



Core Challenges of Social Robot Navigation: A Survey

CHRISTOFOROS MAVROGIANNIS, Paul G. Allen School of Computer Science & Engineering, University of Washington, USA

FRANCESCA BALDINI, Honda Research Institute and California Institute of Technology, USA

ALLAN WANG and DAPENG ZHAO, The Robotics Institute, Carnegie Mellon University, USA

PETE TRAUTMAN, Honda Research Institute, USA

AARON STEINFELD and JEAN OH, The Robotics Institute, Carnegie Mellon University, USA

Robot navigation in crowded public spaces is a complex task that requires addressing a variety of engineering and human factors challenges. These challenges have motivated a great amount of research resulting in important developments for the fields of robotics and human-robot interaction over the past three decades. Despite the significant progress and the massive recent interest, we observe a number of significant remaining challenges that prohibit the seamless deployment of autonomous robots in crowded environments. In this survey article, we organize existing challenges into a set of categories related to broader open problems in robot planning, behavior design, and evaluation methodologies. Within these categories, we review past work and offer directions for future research. Our work builds upon and extends earlier survey efforts by (a) taking a critical perspective and diagnosing fundamental limitations of adopted practices in the field and (b) offering constructive feedback and ideas that could inspire research in the field over the coming decade.

CCS Concepts: • **Computing methodologies** → **Simulation evaluation; Reinforcement learning; Robotic planning**; • **Computer systems organization** → **Robotics; Robotic control**; • **Human-centered computing** → *User studies*;

Additional Key Words and Phrases: Social robot navigation, motion planning, motion prediction, multiagent systems, social robotics, benchmarking

ACM Reference format:

Christoforos Mavrogiannis, Francesca Baldini, Allan Wang, Dapeng Zhao, Pete Trautman, Aaron Steinfeld, and Jean Oh. 2023. Core Challenges of Social Robot Navigation: A Survey. *ACM Trans. Hum.-Robot Interact.* 12, 3, Article 36 (April 2023), 39 pages.

<https://doi.org/10.1145/3583741>

This work was funded under grants from the Honda Research Institute USA, the National Science Foundation (NSF IIS-1734361), the National Institute on Disability, Independent Living, and Rehabilitation Research (NIDILRR 90DPGE0003), the U.S. Army Ground Vehicle Systems Center & Software Engineering Institute at Carnegie Mellon University (PWP 6-652A4 DDCD GVSC), the Air Force Office of Scientific Research (AFOSR FA2386-17-1-4660), and the Ministry of Trade, Industry and Energy (MOTIE) and Korea Institute of Advancement of Technology (KIAT) through the International Cooperative R&D program: P0019782, Embedded AI Based fully autonomous driving software and Maas technology development.

Authors' addresses: C. Mavrogiannis, Paul G. Allen School of Computer Science & Engineering, University of Washington, 3920 E. Stevens Way NE Seattle, WA 98195, USA; email: cmavro@cs.washington.edu; F. Baldini, Honda Research Institute and California Institute of Technology, 1200 E California Blvd, Pasadena, CA 91125, USA; email: fbaldini@caltech.edu; A. Wang, D. Zhao, A. Steinfeld, and J. Oh, The Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA, 15213, USA; emails: allanwan@andrew.cmu.edu, tim.eagle.zhao@gmail.com, {steinfeld, jeanoh}@cmu.edu; P. Trautman, Honda Research Institute, 70 Rio Robles, San Jose, CA 95134, USA; email: peter.trautman@gmail.com.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2023 Copyright held by the owner/author(s).

2573-9522/2023/04-ART36 \$15.00

<https://doi.org/10.1145/3583741>

1 INTRODUCTION

The vision of mobile robots navigating seamlessly through public spaces has inspired a significant amount of research over the past three decades. While early efforts were driven by engineering principles, it quickly became evident that the problem is inherently interdisciplinary, requiring insights from fields such as human-robot interaction, psychology, design research, and sociology. This understanding resulted in the emergence of a dedicated field of research, which is commonly referred to as *social robot navigation*, broadly addressing all algorithmic, behavioral, evaluation, and design aspects of the problem.

Early work focused on the development of tour guide robots. The RHINO [26] and MINERVA [243] robots, deployed in Museums in Germany and the United States in the late 1990s, exemplified the vision of integrating robots in human environments. Leveraging state-of-the-art probabilistic methods for localization and mapping [53, 212, 213], these robots engaged with thousands of users under unscripted interaction settings, pioneering the nascent field of human-robot interaction in natural spaces [223, 246]. These deployments generated enthusiasm in the robotics community and motivated further research in many important areas including state estimation, human behavior prediction, robot motion planning, crowd simulation, user experience design, and more.

Over the more than 20 years following the RHINO and MINERVA deployments, many exciting research developments have emerged, for instance: Many lab studies and field deployments have contributed important empirical knowledge about robot navigation in crowds [18, 72, 124, 131, 165, 249]; the study and simulation of crowds [69, 105, 190, 282] has informed strategies on collision avoidance [65, 76, 261]; forecasting techniques from the motion tracking community are increasingly integrated into the social navigation stack to provide high-accuracy models of crowd dynamics [2, 90, 162]; (inverse) reinforcement learning techniques have been targeting a better data-driven understanding of human navigation objectives [39, 61, 140, 265, 288]; considerations of human comfort like the personal space have become central to the research agenda of roboticians [114, 132, 191]; inspiration from sociology [87, 279] and psychology [49] studies has shaped the design of novel control techniques [135, 165]. Recently, the interest in the field has been propelled even further as a result of the overlap of social navigation with neighboring, high-pace fields such as autonomous driving. Further, recent and ongoing large projects such as SPENCER [232] or CROWDBOT [48] funded by the European Commission are indicative of the momentum and growth of the field.

