navigation [4–6]. Despite recent advancements, though, existing agents cannot typically exhibit the level of sophistication and decision making that humans do in multi-agent settings. While the behavior of the agents can be improved through limited one-way communication [7, 8], and accounting for social norms [9–11], here we explore an alternative, data-driven approach that takes advantage of publicly available human trajectory datasets. Our approach starts with expert demonstrations obtained from such trajectories and learns human-like navigation policies in fully decentralized multi-agent environments using reinforcement learning.

The idea of imitating expert behaviors is not something new. For example, behavior cloning techniques have been successfully applied to autonomous driving [12], motion planning for autonomous ground robots [13], and distributed robot navigation [4] among others. However, pure imitation learning techniques are severely limited by the quality of the expert training dataset and cannot scale well to the multi-agent navigation domain due to its open-ended nature.

Combining expert demonstrations with reinforcement learning can address the problem of insufficient training samples, with typical approaches relying on bootstrapping reinforcement learning with supervised learning [14, 15],

inverse reinforcement learning [16], and generative adversarial imitation learning [17]. However such techniques are not directly applicable to the task of human-like collision avoidance learning since they typically assume reliable and representative expert demonstrations for the task in hand. Unfortunately, the experts (pedestrians) in human trajectory datasets are biased to some degree due to the fact that i) we only have access to a limited number of interaction data, which does not necessarily capture the behavior that the same expert will exhibit in a different setting; and ii) trajectory datasets cannot capture the non-deterministic nature of human decision making, i.e., there are more than one trajectories that a human expert can take in the same setting.

To address these issues, we propose to use *knowledge distillation* [18] to learn a human-like navigation policy through expert policies extracted from human demonstrations. Given the imperfect nature of the experts, we avoid directly optimizing over them by adding an extra term to the objective function, as typically proposed in the literature [19, 20]. Instead, we utilize the expert policies to shape the reward function during reinforcement learning while promoting goal-directed and collision-free navigation. The resulting trained agents can surpass the experts, and achieve better performance than pure learning-based agents without any expert reward signal and planning-based agents, while behaving in a human-like and more adept manner.

Overall, this paper makes the following contributions:

- 1) We introduce a reinforcement learning approach for human-inspired multi-agent navigation that exploits human trajectory demonstrations and knowledge distillation to train a collision avoidance policy in homogeneous, fully decentralized settings.
- 2) We experimentally show that the trained policy enables agents to take human-like actions for collision avoidance and goal-directed steering, and can generalize well to unseen scenarios not provided by the demonstrations, surpassing the performance of the experts and that of pure reinforcement-learning based agents or planning-based agents.
- 3) We provide related code and pre-trained policies that can be used by the research community as baselines to facilitate further research and advance the field of human-inspired multi-agent navigation.

II. RELATED WORK

A. Multi-Agent Navigation

State-of-the-art techniques for decentralized multi-agent navigation can be broadly classified into local planning ap-

*This work was supported in part by the National Science Foundation under Grant No. IIS-2047632 and by an Amazon Research Award.

The authors are with the School of Computing at Clemson University, South Carolina, USA. {peix, ioannis}@clemson.edu

proaches and learning-based methods. Existing local planning approaches that rely on social forces, energy-based formulations, and rule-based techniques have been successfully applied to a variety of domains and have been shown to generate human-like collision avoidance behavior [21–23]. In robotics, geometric local planners based on the concepts of velocity obstacles and time to collision [1, 2, 24] are widely applicable as they provide formal guarantees about the collision-free behavior of the agents and can be extended to account for motion and sensing uncertainty allowing implementation on actual robots [3, 25, 26]. However, despite their robustness, local planning approaches typically require careful parameter tuning which can limit their applicability to unseen environments. In addition, such approaches typically rely on user-defined assumptions about the optimal motion principles that govern the interactions between the agents.

Learning-based decentralized methods typically employ a reinforcement learning paradigm to address the problem of limited training data, allowing agents to learn navigation policies through interactions with the environment [27, 28]. Such methods do not make assumptions explicitly about what the optimal policy should be, but rather let the agents learn that policy through trial and error based on a reward function. Despite lacking theoretical collision avoidance guarantees, such approaches allow for better generalization to new environments and conditions as compared to local planning methods, with some of the most recent works enabling crowd-aware navigation for physical robots [6, 29, 30] as well as fully distributed multi-robot navigation [5, 31]. Our work is complementary to such learning-based methods, as we consider a reinforcement learning framework combined with imitation learning to train a human-like navigation policy applicable to homogeneous multi-agent navigation settings.

B. Imitation Learning

Below we consider imitation learning approaches that rely on limited number of expert demonstrations collected offline. Under this assumption, prior work has focused on extracting an action policy to generate expert-alike controls. Behavior cloning methods learn policies by directly matching the state-action pair as the input and output of the policy to the expert demonstrations through supervised learning [12, 32]. Such approaches have also been explored for motion planning and distributed multi-robot navigation where expert demonstrations are based on simulations [4, 13]. Similar ideas are also applicable when the expert is represented by a distribution policy such that we can perform policy distillation to minimize the divergence or crossing entropy between the target policy and the expert policies [33, 34]. Inverse reinforcement learning (IRL) [16] methods estimate a parameterized reward function based on the expert trajectories and perform training using reinforcement learning. IRL has been successfully applied for robot navigation through crowded environments in [35–37]. Generative adversarial imitation learning (GAIL) [17] have also been recently exploited for socially compliant robot navigation using raw depth images [38].

While highly relevant, the aforementioned methods have

difficulties when applied to the task of multi-agent human-like collision avoidance learning in general environments. A typical assumption in imitation learning is that the expert policies or demonstrations are reliable. However, the demonstrations provided by pedestrian datasets are biased to some degree as a dataset can only contain a limited number of human-human interactions in certain situations. Also, there is a lot of uncertainty in human decision making which cannot be captured by trajectory datasets. To address these issues, we leverage the idea of knowledge distillation [18], and perform optimization implicitly through reward shaping during reinforcement learning based on the imperfect expert policies learned from human pedestrian trajectory demonstrations.