In this article, motivated by the great progress and interest in the field, we organize the literature into a set of core challenges that represent major classes of ongoing research and open problems. In particular, we focus on:

- Planning challenges: *How should a robot plan its motion through a crowd of navigating humans to ensure safety and efficiency?*
- Behavioral challenges: *What types of social signals can be inferred from human behavior, and what types of behaviors should a robot exhibit to ensure social norm compliance?*
- Evaluation challenges: *How can we evaluate a social robot navigation system?*

In our view, these categories contain some of the most important and demanding open questions that the community needs to address to make important progress in the field. For each class of challenges, we review relevant literature, diagnose issues, and offer research directions to address them. It should be noted that our perspective is driven by computational considerations and empirical insights. As such, the proposed categorization, the selection of the literature, and the discussion is done from a computational perspective. While this allows us to cover a wide range of important topics, there are additional challenges from the broader space of Robotics and **Human-Robot**

Interaction (HRI) that are outside of the scope of this review. Some of them (e.g., robot design, perception pipelines) are briefly mentioned in Section 5.1.

1.1 Contributions

Overall, this article contributes the following:

- A holistic discussion of the literature in the area of social robot navigation spanning three decades of work, featuring a discussion of past and active trends.
- An enumeration of fundamental challenges in social robot navigation, representing our perspective on the most important open problems that require further research.
- A discussion of possible research directions that could address the challenges identified.

This article builds upon and complements past surveys on the topic and related areas [33, 142, 146, 209, 214, 241]. In particular, the survey of Rios-Martinez et al. [209] specifically focuses on the sociological aspects underlying the development of social robot navigation frameworks. Charalampous et al. [33] focus on the perception challenges arising in social-navigation research, whereas in the survey of Thomaz et al. [241], the focus is on computational human-robot interaction more broadly. The survey of Lasota et al. [146] focuses on the aspect of safety for human-robot interaction broadly. Finally, the survey of Rudenko et al. [214] specifically focuses on the problem of human motion prediction, leaving out the challenges underlying embodiment on robot platforms or robot navigation.

Our survey is timely: The recent widespread interest in autonomous driving has placed social navigation into the radar of a wider audience. We aspire that this survey will enable newcomers to grasp the foundations of the field and inspire researchers already in the field to revisit important milestones. Our survey is inspired by and close in principle with the survey of Kruse et al. [142] from 2013. In fact, we are following up on their conclusions, calling for a “holistic theory of human-aware navigation.” While we are still far from a unified theory on the topic, we contribute: a unique bottom-up perspective, centered on the classification of open challenges along a set of major axes and a selection of proposed directions that will hopefully constitute a valuable resource for researchers and practitioners in the decade to come.

This article is organized as follows: Section 2 delves into the planning and control problems arising in social robot navigation applications, highlighting the coupling between prediction and planning in crowded human environments; Section 3 focuses on the behavioral aspects underlying the motion of robots in public spaces, ranging from the theory of proxemics to the importance of intention and formation recognition; Section 4 discusses methodologies of evaluation and points out fundamental issues in simulation and experimental validation; and finally, Section 5 summarizes our key insights and lists remaining open questions.

2 CORE PLANNING CHALLENGES OF SOCIAL NAVIGATION

The problem of planning safe, comfortable, and efficient robot motion within or next to human crowds is challenging. The dynamic nature of the environment demands the integration of models for human motion into algorithms for robot decision-making. By studying the literature of the past decades, we observe two main trends of the planning paradigms: approaches that *decouple prediction and planning* and approaches that *couple prediction and planning*. The former approaches, often referred to as “dynamic obstacle avoidance,” generate estimates of how human agents will be moving in the near future and plan robot motion that avoids intersecting with their expected paths. In other words, these approaches assume that human motion prediction is unaffected by what the robot will be doing next. In contrast, the latter approaches entangle human motion prediction and robot motion planning by recognizing that robot motion affects human

behavior and vice versa. These approaches effectively treat robot planning as a component of a *joint motion prediction* that includes estimates of how all agents will be moving. In the following paragraphs, we review representative papers on planning for social robot navigation, following the decomposition of coupled/decoupled approaches.

2.1 Decoupled Prediction and Planning

Among the decoupled approaches, a further classification can be made in terms of how human agents are modeled. Some approaches treat humans as dynamic, non-responsive obstacles; others incorporate interaction-agnostic models of uncertainty; finally, some approaches incorporate dynamic obstacle avoidance considerations into planning through the use of social objectives.

2.1.1 Humans as Dynamic, Non-responsive Obstacles. Early works were driven by system design for specific real-world applications, such as autonomous museum tour guidance. For instance, *RHINO* [26] and *MINERVA* [243] were some of the first known robotic systems to interact with human visitors in real-world museums. Deployed in Bonn, Germany, and Washington, D.C., in 1997 and 1998, respectively, these robots successfully served thousands of tour guide requests and navigated alongside museum visitors over several weeks. The demonstrated capabilities of these systems motivated additional efforts in the tour-guide domain such as Robox [227], Mobot [187], Rackham [45], and CiceRobot [35]. Finally, progress in the field was incentivized through competitions such as the AAI Mobile Robot Challenge, organized by the **American Association for Artificial Intelligence (AAAI)** [172, 176].

While these systems achieved important milestones in terms of autonomous navigation capabilities and serving guidance requests, they were agnostic to the fact that dynamic agents around the robot were humans. In fact, their underlying navigation frameworks [74, 179] treated humans as non-reactive obstacles without explicitly modeling human navigation, social engagement, or interaction with and between humans. Abstracting away these processes has generally proved to provide a practical solution in terms of collision avoidance, which has inspired a significant amount of research over the years [17, 72, 196]. However, the strong assumptions of those models tend to produce robot behaviors that negatively impact user experience.

Another prevalent paradigm in this area involves planning *around* future estimates of human motion to avoid hindering their motion. For example, Ziebart et al. [288] and Vasquez et al. [265] recover reward functions describing the motion of humans in a cluttered environment using **inverse reinforcement learning (IRL)** and use it to inform the costmap of a search-based planner. In a similar approach, Luber et al. [161] learn a set of dynamic navigation prototypes describing human motion and use them to inform the costmap of an A* planner. Henry et al. [107], also using IRL, focus on navigation in large-scale crowded scenes and demonstrate humanlike self-organization behaviors like lane formation. Brito et al. [22] propose a nonlinear model predictive control framework that generates motion avoiding approximated collision regions. These approaches indeed introduce social considerations into the planning process by having the robot either directly mimic human navigation traits or avoid humans by leveraging models of how humans tend to navigate in known spaces. However, they tend to ignore the effect that the robot's own motion can have on human motion itself. Failing to account for human reactions to robot behavior often results in robot motion that surprises the human, who in turn reacts unpredictably, contributing to a short oscillatory interaction, coined as the "reciprocal dance" [69].

2.1.2 Modeling Uncertainty. The practical issues observed during the first wave of robot deployments in crowded human environments motivated a shift in the navigation architecture. In particular, the second wave of work in the area focused on the development of methodologies for reasoning about the inherent uncertainty of the domain. For instance, Du Toit and Burdick [58]

presented a receding-horizon control framework that incorporated predictive uncertainty in the robot's decision-making. Thompson et al. [242] introduced a probabilistic model of human motion based on individual human intent inference, designed to assist in motion planning problems. Joseph et al. [115] proposed a Bayesian framework for reasoning about individual human motion patterns to inform a motion prediction pipeline. Similarly, Bennewitz et al. [18] extracted patterns of human motion in a crowded environment and derive a hidden Markov model to perform online human motion prediction. Unhelkar et al. [259] introduced a motion prediction framework that makes use of biomechanical features to anticipate human turning actions.

Despite the principled introduction of models for reasoning about uncertainty, these works tend to treat human agents as *individual non-interactive entities*. The lack of explicit coupling over the possible motion of human agents often results in an *uncertainty explosion*, as observed by Du Toit and Burdick [58]. Indeed, overapproximating the uncertainty over the unfolding crowd motion might lead the robot to conservatively mark all of its possible paths as unsafe. This might lead to the “freezing robot problem” [249]: The robot decides that the safest action to take is to stay in place, blocking or hindering the motion of surrounding humans.

2.2 Coupled Prediction and Planning

Artifacts such as the “freezing robot” [249] or the “reciprocal dance” [69] inspired a recent wave of approaches in the area of planning for social robot navigation that follow the paradigm of coupled planning and prediction. Researchers began developing mechanisms modeling the effects of robot motion on the motion of surrounding humans navigating in the same environment. Inspired by human navigation [279], many approaches relied on the development of models of **cooperative collision avoidance (CCA)**, capturing multiagent interactions in crowd domains. Within this body of work, we observe two main paradigms. Some approaches *explicitly* couple prediction and planning into the same computational process by jointly reasoning about the future states of the robot and surrounding humans. In contrast, *implicit* approaches incorporate interaction awareness into the design of objectives (e.g., heuristics, cost/reward functions) driving decision-making.

2.2.1 Explicit Approaches. By explicit, we refer to approaches that employ explicit predictions of human behavior, typically in the form of trajectories, probability distributions, or expected occupancy maps. Note that human motion prediction is a relevant and active field that has received much attention recently [2, 90, 178, 214, 270, 287]. However, within the scope of this survey, we focus on works that are primarily concerned with the navigation tasks. Thus, in the following paragraphs, we discuss literature that focuses on navigation that is informed by models of explicit prediction.

A body of work proposes interaction-aware navigation algorithms using behavior prediction models based on intent inference. Specifically, several works condition prediction and motion generation on estimates of other agents' goals acquired from past behavior observation. For instance, Bai et al. [9] and Reference [11] reason about the uncertainty of human intent by employing **partially observable Markov decision processes (POMDP)** and **mixed observability Markov decision processes (MOMDP)** frameworks. Some other works employ motion indicators to improve intent inference, such as anticipatory biomechanical turn indicators [259] or libraries of motion-level indicators [200]. Similarly, Park et al. [195] employ a data-driven model that enables a robot to avoid blockage due to human activities (e.g., engagement or obstruction) during navigation. Although beyond the scope of our survey, we mention that similar trends can be found in the autonomous driving literature, where behavior prediction is often conditioned on agents' goal estimates [42, 134, 155, 156, 206, 280, 286] or multiagent interaction modes [210, 239].

Trautman et al. [249], observing how cooperation shapes human navigation in pedestrian spaces, introduced *interacting Gaussian processes*, a mechanism explicitly accounting for human cooperation for human motion prediction by introducing interdependencies across individual Gaussian processes tracking the motion of pedestrians. Some works employ geometric or topological representations to model the coupling among the trajectories of multiple navigating agents. For instance, Mavrogiannis et al. employ topological representations such as braids [21, 166, 169] and topological invariants [20, 170, 171] to introduce a model of multiagent dynamics into the robot decision-making. Using these abstractions, they perform symbolic inference for multiagent motion prediction and contribute motion planners for legible and socially compliant motion generation.

Another body of work models multiagent collision avoidance in social robot navigation by leveraging tools from game theory. Turnwald and Wollherr [256] model the problem as a non-cooperative, static, multiplayer game and iteratively solve for Nash equilibria corresponding to collision-free strategies. Fridovich-Keil et al. [77] treat the problem as a multiplayer, general-sum differential game that they solve using an iterative approach resembling **iterative linear quadratic regulator (ILQR)**.