III. APPROACH

We propose a reinforcement learning framework for collision-avoidance and human-inspired multi-agent navigation. We assume a homogeneous, fully decentralized setup consisting of holonomic disk-shaped agents that do not explicitly communicate with each other but share the same action policy to perform navigation based on their own observations. To generate human-like actions, we take advantage of publicly available human trajectory data and perform reward shaping through knowledge distillation in the learning process. Overall, our approach adopts a two-stage training process: (1) supervised learning on expert demonstrations from human trajectory data, and (2) reinforcement learning with knowledge distillation to generate a general action policy.

A. Problem Formulation

We consider a decentralized multi-agent control setting, where each agent acts *independently*, but all agents share the same action policy $\mathbf{a}_{i,t} \sim \pi_\theta(\cdot|\mathbf{o}_{i,t})$, with $\mathbf{a}_{i,t}$ denoting the action that the agent i samples at a given time step t , $\mathbf{o}_{i,t}$ is the observation that the agent receives at t , and θ denotes the parameters of the policy. The interaction process between an agent and the environment, which also includes other agents, is a partially observable Markov decision process, where each agent only relies on its own observations to make decisions.

For an N -agent environment, let $\mathcal{M}(\mathcal{S}_{t+1}|\mathcal{S}_t, \mathcal{A}_t)$ be the transient model of the Markov decision process with state space \mathcal{S} and action space \mathcal{A} . Then, at time t , $\mathbf{o}_{i,t} \in \mathcal{S}_t$ is a partial observable state for the i -th agent and $\mathcal{A}_t = \{\mathbf{a}_{i,t} | i = 1, \dots, N\}$. Given a time horizon T and the trajectory $\tau = (\mathcal{S}_1, \mathcal{A}_1, \dots, \mathcal{A}_{T-1}, \mathcal{S}_T)$, we optimize the parameter θ by maximizing the cumulative reward

$$J(\theta) = \mathbb{E}_{\tau \sim p_\theta(\tau)} [\mathcal{R}_t(\tau)] \quad (1)$$

where $\mathcal{R}_t(\tau) = \sum_{t=1}^{T-1} \sum_{i=1}^N \gamma^{t-1} r_{i,t}(\mathcal{S}_t, \mathbf{a}_{i,t}, \mathcal{S}_{t+1})$ is the total cumulative reward achieved by all agents with the step reward $r_{i,t}(\mathcal{S}_t, \mathbf{a}_{i,t}, \mathcal{S}_{t+1})$ and discount factor γ , and $p_\theta(\tau)$ is the state-action visitation distribution satisfying $p_\theta(\tau) = p(\mathcal{S}_1) \prod_{t=1}^{T-1} \mathcal{M}(\mathcal{S}_{t+1}|\mathcal{S}_t, \mathcal{A}_t) \prod_{i=1}^N \pi_\theta(\mathbf{a}_{i,t}|\mathbf{o}_{i,t})$.

Given that the decision process of each agent is completely independent under the fully decentralized context, the problem in Eq. 1 can be solved using a standard reinforcement learning

setup by optimizing the combined cumulative reward of all agents based on their own trajectories, i.e.:

$$J(\theta) = \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\tau_i \sim p_\theta(\tau_i)} \left[\sum_{t=1}^T \gamma^{t-1} r_{i,t}(\mathcal{S}_t, \mathbf{a}_{i,t}, \mathcal{S}_{t+1}) \right] \quad (2)$$

where $\tau_i = (\mathbf{o}_{i,1}, \mathbf{a}_{i,1}, \dots, \mathbf{a}_{i,T-1}, \mathbf{o}_{i,T})$ is the observation-action trajectory for the i -th agent itself.

B. Observation and Action Space

The observation space of each agent i is defined using a local coordinate system based on the agent's current position, $\mathbf{p}_{i,t} \in \mathbb{R}^2$, and velocity, $\mathbf{v}_{i,t} \in \mathbb{R}^2$, as: $\mathbf{o}_{i,t} = [\mathbf{o}_{i,t}^{n_0}, \mathbf{o}_{i,t}^{n_1}, \dots, \mathbf{o}_{i,t}^{n_m}, \mathbf{o}_{i,t}^g, \mathbf{o}_{i,t}^v]$, where

- $\mathbf{o}_{i,t}^{n_j} = [\mathbf{p}_{j,t} - \mathbf{p}_{i,t}, \mathbf{v}_{j,t} - \mathbf{v}_{i,t}]$ denotes the local state of the j -th neighbor. We assume that the agent has an observing radius r , and another agent j is considered as its neighbor if $j \in \{j : \|\mathbf{p}_{j,t} - \mathbf{p}_{i,t}\| \leq r, \forall j \neq i\}$.
- $\mathbf{o}_{i,t}^{n_0} = [\mathbf{0}, \mathbf{0}]$ is the neighborhood representation for the agent itself. This is a dummy neighbor representation to help data process when there is no neighbor in the observable range.
- $\mathbf{o}_{i,t}^g = \mathbf{g}_i - \mathbf{p}_{i,t}$ is the relative goal position where $\mathbf{g}_i \in \mathbb{R}^2$ is the goal position of the agent defined in the global coordinate system.
- $\mathbf{o}_{i,t}^v = \mathbf{v}_{i,t}$ is the current agent's velocity.

The action space is a continuous 2D space denoting the expected velocity of the agent at the next time step, i.e. $\mathbf{a}_{i,t} = \hat{\mathbf{v}}_{i,t+1}$. When applied on the agent, the expected velocity is scaled to satisfy the maximal speed v_i^{max} at which agent i can move:

$$\mathbf{v}_{i,t+1} = \begin{cases} \hat{\mathbf{v}}_{i,t+1} & \text{if } \|\hat{\mathbf{v}}_{i,t+1}\| \leq v_i^{max} \\ v_i^{max} \hat{\mathbf{v}}_{i,t+1} / \|\hat{\mathbf{v}}_{i,t+1}\| & \text{otherwise.} \end{cases} \quad (3)$$

C. Reward

We employ a reward function that gives a high reward signal $r_{arrival}$ as a bonus if an agent reaches its goal, and a negative reward signal $r_{collision}$ as a penalty to those that collide with any other agent. These reward terms urge the agent to reach its goal in as few steps as possible while avoiding collisions. As opposed to introducing additional reward terms to regularize the agent trajectories, such as penalties for large angular speeds [5] and uncomfortable distances [30] or terms that promote social norms [10], our work mainly relies on human expert policies, $f_e(\mathbf{o}_{i,t})$, extracted from real crowds trajectories to shape the reward.

In particular, we seek for the agents to exhibit navigation behavior in the style of the experts without explicitly dictating their behaviors. However, expert policies obtained by learning from human trajectory demonstrations are typically not general enough to lead agents to effectively resolving collisions in unseen environments. As such, we avoid directly optimizing the action policies through cloning behaviors from imperfect experts by introducing auxiliary loss terms in the objective function (Eq. 2). Instead, we propose to

use *knowledge distillation* and exploit the behavior error $\frac{1}{K} \sum_{e=1}^K \|\mathbf{a}_{i,t} - f_e(\mathbf{o}_{i,t})\|$ averaged over K expert policies to shape the *reward function*.

During inference, expert policies would become unreliable if the agent's speed is significantly different from the demonstrations. Humans typically prefer to walk at a certain *preferred speed* depending on the environment and their own personality traits, and their behavior patterns may differ a lot at different walking speed. As such, we introduce a velocity regularization term in the reward function to encourage each agent to move at a preferred speed v^* similar to the ones observed in the demonstrations. To promote goal seeking behavior and avoid meaningless exploration led by imperfect experts, v^* is used to scale the goal velocity that points from an agent's current position to its goal.

Given an environment with N agents, the complete reward function of our proposed imitation learning approach with knowledge distillation from K experts is defined as follows

$$r_{i,t}(\mathcal{S}_t, \mathbf{a}_{i,t}, \mathcal{S}_{t+1}) = \begin{cases} r_{arrival} & \text{if } \|\mathbf{g}_i - \mathbf{p}_{i,t}\| \leq R \\ r_{collision} & \text{if } \|\mathbf{p}_{j,t+1} - \mathbf{p}_{i,t+1}\| \leq 2R \\ w_e r_{i,t}^e + w_v r_{i,t}^v & \text{otherwise} \end{cases} \quad (4)$$

where $j = 1, \dots, N$ with $j \neq i$, R is the agent radius, and w_v and w_e are weight coefficients. The knowledge distillation reward term $r_{i,t}^e$ and the velocity regularization term $r_{i,t}^v$ are computed as:

$$r_{i,t}^e = \exp \left(\frac{\sigma_e}{K} \sum_{e=1}^K \|\mathbf{a}_{i,t} - f_e(\mathbf{o}_{i,t})\| \right), \quad r_{i,t}^v = \exp \left(\sigma_v \|\mathbf{v}_{i,t+1} - \mathbf{v}_{i,t+1}^*\| \right) \quad (5)$$

where σ_e and σ_v are scale coefficients, and $\mathbf{v}_{i,t+1}^* = v^*(\mathbf{g}_i - \mathbf{p}_{i,t}) / \|\mathbf{g}_i - \mathbf{p}_{i,t}\|$ is the goal velocity of the agent i having a magnitude equal to the preferred speed v^* .

D. Policy Learning

The learning process of our approach has two stages: (1) obtain expert policies by supervised learning from real human pedestrian trajectories, and (2) train a shared action policy in multi-agent environments through reinforcement learning to maximize Eq. 2 with the reward function defined in Eq. 4.

Supervised Learning. Expert policies are trained by supervised learning on human pedestrian trajectory datasets where the observation-action training data $\{(\mathbf{o}_{i,t}, \mathbf{a}_{i,t})\}$ is extracted as described in Section III-B. The goal position is defined by the last position of each pedestrian appearing in the datasets. We use a neural network denoted by $f_e(\cdot | \phi_e)$ with parameters ϕ_e to perform action prediction, which is optimized by minimizing the mean squared error between the predicted and the ground truth action:

$$\mathcal{L}_{sup} = \mathbb{E}_{(\mathbf{o}_{i,t}, \mathbf{a}_{i,t})} [\|f_e(\mathbf{o}_{i,t} | \phi_e) - \mathbf{a}_{i,t}\|^2] \quad (6)$$

Trajectory data in human datasets is typically recorded by sensors at fixed and rather sparse time intervals. Training on such limited, discrete data is easily prone to overfitting, as the network simply remembers all the data. To alleviate

this issue, we perform data augmentation during training by randomly sampling a continuous timestamp and using linear interpolation to estimate corresponding observation-action data from discrete trajectories. In addition to linear interpolation, we adopt two more types of data augmentation: scenario flipping and rotation. During training, the x- and/or y-axis of each observation-action datum is flipped by a stochastic process along with randomly rotating the observation-action local system. This helps increase the generality of the training samples without destroying the relational information between human-human interactions captured in the dataset.

While multiple expert policies can be trained by using different datasets, the action patterns of pedestrians recorded at different times and/or under different scenarios may vary a lot. This could lead to too much divergence when we ensemble the expert policies in the reward function of the reinforcement learning stage (cf. Eq 4). As such, in our implementation, we extract expert policies only from one pedestrian trajectory dataset. Under this setting, the deviation caused by stochastic factors during expert training can be effectively eliminated by data augmentation (Section IV-E). This allows us to employ just a single expert policy in all of our experiments and obtain a final action policy that is comparable to the one obtained with multiple experts but at a significant training speedup.