Finally, some works leverage insights about the structure of selected navigation contexts. For example, Chatterjee and Steinfeld [34] and Wang and Steinfeld [274] propose a “Gestalt” approach using frameworks that cluster navigating pedestrians into groups. Wang et al. [273] leverage this notion of grouping into a reactive control algorithm that generates group-aware robot motion in crowded scenes. Nanavati et al. [183] focus on leader-follower tasks involving a guiding robot and a following human and extract an empirical model of human following behavior that enables a robot to plan effective maneuvers leading humans to their destinations.

2.2.2 Implicit Approaches. By implicit, we refer to approaches that do not make use of explicit predictions of other agents’ behaviors (in the form of, e.g., trajectories, occupancy distributions) but rather leverage insights about the principles underlying human crowd navigation to guide robot decision-making. Often, such approaches leverage human data to inform the design of policies or objectives driving motion generation.

A class of works extracts social objectives of human navigation from demonstrations. The works of Kim and Pineau [131] and Kretschmar et al. [140] employ **inverse reinforcement learning (IRL)** as a technique to recover models of human navigation objectives and use them to infuse humanlike navigation traits to robot navigation. A number of works have proposed deep reinforcement learning models for prediction in crowd navigation domains. For instance, Chen et al. [39], Fan et al. [63] employ CADRL [40], a deep reinforcement learning framework for socially aware multiagent collision avoidance. Everett et al. [61] relaxes prior assumptions with an actor-critic variant [8], simultaneously learning policies and agent motion models. Furthermore, Tai et al. [237] employ generative adversarial imitation learning, trained on a dataset generated using the social force model [105]. Finally, Chen et al. [37] employ attention-based reinforcement learning to produce interaction-aware collision avoidance behaviors. Tsai and Oh [254] introduce a generative adversarial network architecture for social navigation where an agent is trained to navigate using multiple objectives including private intention and social compliance.

2.3 Open Problems and Directions for Future Work

Despite substantial and sustained interest in planning framework design for social robot navigation, several fundamental problems remain.

2.3.1 Intractability of Coupled Approaches. Much of the motivation for coupled prediction and planning stems from the inability of a robot to move through a dense crowd unless the robot accounts for human cooperation [248]. However, the “limit” nature of the freezing robot

problem—limit in the sense that it only applies in dense enough crowds—provides the impression that problems only arise after a density threshold is surpassed. This is not necessarily true: Even for one robot moving past one human, if the human is close (in velocity or position space) to the robot, then the agent density is high. Thus, for a robot that ignores coupling, a nearby human can cause overly evasive maneuvers (or stopping without reason) that can make navigation unsafe. This suggests that human-robot coupling (or, equivalently, coupled prediction and planning) is much more than just a “limit” phenomenon, but instead a core abstraction for social navigation architectures. Note that recent research in neighboring fields such as autonomous driving and human motion prediction is also dominated by approaches that strive to leverage models of coupled interaction [31, 152, 154, 167, 207, 210, 239].

Unfortunately, coupled models—where each agent’s movement is dependent on and influences every other agent’s movement—are, in general, computationally intractable. This complexity is easily visualized: Consider n_t agents, each with a discretized space of actions along 8 planar directions. For a prediction horizon of T timesteps, the system has $O(8^{n_t T})$ states. Even for simple cooperative collision avoidance objectives (such as described in Reference [250]), 10^6 Monte Carlo joint samples consistently produce unsafe and inefficient solutions [250]; this is hardly surprising: The prediction horizon is 3 seconds ($T = 20$) and $n_t = 7$, so the joint space is of size $8^{20 \times 7} = 8^{140}$. Importantly, this Monte Carlo scheme is not *random*—samples come from probability distributions governing individual agent intent. What this suggests instead is that reasoning over the *decoupled* space of agent intents will not recover a feasible *joint* solution. Although Trautman and Patel [250] provided an efficient means to recover local optima by optimizing according to the joint objective, the structure of the global optimality landscape was left as an open question, precisely because general tools to analyze non-convex objectives do not exist. Based on this result, we expect that joint objective functions for sampling-based motion planning (e.g., PRM [126], RRT [148], or RRT* [119]-based frameworks) will fail in much the same way, since there does not exist a method to draw samples from the joint objective—only from the individual agent models. Even if sampling-based motion planning increased performance by multiple orders of magnitude, this still would not be enough to overcome the complexity of this objective function.

Broadly, we expect—and it has been proven—that coupling in existing planning formalisms—e.g., interactive POMDPs [9, 99], coupled dynamical systems [275] (in the best case, social navigation is related to n-body classical physical simulation, which requires supercomputers [204]), game-theoretic approaches [52, 77], braid theoretic [169], or joint probabilistic inference [250]—is inherently non-convex (whether some objective function might reduce the computational complexity is an open and important research question). In general, for continuous-time Bayesian networks (similar to crowd navigation), complexity is NP-hard [47, 234]. For reinforcement-learning approaches, e.g., Chen et al. [37, 39], there is no reason to expect this complexity issue to disappear; rather, we expect that the complexity will be reflected in training burden. For learning-based approaches—e.g., imitation learning [237]—we would likewise expect data needs to reflect state space complexity.

Beyond the non-convexity of “simple” interacting agents (e.g., physical particles governed by gravitational equations), social navigation is also multi-objective (at the very least, we are typically concerned with balancing safety and efficiency). The tradeoff between safety and efficiency in coupled systems likely involves a Pareto front. That is, not only is this problem non-convex, but there are likely many global optima with the same objective value. This suggests that optimizing for safety and efficiency is an under-specification, and so additional constraints will be required. It also motivates elevating multi-objective non-convex optimization to the forefront of algorithmic social navigation research. We simply cannot provide basic guarantees that a coupled algorithm is even doing what we *intend* it to do (apart from whether or not our objective or our dataset appropriately reflects reality) until we can prove that (a) we are acting

according to a global optimum and (b) if a Pareto front exists, then the chosen optimum trades competing objective criteria in the desired manner.