Reinforcement Learning: To perform reinforcement learning, we exploit the DPPO algorithm [39], which is a distributed training scheme that relies on the PPO algorithm [40]. DPPO optimizes Eq. 2 using policy gradient method [41] by maximizing $J(\theta) = \frac{1}{N} \sum_i^N \mathbb{E}_t \left[\log \pi_\theta(\bar{\mathbf{a}}_{i,t} | \bar{\mathbf{o}}_{i,t}) \hat{A}_{i,t} \right]$, where $\hat{A}_{i,t}$ is an estimation to the discounted cumulative reward term with bias subtraction, which in our implementation is computed by the generalized advantage estimation (GAE) [42].

To help exploration, we also introduce differential entropy loss to encourage exploration and avoid premature convergence, resulting the following objective function:

$$\mathcal{L}_{re}(\theta) = \frac{1}{N} \sum_i^N \mathbb{E}_t \left[\log \pi_\theta(\bar{\mathbf{a}}_{i,t} | \bar{\mathbf{o}}_{i,t}) \hat{A}_{i,t} + \beta \mathcal{H}(\pi_\theta(\cdot | \bar{\mathbf{o}}_{i,t})) \right] \quad (7)$$

where \mathcal{H} represents the entropy of the action policy, and β is a scalar coefficient of the policy entropy loss. The parameters $\bar{\mathbf{o}}_{i,t}$ and $\bar{\mathbf{a}}_{i,t}$ denote the observation and action, respectively, after performing goal direction alignment, i.e. after rotating the local coordinate system described in III-B such that the goal is always at a fixed direction (the positive direction of x axis in our implementation). By this method, we can remove one dimension of $\mathbf{o}_{i,t}^g$ and use the distance to the goal position instead, i.e., $\bar{o}_{i,t}^g = \|\mathbf{g}_i - \mathbf{p}_{i,t}\|$. This trick helps the learning, as it reduces the complexity of state space and is beneficial to exploration. Intuitively, agents in the goal-aligned local system are more likely to encounter similar observations and thus exploit previous experience more effectively.

We do not apply the goal-alignment trick during expert policy training, since it could result in too much overfitting. In human trajectory datasets, pedestrians move mainly towards their goals. With the alignment trick the output actions

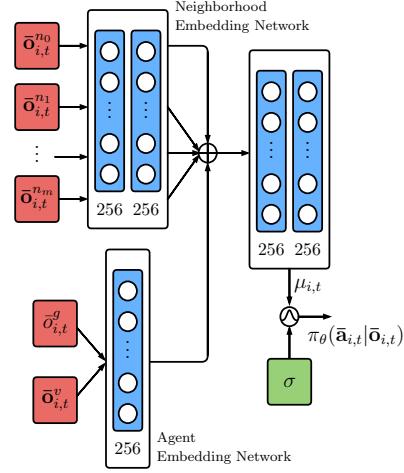


Fig. 1. Network architecture employed as the policy network during reinforcement learning. The policy is defined as a Gaussian distribution with $\mu_{i,t}$ as the mean value and an input-independent parameter σ as the standard deviation. The same architecture is adopted for the value network as well as the expert policy learning.

(velocities) would mostly respond to the goal direction. As a result, expert policies would prefer orientation adaption more than speed adaptation which can often lead to understeering. *Network architecture.* We use a bivariate Gaussian distribution with independent components as the action policy, where the mean values are provided by the policy network with architecture shown in Fig. 1, and the standard deviation is a 2-dimension observation-independent parameter. Neighbor representations, $\bar{o}_{i,t}^{n_0}$ and $\bar{o}_{i,t}^{n_1}, \dots, \bar{o}_{i,t}^{n_m}$, are simply added together with the agent representation, $[\bar{o}_{i,t}^g, \bar{o}_{i,t}^v]$, after the embedding networks. $\bar{o}_{i,t}^{n_0} = \mathbf{o}_{i,t}^{n_0} = [\mathbf{0}, \mathbf{0}]$ denotes the neighborhood representation of the agent itself and is always kept such that the network can work normally when there is no neighbors observed. This network architecture can support observations with an arbitrary number of neighbors and the embedding result does not rely on the input order of the neighbor representation sequence. We use the same architecture for the value network to perform value estimation for the GAE computation, and to train the expert policies through supervised learning without goal alignment observations.

IV. EXPERIMENTS

A. Simulation Environment Setup

Our simulation environment has three scenarios shown in Fig. 2. In the circle scenarios, agents are roughly placed on the circumference of a circle and have to reach their antipodal positions; in the corridor scenarios, agents are randomly initialized at the two sides, oriented either vertically or horizontally, and have random goals at the opposite sides; and in the square crossing scenarios, agents are placed at the four sides and need to walk across the opposite sides to reach randomly placed goals. To increase the generality, all scenarios are generated randomly during simulation. The circle scenario has a random radius varying from 4m to 6m. The other two scenarios have widths and heights in the range of 8m to 12m.

TABLE I
QUANTITATIVE EVALUATION

		20-Circle	24-Circle	20-Corridor	24-Corridor	20-Square	24-Square
Success Rate (higher is better)	ORCA	0.91 ± 0.15	0.71 ± 0.17	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00
	SL	0.40 ± 0.17	0.23 ± 0.12	0.64 ± 0.15	0.59 ± 0.14	0.64 ± 0.15	0.58 ± 0.13
	RL	0.97 ± 0.05	0.96 ± 0.05	0.93 ± 0.08	0.92 ± 0.08	0.91 ± 0.09	0.88 ± 0.11
	Ours	1.00 ± 0.01	0.99 ± 0.02	1.00 ± 0.00	0.99 ± 0.03	0.99 ± 0.03	0.97 ± 0.06
Extra Distance (lower is better)	ORCA	0.17 ± 0.33	0.27 ± 0.41	0.13 ± 0.47	0.20 ± 0.71	0.23 ± 0.49	0.40 ± 0.80
	SL	3.80 ± 4.67	6.37 ± 6.39	6.06 ± 8.49	7.29 ± 9.32	5.83 ± 7.85	6.92 ± 9.15
	RL	0.81 ± 0.69	0.93 ± 0.73	0.25 ± 0.44	0.31 ± 0.53	0.43 ± 0.61	0.48 ± 0.59
	Ours	0.34 ± 0.30	0.49 ± 0.29	0.09 ± 0.08	0.01 ± 0.09	0.03 ± 0.13	0.04 ± 0.12
Energy Efficiency (higher is better)	ORCA	0.34 ± 0.11	0.27 ± 0.17	0.35 ± 0.07	0.35 ± 0.09	0.32 ± 0.10	0.30 ± 0.11
	SL	0.06 ± 0.20	0.04 ± 0.20	-0.02 ± 0.20	-0.02 ± 0.19	-0.02 ± 0.20	-0.04 ± 0.19
	RL	0.29 ± 0.05	0.28 ± 0.06	0.30 ± 0.04	0.30 ± 0.05	0.29 ± 0.06	0.29 ± 0.06
	Ours	0.36 ± 0.05	0.36 ± 0.05	0.36 ± 0.07	0.37 ± 0.05	0.36 ± 0.06	0.36 ± 0.05