2.3.2 Informing Planning with Human Models. A significant decision for any planning framework in social navigation involves the definition of a motion model for co-navigating humans. As discussed in Section 2.1.1, many works make the assumption that other agents move as dynamic obstacles with primitive or nonexistent models of intelligence. A notable example is the constant-velocity model, according to which an agent will move deterministically along their current direction of motion with a constant speed throughout the planning horizon. Under a sufficiently small replanning cycle, this technique may prove effective, especially in situations where humans tend to follow consistently goal-directed, compliant motion. Human motion in pedestrian domains is more nuanced. Pedestrians are not always cooperative [279]; they can be distracted (by, e.g., looking at their phone), they can be in a rush, paying little attention to the possible existence of a robot, or they might have different perceptions or preferences of personal space, and so on. Another prevalent example involves the assumption that all humans are identical, represented as disks of equal diameter. This is clearly not the case; people walk in different ways, depending on their height and size, whether they are completing additional tasks like carrying a suitcase, or whether they have a disability.

Enabling a robot to robustly navigate in so rich environments would require the incorporation of high-fidelity, high-granularity models of human motion and of human perceptions of robots into planning algorithms and the development of mechanisms for adaptation to the diversity of the crowd. While the human-robot interaction community has been working in parallel on understanding human behavior around robots, as discussed in Section 3, converting this understanding into practical models ready to be incorporated into planning algorithms is not trivial and requires further investigation.

2.3.3 Informing Planning with Context Models. The majority of prediction and planning approaches do not explicitly account for the context of the environments in which robots are deployed. Typically, it is assumed that the robot is deployed in an open-space environment, where there are no hard constraints on maneuvering due to obstacles. However, we are aware of the importance that the shape of a space (e.g., open space or narrow hallway), the timing (e.g., rush hour in a metro station or a museum or lunch time at a cafeteria), the semantic map (e.g., at airports there are often queues next to the gates), or just the occasion (e.g., a public lecture, a demonstration or a party) play on the type of crowd interactions. While different types of context such as environmental or task-specific have been studied to design general intelligent agent behavior [222], the notion of context is not typically integrated into the design of planning algorithms. Some works have looked at very specific selected contexts like grouping behavior [273], but the extension to different crowd formations requires further investigation. Finally, there is limited work in human motion prediction [216], which accounts for environment geometry constraints, but the extension to navigation has not been thoroughly explored. An important direction for future work involves the formalization of the notion of context in social navigation: defining what types of crowd conditions, behaviors, and settings the robot needs to be aware of during planning. Given such a formalism, it would be crucial to develop mechanisms for context classification that would enable a robot to adapt to the typically dynamic context of pedestrian scenes. Note that while similar problems might be encountered in the area of autonomous driving, the space of behavior in social navigation is so rich that it constitutes a dedicated investigation on its own. In driving, human motion can be assumed to mostly abide by a certain level of structure; in contrast, social navigation is practically rule-free, only governed by non-uniform social conventions and affected by complex and hard-to-model mechanisms connecting human perception and action.

3 CORE BEHAVIORAL CHALLENGES OF SOCIAL NAVIGATION

Section 2 focused on functional aspects of navigating within a crowded environment, motivated by safety and efficiency considerations. However, merely addressing human safety and robot task efficiency is not sufficient for producing socially competent robot behavior in crowded spaces. Humans value higher-level considerations such as proxemics and formations and are attuned to signals stemming from body posture, gestures, and eye gaze of others. As such, analyzing human behavior is not only important for better planning performance, but also for better human comfort and trust. In this section, we review a significant body of work from human-robot interaction focusing on understanding such mechanisms, often via mathematical expression conversions and human user studies, that aim to augment planning and prediction algorithms for social navigation.

We partition the works in three sections with increasing order of magnitude and social signal richness. The first section addresses human behavior on an individual level. The second section also addresses human behavior on an individual level but contains richer social signals. The third section addresses human behavior on a macro level and contains the highest-level social signals.

3.1 Proxemics

The notion of personal space from the theory of proxemics [96, 97] describes a space around a pedestrian that, if intruded, might cause the pedestrian to feel discomfort. Proxemics motivates the design of models and algorithms to enable robots to avoid getting too close to pedestrians. Early work focuses on simple models that consider each pedestrian individually. Hall [97] classifies personal space into concentric circular zones. Each zone represents a different level of affinity. Later, egg shape models were proposed by Hayduk [101] and Kirby [132] from the observation that pedestrians value frontal space more than their personal space in other directions. The parameters of the egg shape models are then experimentally determined by Barnaud et al. [14]. Later works believe that personal spaces are asymmetrical [82, 276] and dynamic [102]. Amaoka et al. [5] additionally considers gender and age when modeling personal spaces.

Regarding interaction with a robot, some works on proxemics focus on stationary subject approaching behavior. Takayama and Pantofaru [238] studied the spatial signals around humans in both human-approach-robot and robot-approach-human scenarios. They analyzed which factors from the environmental context can affect the humans' proxemics behavior. Torta et al. [247] designed an empirical expression from user study experiments to model human's proxemics zones when approached by a robot. Mead and Matarić [174] used a probability-based model to model social signals around stationary pedestrians when approached by a robot. These social signals are later used to analyze gesture and speech behavior [175].

In practice, proxemics signals are sometimes used as supplementary components to social navigation models. In a loose and simple form, they are manifested as potential fields around pedestrians. Social force-based models [105] define repulsive forces from other pedestrians. This can be viewed as one form of potential field-based proxemics modeling. In Svenstrup et al. [235], rapidly exploring random trees are combined with a potential field [130, 138] using parameters based on proxemics. Pradhan et al. [201] take a similar proxemics potential function-based approach.

Other works treat proxemics signals as hard boundaries around pedestrians that should be respected. Lam et al. [145] defined various 'sensitivity fields' with hard-boundaries that allow the robot to identify its social situation and execute the appropriate navigation strategy. Pacchierotti et al. [191] defined a lateral passing distance in scenarios where a robot and a pedestrian walk head-to-head towards each other. Pandey and Alami [193] used a two-level proxemics model to determine the reactive strategies for a robot to navigate around humans in which the robot should completely avoid these proxemics regions. In Bera et al. [19], the authors used data-derived parameters to define a dynamic rigid personal space and social space around pedestrians.

Last but not least, some works model proxemics rules as a term inside a cost function to be optimized. In Chung and Huang [44], the authors optimize an overall cost function that consists of several weighted cost terms to generate trajectories for the robot to follow. One of the cost terms is defined by proxemics-based rules. In Khambhaita and Alami [128], such an optimization is performed for both the robot and an interacting pedestrian reciprocally.

3.2 Intentions

Intentions often give rise to rich signals that offer hints to the future behavior of an agent. Literature from psychology and cognitive science [10, 49, 277] has shown that humans tend to interpret observed actions of others as goal-directed. Inspired by such findings, the human-robot interaction community has focused on engineering intent into robot motion, demonstrating improved teamwork [57], inference of robot capabilities [144, 272], and crucially lower planning effort [28] and smooth co-navigation [165]. While much of the intention information is often encoded in robot motion, alternative modalities like lights [13, 68], gestures [100], facial expressions, or gaze [1] may enhance the effectiveness of signaling. While these modalities are valuable in communicating intentions, they require robot designs that go beyond the general-purpose mobile robots that most researchers and practitioners have access to. In contrast, all mobile robots may leverage their base motion. For this reason, we focus on works that extract intentional signals from robot (base) motion.

There has been a variety of approaches towards enabling robots to communicate their intentions. Sisbot et al. [228] developed a series of cost functions that enable a cost-based planner to account for the accessibility, vision field, and human preferences of surrounding humans. Dragan et al. [56] defined legibility as the property of motion that enables an observer to quickly and confidently predict the goal of an actor. Lichtenthaler et al. [157] showed that legible motions increase the perceived safety by pedestrians. Kruse et al. [141] designed a human-behavior-imitating cost function to be optimized to produce legible motions for mobile robots. Knepper and Rus [135] distilled the concept of *civil inattention* observed by Goffman [86] into a multirobot path planner. Szafer et al. [236] designed a set of expressive curvatures for an aerial robot to follow to signal its navigation goals. Mavrogiannis et al. [165] used the physical quantity of angular momentum as a heuristic that signals intent over a passing side.

In addition to future behavior hints, intentions offer more implicit, context-level signals that can further aid pedestrian's interpretation of robot motions. In a shopping mall setting, Satake et al. [219], Satake et al. [220] studied strategies for robots to successfully approach shoppers and ask them whether they need assistance. The authors found that the robots needed to successfully signal the pedestrians that they are being approached or the pedestrians would miss the robot. Saerbeck and Bartneck [217] used various combinations of different accelerations and path curvatures to invoke emotional interpretations from pedestrians. Similarly, Knight and Simmons [136] produced expressive robot motions by drawing inspiration from Laban efforts. These expressive motions conveyed the robot's contextual state, such as being happy or in a hurry. Finally, Walker et al. [272] proposed a general framework for learning to map robot motion to a space of behavioral attributions like competence, curiosity, and brokenness.

3.3 Formations and Social Spaces

The spatial behavior of pedestrians and their encompassed social spaces contain rich social signals. These signals can be used to inform a robot that social areas are inappropriate to trespass and offer additional insights assisting with pedestrian motion tracking and prediction.

Much of the work on pedestrian spatial behavior is focusing on grouping. Grouping is thought by some as a synchronization process where a few individuals attune their behavior to that of the

people around them. Some [110, 111, 159] have investigated how robots can adjust their behavior to fit into human groups. From the perspective of navigation, researchers are more interested in groups rather than the grouping processes. Current works on groups can be divided into static and dynamic settings.

In static groups, people do not move and form interactions, such as conversation groups. Social space formed within static group formations should be respected by mobile robots. Vázquez et al. [266] and Vroon et al. [271] have studied how a robot may affect the spatial behavior of static groups by performing certain actions. A common tool used to analyze static groups is the F-formation Kendon [127], which was used by Kuzuoka et al. [143] to also analyze how a robot's action can affect static group spatial arrangements. Apart from studying a robot's effect on static group spatial arrangements, researchers are also interested in detecting these groups. Although sometimes not explicitly related to navigation, these static groups should be counted as impassable spaces and should be avoided during navigation. Again, some work has used F-formations to detect these groups [109, 224, 267]. Recent work has also used deep learning-based models to detect static groups from images and videos [3, 43]. In social navigation, static group space is also referred to as a type of interaction space [38, 268] or dynamic social zone [253] that guides the robot to respect static groups.

Dynamic groups are groups implicitly formed by moving pedestrians. Similar to static groups, social spaces are present within dynamic groups and are best not to be interrupted. Similar to static groups, Fiore et al. [70] and Shiomi et al. [225] have studied how robots can affect dynamic group spatial behavior, as commonly seen in museum tour guide robot settings. Garrell and Sanfeliu [79] also studied appropriate robot spatial behavior when accompanying pedestrians. Apart from spatial behavior, most of the focus when studying dynamic groups is on detection or tracking. Detection and tracking of groups are examined from both exo-centric and ego-centric perspectives, where exo-centric assumes known pedestrian properties, while ego-centric introduces additional challenges associated with perception.