During each simulation, 6 to 20 agents, sharing the same action policy under optimization, are placed into the scenarios and are given stochastic goal positions. The simulation time step size is 0.04s, while agents receive control signals every 0.12s. All agents are modeled as disks with a radius of 0.1m, and have a preferred speed v^* of 1.3m/s with a maximum speed limit of 2.5m/s to imitate human walking patterns in the dataset for expert policy training. During simulation, collided agents are kept as static obstacles to other agents, while agents that arrive at their goals are removed from the environment. Each simulation episode terminates if all non-collided agents reach their goals or a time limit of 120s is met.

B. Training Details

We exploit the *students* dataset [43] for supervised learning of expert policies. This dataset records the trajectories of 434 pedestrians in an open campus environment, providing a more general setting to learn typical expert policies as compared to environments in other publicly available datasets consisting of too many obstacles and/or terrain restrictions. Each pedestrian is considered as a disk agent with a radius of 0.1m, estimated by the typical minimal distance between trajectories; the goal position is chosen as the last position of each pedestrian's trajectory; velocities are estimated by finite differences.

Before training, we preprocess the raw data to remove pedestrians that stand still or saunter without clear goals, as trajectories of such pedestrians cannot help learn navigation policies and would become training noise. After data cleansing, we kept the observation-action records from 300 out of the 434 pedestrians for expert training (Eq. 6), while the rest of the pedestrian trajectories are only used as part of the neighborhood representation of the active pedestrians. During the following experiments, we, by default, use only one expert policy trained with data augmentation (see Section IV-E for related analysis). The default values of the reward function weights are: $w_e = 0.02$, $w_v = 0.08$, $\sigma_e = \sigma_v = 0.85$. We refer to <https://motion-lab.github.io/KDMA> for all hyperparameters used, along with related videos and code.

C. Quantitative Evaluation

We compare our method to three approaches: the geometric-based method of Optimal Reciprocal Collision Avoidance

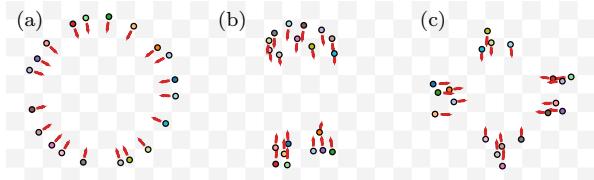


Fig. 2. Our environment consists of: (a) a circle crossing scenario, (b) a corridor scenario, and (c) a square crossing scenario. During each simulation episode, a scenario is randomly chosen and 6 to 20 agents are put into the scenario randomly. Red arrows indicate goal directions.

(ORCA) [2], a supervised learning approach (SL), and a reinforcement learning approach without knowledge distillation (RL). To prevent ORCA agents from staying too close to each other, we increase the agent radius by 20% in the ORCA simulator. SL denotes the performance of the expert policy obtained in our approach. RL uses the reward function from [5], which optimizes agents to reach their goals as fast as possible without collisions. Since RL performs optimization based on the maximal agent speed but our performance evaluation are computed based on a preferred speed of 1.3m/s, for fairness, we set the maximal speed in RL training and testing cases as 1.3m/s instead of the default value of 2.5m/s.

We perform 50 trials of testing, which are the same across different methods, using the three scenarios shown in Fig. 2. We run comparisons using three evaluation metrics and report the results in Table I. Reported numbers are the mean ± std. *Success rate* denotes the percentage of agents that reached their goals. *Extra distance* is the additional distance in meters that an agent traversed instead of taking a straight line path to the goal. This metric is recorded only for agents that reached their goals. *Energy efficiency* measures the ratio between an agent's progress towards its goal and its energy consumption at each time step [8]. The energy consumption is computed following the definition in [8] such that the optimal speed in terms of energy expended per meter traversed is 1.3m/s.

As shown in Table I, our approach can achieve as good as or better results than ORCA in terms of success rate and energy efficiency, while it significantly outperforms RL and SL in most of the scenarios. Given that the expert policies obtained by SL are used in the second stage of our approach, it is



Fig. 3. Trajectories of different methods in 2-agent head-to-head scenario.

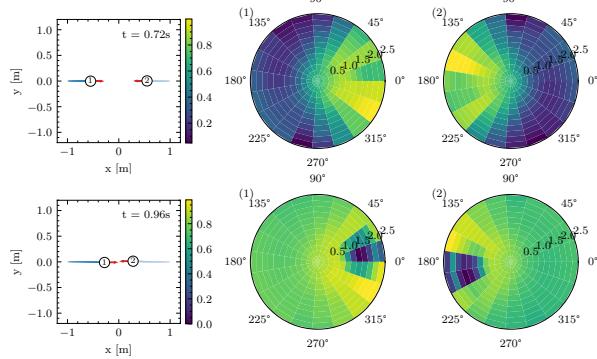


Fig. 4. Value heatmaps of our approach in the 2-agent interaction scenario. Brighter radial bins indicate the action decisions that are considered better by the agents. Dark bins indicate the actions that agents do not prefer.

clear that agents with knowledge distillation can surpass the experts to achieve collision-free and energy-efficient behavior patterns. Regarding the extra distance traversed, agents with our approach favor shortest paths in the corridor and square scenarios. In the circle scenarios, ORCA outperforms the methods as its agents prefer to take straight line paths to the goal that passes through the center of the environment by mainly adapting their speeds. However, as discussed below these are not typically the types of trajectories that humans would exhibit. In addition, we note that the behavior of ORCA agents varies a lot depending on the choices of parameters used, including the size of the simulation time step and the safety buffer added to the radius of the agent.