From exo-centric perspective, common approaches include probabilistic progression-based methods [16, 32, 81, 199, 283, 284], graph-based approaches [30, 67, 129, 186], clustering methods [80, 231], and social force model-based approaches [149, 173, 260]. From ego-centric perspectives, grouping becomes more difficult, because sensor errors and pedestrian occlusions are present. Common models use Multiple Hypothesis Tracking models [147, 160, 180], which focus on data association between group selection hypothesis on two consecutive frames, and clustering-based approaches [34, 240] on features abstracted from sensor inputs.

In a navigation setting, taking dynamic groups into consideration has shown improvements in performances. Wang et al. [273] and Katyal et al. [125] have shown that treating groups as obstacle representations significantly benefits the safety of the robot planner. Sathyamoorthy et al. [221] show that by allowing intrusions of groups with weak cohesion, the robot can unstick itself from "the freezing robot" situations and achieve better efficiency.

The literature presented so far studies groups in a generic multi-agent setting. Other researchers focus on specific types of grouping behavior or actions, sometimes with specific application scenarios in mind. Researchers studied how a robot should follow pedestrians [85, 89, 116, 285] and conversely, how a robot should guide pedestrians [66, 78, 183, 192]. Others focus on how a robot should stay in group formations [4, 50, 163]. Last, some researchers focus on predicting group splits and merges [274].

Researchers have also focused on meta-level formations in high-density crowds. Henry et al. [107] looked at lane formations in narrow corridors when crowded pedestrians walk in opposing directions. Van Den Berg et al. [264] observed such behavior in simulation from their navigation

model. Cheriadat and Radke [41] aimed to detect dominant pedestrian flows in surveillance videos in dense crowds.

3.4 Open Problems and Directions for Future Work

One of the barriers to designing socially competent robot navigation systems is the limited *understanding* of human pedestrian behavior. Principled knowledge is important for defining social rules that guide the robot's social navigation strategy. How do we model personal space and how well are the robots respecting the space? How should a robot communicate its navigation intent to nearby pedestrians? How should a robot identify and negotiate interaction space formed by several pedestrians? With principled answers to these questions, the robot can incorporate these behavioral guidelines either as modules to be explicitly plugged into a navigation framework or as cost terms to implicitly encourage the robot to follow the guidelines. Following social behavior guidelines defined by this principled knowledge allows a robot to abide by human navigation expectations. To a greater extent, this principled knowledge can help establish metrics that measure the socialness of robot's navigation behavior. As mentioned in Section 4.1, most of the metrics present in current literature focus on efficiency and safety of the robot, where socialness is either implied from safety or abstracted from human participants' overall ratings during user studies. Principled knowledge in human behavior can help with the design of measurable socialness metrics, such as how often the robot intruded personal/interaction space or how legible the robot's trajectories are in communicating its intended navigation goal.

Nevertheless, obtaining this principled knowledge presents several significant challenges. At the current stage, research in pedestrian behavior understanding takes a trial-and-error-based approach. For most of the literature mentioned above, researchers design a model for a specific aspect of pedestrian behavior and often show in a small-scale lab study that the model shows promise. While the lab studies and validation methods used to establish scientific significance of these models are sufficient, lab studies are often heavily controlled and lack authenticity when compared to real world scenarios. Section 4.2.4.2. provides more discussion on lab studies in social navigation model evaluations. To turn these models into principled knowledge, they need to be tested extensively in more generalized, real-world environments. To further expand on this, pedestrian environments tend to feature a much richer selection of human activity, such as distracted pedestrians or pedestrians in a hurry as mentioned in Section 2.3.2. Enabling a behavior model to robustly comply in such rich environments would require the model to be able to adjust its behavior based on various types of human pedestrian behavior. Additionally, for the study of a particular pedestrian behavior to become principled knowledge, the behavior model's significance needs to be tested in a practical setting. It is also often unclear how much the inclusion of such particular behavioral aspects benefit social navigation. For example, the inclusion of personal space geometry or executing legible robot motions are known to improve human comfort, but when *integrated into a social navigation framework*, how much it contributes to the human's perception of the robot's socialness has yet to be measured. Future work is recommended on integrating social navigation planners with these behavioral aspects to analyze their practical effectiveness.

4 CORE EVALUATION CHALLENGES OF SOCIAL NAVIGATION

Social navigation is a sequential decision-making problem with multiple objectives. In the navigation problem without the social interaction part, the primary objective of an agent is to reach their destination. Considering this objective only, an agent's navigation performance can be evaluated based on how successful it is at arriving at its destination, e.g., using the physical distance between an agent's position and a destination. Among the solutions that have reached a destination, it is generally subjective to rank alternatives according to multiple preferences or soft constraints.

Table 1. Evaluation Metrics for Social Navigation

Metric	Description
Arrival rate	Percentage of agents that successfully reach the goal
Path length	The distance traveled by an agent
Collision rate	Percentage of cases with at least 1 collision
Failure rate	Percentage of agents that fail to reach destinations
Average time to goal	Average time to reach goals
Social score	Combined score of success rate and comfort
Average acceleration	Acceleration averaged by time
Average energy	Integral of squared velocity of an agent, averaged by timestep
Path irregularity	Total amount of rotation beyond a straight line path
Path efficiency	Ratio between straight line path and actual path
Time spent per unit path length	Average time spent
Topological complexity	Amount of entanglement among agents' trajectories
Speed efficiency	Ratio of nominal and actual speed

Table 2. Dataset Comparison According to the Metrics in Table 1 Relevant to Recorded Data

Metric	ETH	HOTEL	UNIV	ZARA1	ZARA2
Avg. acceleration (m^2/sec)	0.092	0.044	0.027	0.019	0.023
Avg. energy (m^2/sec^2)	6.296	1.765	0.747	1.445	1.473
Path irregularity ($degree/m$)	0.914	3.973	3.708	2.202	4.702
Path efficiency (%)	98.3	93.3	91.2	98.1	96.9
Number of agents	360	389	967	148	204
Avg. Agent Duration (sec)	5.63	5.80	15.04	13.53	18.66
Avg. Pedestrian Density ($./m^2$)	0.313	0.312	0.398	0.354	0.384
Max. Pedestrian Density ($./m^2$)	0.583	0.750	0.750	0.625	0.75
Min. Pedestrian Density ($./m^2$)	0.250	0.250	0.250	0.250	0.25

Social compliance, for example, is a soft constraint that is comprised of many aspects including safety, comfort, naturalness, and other societal preferences. The quality of social navigation performance can thus be evaluated differently, depending on where the focus is among those potentially competing objectives.