D. Comparisons to Human Reference Trajectories

Figure 3 compares the trajectories generated by different methods in a symmetric head-to-head interaction scenario to reference data obtained from motion captured human experiments [44]. Due to the symmetric nature of the scenario, ORCA agents fail to reach their goals and end up standing still next to each other in the middle of the environment. RL agents maintain a high walking speed close to their maximal one and solve the interaction by performing an early orientation adaptation as opposed to the reference data where the avoidance happens later on. Trajectories generated by our approach more closely match the reference trajectories. To further analyze the decision making of the agents in our approach, Fig. 4 depicts the corresponding action value heatmaps obtained by estimation through the value network of DPPO. As shown, at 0.72s the agents implicitly negotiate to pass each other from the right side; and later on at the 0.92s mark every action that can lead to an imminent collision is forbidden to both agents, with agent 1 preferring to turn a bit to the right while moving at a speed of at least 1m/s while agent 2 decides to veer a bit to the right and accelerate.

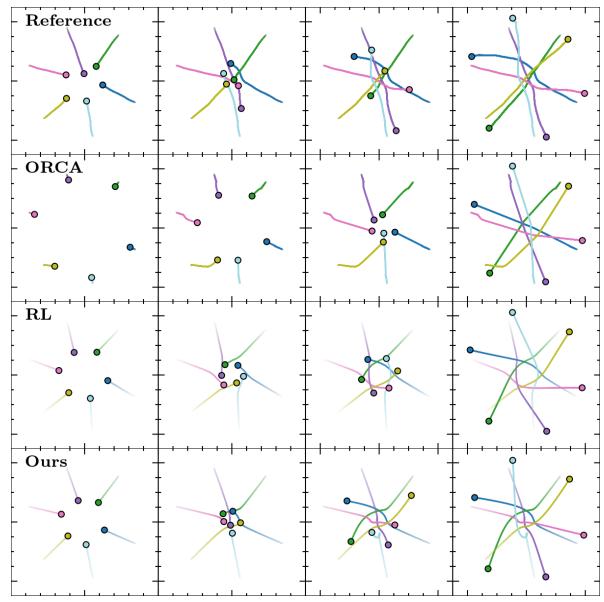


Fig. 5. Agent trajectories in a 6-agent circle scenario. Each column is captured at the same time fraction of the whole trajectories.

Figure 5 shows trajectories from a 6-agent circle scenario, where the reference human trajectories were obtained from [45]. In this scenario, ORCA agents mainly resolve collisions by adapting their speeds; they move slow at the beginning due to the uncertainty about what their neighbors are doing and eventually do a fast traverse mostly along straight lines towards their goals. RL agents, on the other hand, prefer to travel at their maximum speed towards the center of the environment, and then adapt their orientations in unison resulting in a vortex-like avoidance pattern. Agents in our approach exhibit more variety by using a combination of both speed and orientation adaptation to balance the tradeoff between collision avoidance and goal steering. As such, the resulting trajectories can capture to some degree the diversity that humans exhibit in the reference data.

To statistically analyze the quality of generated trajectories, we reproduce the *students* scenario with different methods and run novelty detection using the *k*-LPE algorithm [46]. As a comparison, we also introduce another baseline, the PowerLaw method [23], which is a state-of-the-art crowd simulation method elaborated from the analysis of pedestrian datasets including the *students* one. Figure 6 shows how the corresponding agent trajectories compare to the trajectories in the *students* reference dataset based on the metrics of the agent's linear speed, angular speed, and distance to the nearest neighbor. As it can be seen, the majority of trajectories generated by our approach have low anomaly scores as compared to ORCA and RL. Figure 7 further highlights that agents in our approach can more closely match the velocity distribution of the ground truth data. We note that PowerLaw provides high fidelity simulation as well. However, to guarantee numerical stability, a very small simulation time step is typically required during each sensing-acting cycle, which may not always be applicable to real robots.

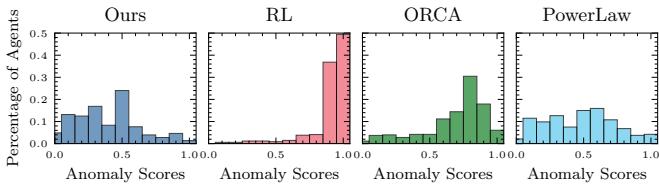


Fig. 6. Distributions of anomaly scores on *students* scenario.

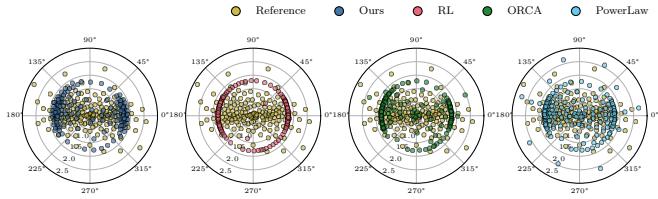


Fig. 7. Agent velocity distributions on *students* scenario.

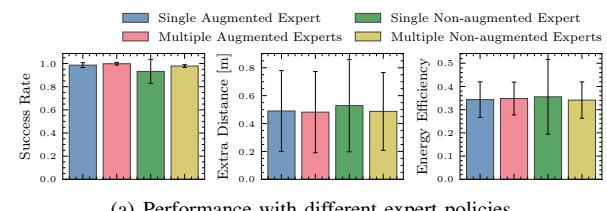
E. Sensitivity Analysis

Figure 8(a) shows the performance of action policies trained with different experts in the 24-Circle scenario. As it can be seen, the performance varies a lot when employing a single expert trained without data augmentation. However, the uncertainty caused by stochastic factors of supervised training can be effectively eliminated by introducing data augmentation during expert policy training. Using either multiple augmented experts or a single expert with data augmentation results in comparable performances. As such, and taking into account the fact that the inference time increases dramatically as more expert policies are exploited simultaneously, we employed a single expert policy with data augmentation in all of our experiments. In our tests, using two V100 GPUs, it took approximately six hours to perform training with three expert policies, while training with a single expert took only around three and a half hours.

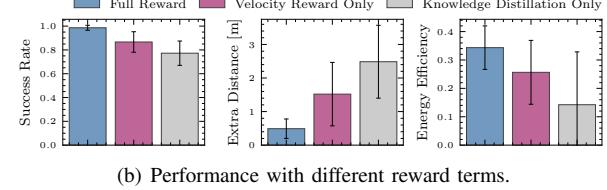
Figure 8(b) compares the training results obtained by using the full reward function that considers both the knowledge distillation term and the velocity regularization term (Eq. 5) to results obtained by using only one of the two terms. As shown in the figure, the expert policy by itself is not quite reliable, as it leads to a worse performance across all three metrics when only the knowledge distillation term is used. However, when combined with the velocity reward term, the expert can help agents achieve the best overall performance.

V. LIMITATIONS AND FUTURE WORK

We propose a framework for learning a human-like general collision avoidance policy for agent-agent interactions in fully decentralized multi-agent environments. To do so, we use knowledge distillation implicitly in a reinforcement learning setup to shape the reward function based on expert policies extracted from human pedestrian trajectory demonstrations. Our approach can help agents surpass the experts, and achieve better performance and more human-like action patterns, compared to using reinforcement learning without knowledge distillation and to existing geometric planners. During testing, agents use the trained navigation policy deterministically. In



(a) Performance with different expert policies.



(b) Performance with different reward terms.

Fig. 8. Sensitivity analysis based on performance in the 24-agent circle. (a) Multiple experts testing cases use three experts simultaneously for reward computation. Single expert results are the average of three training trials each of which uses only one of the corresponding three experts. (b) Comparisons of different reward terms using a single expert with data augmentation.

the future, to account for the uncertainty of human navigation decisions, we would like to test stochastic actions by keeping the executed policy as a Gaussian distribution instead of using only the mean value.

A limitation of our approach is that we assume all agents are of the same type, and particularly are holonomic disk-shaped agents. While, in theory, the same trained policy can be ported to other types of agents, the resulting behavior can be very conservative as the true geometry of the target agent will be ignored during training. In addition, we assume that trained agents can behave similarly to the expert pedestrians that provide the demonstrations, ignoring the kinodynamic constraints of specific robot platforms. Even though this issue can be addressed by relying on a controller to convert human-like velocity commands to low-level robot control inputs, it also opens an interesting direction for future work that focuses on mining robot-friendly action patterns rather than action commands from human trajectories. Another avenue for future work is extending our method to mixed settings where agents can interact with humans. The recent works of [11, 30, 47] on socially-aware robot navigation and of [48, 49] on generating socially compliant behaviors in multi-agent environments can provide some interesting ideas towards this research direction.

REFERENCES

- [1] P. Fiorini and Z. Shiller, "Motion planning in dynamic environments using velocity obstacles," *The International Journal of Robotics Research*, vol. 17, pp. 760–772, 1998.
- [2] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics research*. Springer, 2011, pp. 3–19.
- [3] J. Alonso-Mora, A. Breitenmoser, M. Rufli, P. Beardsley, and R. Siegwart, "Optimal reciprocal collision avoidance for multiple non-holonomic robots," in *Distributed autonomous robotic systems*. Springer, 2013, pp. 203–216.
- [4] P. Long, W. Liu, and J. Pan, "Deep-learned collision avoidance policy for distributed multiagent navigation," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 656–663, 2017.
- [5] 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 *IEEE International Conference on Robotics and Automation*, 2018, pp. 6252–6259.

[6] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *IEEE International Conference on Robotics and Automation*, 2019, pp. 6015–6022.

[7] D. Hildreth and S. J. Guy, "Coordinating multi-agent navigation by learning communication," *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, vol. 2, no. 2, pp. 1–17, 2019.

[8] J. Godoy, S. J. Guy, M. Gini, and I. Karamouzas, "C-nav: Distributed coordination in crowded multi-agent navigation," *Robotics and Autonomous Systems*, vol. 133, p. 103631, 2020.

[9] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese, "Social LSTM: Human trajectory prediction in crowded spaces," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 961–971.

[10] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 1343–1350.

[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] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.

[13] M. Pfeiffer, M. Schaeuble, J. Nieto, R. Siegwart, and C. Cadena, "From perception to decision: A data-driven approach to end-to-end motion planning for autonomous ground robots," in *IEEE International Conference on Robotics and Automation*, 2017, pp. 1527–1533.

[14] J. Peters and S. Schaal, "Reinforcement learning of motor skills with policy gradients," *Neural networks*, vol. 21, no. 4, pp. 682–697, 2008.

[15] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," *arXiv preprint arXiv:1709.10087*, 2017.

[16] A. Y. Ng and S. Russell, "Algorithms for inverse reinforcement learning," in *17th International Conf. on Machine Learning*. Citeseer, 2000.

[17] J. Ho and S. Ermon, "Generative adversarial imitation learning," in *Advances in Neural Information Processing Systems*, 2016, pp. 4565–4573.

[18] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.

[19] D. P. Bertsekas, "Approximate policy iteration: A survey and some new methods," *Journal of Control Theory and Applications*, vol. 9, no. 3, pp. 310–335, 2011.

[20] A. Nair, B. McGrew, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Overcoming exploration in reinforcement learning with demonstrations," in *IEEE International Conference on Robotics and Automation*, 2018, pp. 6292–6299.

[21] C. W. Reynolds, "Steering behaviors for autonomous characters," in *Game Developers Conference*, 1999, pp. 763–782.

[22] D. Helbing, I. Farkas, and T. Vicsek, "Simulating dynamical features of escape panic," *Nature*, vol. 407, no. 6803, pp. 487–490, 2000.

[23] I. Karamouzas, B. Skinner, and S. J. Guy, "Universal power law governing pedestrian interactions," *Physical Review Letters*, vol. 113, no. 23, p. 238701, 2014.

[24] J. Van den Berg, M. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *IEEE International Conference on Robotics and Automation*, 2008, pp. 1928–1935.

[25] D. Bareiss and J. van den Berg, "Generalized reciprocal collision avoidance," *The International Journal of Robotics Research*, vol. 34, no. 12, pp. 1501–1514, 2015.

[26] B. Davis, I. Karamouzas, and S. J. Guy, "NH-TTC: A gradient-based framework for generalized anticipatory collision avoidance," in *Robotics: Science and Systems*, 2020.

[27] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 1343–1350.

[28] Y. F. Chen, M. Liu, M. Everett, and J. P. How, "Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning," in *IEEE International Conference on Robotics and Automation*, 2017, pp. 285–292.

[29] M. Everett, Y. F. Chen, and J. P. How, "Motion planning among dynamic, decision-making agents with deep reinforcement learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2018, pp. 3052–3059.

[30] Y. Liu, Q. Yan, and A. Alahi, "Social nce: Contrastive learning of socially-aware motion representations," *arXiv preprint arXiv:2012.11717*, 2020.

[31] T. Fan, P. Long, W. Liu, and J. Pan, "Distributed multi-robot collision avoidance via deep reinforcement learning for navigation in complex scenarios," *The International Journal of Robotics Research*, vol. 39, no. 7, pp. 856–892, 2020.

[32] D. A. Pomerleau, "ALVINN: An autonomous land vehicle in a neural network," in *Advances in Neural Information Processing Systems*, D. Touretzky, Ed., vol. 1. Morgan-Kaufmann, 1989.

[33] A. A. Rusu, S. G. Colmenarejo, C. Gulcehre, G. Desjardins, J. Kirkpatrick, R. Pascanu, V. Mnih, K. Kavukcuoglu, and R. Hadsell, "Policy distillation," *arXiv preprint arXiv:1511.06295*, 2015.

[34] W. M. Czarnecki, R. Pascanu, S. Osindero, S. Jayakumar, G. Swirszcz, and M. Jaderberg, "Distilling policy distillation," in *The 22nd International Conference on Artificial Intelligence and Statistics*. PMLR, 2019, pp. 1331–1340.

[35] P. Henry, C. Vollmer, B. Ferris, and D. Fox, "Learning to navigate through crowded environments," in *2010 IEEE International Conference on Robotics and Automation*, 2010, pp. 981–986.

[36] H. Kretzschmar, M. Spies, C. Sprunk, and W. Burgard, "Socially compliant mobile robot navigation via inverse reinforcement learning," *The International Journal of Robotics Research*, vol. 35, no. 11, pp. 1289–1307, 2016.

[37] B. Kim and J. Pineau, "Socially adaptive path planning in human environments using inverse reinforcement learning," *International Journal of Social Robotics*, vol. 8, no. 1, pp. 51–66, 2016.

[38] L. Tai, J. Zhang, M. Liu, and W. Burgard, "Socially compliant navigation through raw depth inputs with generative adversarial imitation learning," in *IEEE International Conference on Robotics and Automation*, 2018, pp. 1111–1117.

[39] N. Heess, S. Sriram, J. Lemmon, J. Merel, G. Wayne, Y. Tassa, T. Erez, Z. Wang, S. Eslami, M. Riedmiller *et al.*, "Emergence of locomotion behaviours in rich environments," *arXiv preprint arXiv:1707.02286*, 2017.

[40] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.

[41] R. S. Sutton, D. A. McAllester, S. P. Singh, and Y. Mansour, "Policy gradient methods for reinforcement learning with function approximation," in *Advances in neural information processing systems*, 2000, pp. 1057–1063.

[42] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, "High-dimensional continuous control using generalized advantage estimation," *arXiv preprint arXiv:1506.02438*, 2015.

[43] A. Lerner, Y. Chrysanthou, and D. Lischinski, "Crowds by example," in *Computer graphics forum*, vol. 26, no. 3. Wiley Online Library, 2007, pp. 655–664.

[44] M. Moussaïd, D. Helbing, and G. Theraulaz, "How simple rules determine pedestrian behavior and crowd disasters," *Proceedings of the National Academy of Sciences*, vol. 108, no. 17, pp. 6884–6888, 2011.

[45] D. Wolinski, S. J. Guy, A.-H. Olivier, M. Lin, D. Manocha, and J. Pettré, "Parameter estimation and comparative evaluation of crowd simulations," in *Computer Graphics Forum*, vol. 33, no. 2. Wiley Online Library, 2014, pp. 303–312.

[46] M. Zhao and V. Saligrama, "Anomaly detection with score functions based on nearest neighbor graphs," *Advances in neural information processing systems*, vol. 22, pp. 2250–2258, 2009.

[47] A. J. Sathyamoorthy, J. Liang, U. Patel, T. Guan, R. Chandra, and D. Manocha, "Densecavoid: Real-time navigation in dense crowds using anticipatory behaviors," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 11345–11352.

[48] C. I. Mavrogiannis, W. B. Thomason, and R. A. Knepper, "Social momentum: A framework for legible navigation in dynamic multi-agent environments," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018, pp. 361–369.

[49] C. I. Mavrogiannis and R. A. Knepper, "Multi-agent trajectory prediction and generation with topological invariants enforced by hamiltonian dynamics," in *WAFR*, 2018, pp. 744–761.