In this section, we first review the metrics used to measure various facets of the social navigation performance. Next, we describe three evaluation methodologies: datasets, simulations, and real-world experiments. Finally, we discuss remaining challenges on how to measure success in social navigation.

4.1 Metrics

The metrics discussed in this section are for evaluating the performance of a single agent. We note that some of these metrics in aggregation can also be used to evaluate a population performance. For instance, such statistics collected from the human pedestrian datasets can be used as the references for measuring how closely an algorithmic population, i.e., an algorithm in self-play, resembles human pedestrians in their navigation behavior.

The metrics described in subsections are summarized in Table 1. The statistics of five commonly used pedestrian trajectory datasets according to a subset of the metrics relevant to recorded dataset are shown in Table 2.

4.1.1 Navigation Success Rate. The arrival rate measures how often an agent reaches its goal, which is a primary metric in general robot navigation [188] and has also been used in social navigation [254]. In general, some assumptions are added to measure the arrival rate. For instance, a time constraint can be added to enforce an agent to reach its destination within a certain deadline. In the case of datasets, because the true destinations of pedestrians in a pedestrian dataset are generally unknown, people's destinations need to be assumed. For instance, the location from which an agent exits the scene can be considered as the agent's final destination for the purpose of computing the arrival rate, where the locations can be represented in either a discrete [269] or continuous space [254].

The arrival rate depends not only on an agent's navigation algorithm but also on its environmental context such as the crowd density. Combined with the average speed of agents, the arrival rate can be used to measure the level of congestion of an environment.

4.1.2 Path Efficiency and Optimality. Whereas the success (or arrival) rate evaluates the performance based on the final outcome, the path-efficiency metrics measure the quality of the trajectories. For instance, the length of a path [248] and the time needed for an agent to travel 6 meters [249] have been used to measure the efficiency given a path.

Mavrogiannis et al. [168] proposed a set of metrics for measuring various aspects of trajectory quality. The average acceleration of an agent and the average energy (integral of squared velocity over a trajectory segment) measure subtle changes in pedestrian motions. The path irregularity metric (amount of unnecessary rotation), modified from the work of Guzzi et al. [94], and the path efficiency metric (ratio between Euclidean distance to goal and the actual traveled distance) measure the level of disturbance in pedestrian behavior. Finally, the time spent per unit length, the reciprocal of average speed, and the speed efficiency measure different aspects of efficiency of a trajectory.

The ratio of nominal and actual speed has also been proposed in Reference [76].

4.1.3 Safety/Collision Avoidance. The safety performance can be measured in terms of the number of constraint violations; for instance, the number of collisions [75]. In addition to hard constraint violations, continuous metrics have been proposed to define a robot's closest distance to a human, known as the safety margin [248] or the minimum distance [168].

4.1.4 Behavioral Naturalness. Naturalness is not directly connected to a planning performance, but the concept represents a fitness score for a robot to blend in a human environment. To evaluate how well an agent's behavior resembles a human's, a dataset of recorded human pedestrians can be used. It is a common method to evaluate the naturalness of a robot navigation in terms of the distance between an agent's plans and the actual trajectories taken in the recordings. Two distance metrics that are prominent in existing works are the average displacement error and the final displacement error. Naturally, these metrics are also predominant in the pedestrian prediction literature [214].

Average Displacement Error (ADE), introduced in Reference [198], is the average of Euclidean differences between temporally aligned points in the two paths given by a navigation algorithm and the ground truth path from the dataset, respectively. ADE has been used in References [281] and [98].

Final Displacement Error (FDE) proposed in Reference [2] measures the displacement error only at the final time step. In this regard, it is related to the arrival rate, although the notion of a goal is only implied here.

The prediction accuracy [19] is the percentage of successful predictions in all attempted predictions. A success in this context is typically defined based on a threshold value and another

similarity metric such as the average or maximum displacement error, such that a prediction is regarded as successful if its similarity to the ground truth value is higher than the threshold.

Dynamic Time Warping (DTW) used in Reference [98] finds the optimal way to warp prediction to a ground truth path and considers the cost of performing such a warping as the difference between prediction and truth. In the similar vein, the Hausdorff distance or Fréchet distance can also be considered.

Tsai and Oh [254] proposed the comfort rate following the theory of social force [105]. The original social score definition combines the safety score with the navigation score, e.g., whether an agent has reached the destination. The comfort score only takes the social compliance term into consideration based on the minimum required social distancing constraint.

Mavrogiannis et al. [168, 170] employed the topological complexity index [60] as a measure that quantifies the intensity of *mixing* observed in the motion of multiple navigating agents. The higher the complexity, the more agents directly confront each other as they navigate towards their destinations. Topological complexity was shown to be correlated with a notion of *Legibility* [57] in multiagent navigation scenarios.

4.2 Evaluation Methodologies

In this subsection, we discuss the methodologies used in literature and their pros and cons in evaluating various facets of social navigation.

4.2.1 Datasets. Using a dataset, evaluation is limited to measuring the difference between predictions and the recorded behaviors using a metric such as displacement errors. There are several issues here: (1) recorded behaviors may have been conditioned to a certain context; (2) alternative behaviors that are different from the recording but are equally good can be unfairly penalized; (3) the majority of existing datasets do not have the notion of destinations for the agents in a scene, making it less suitable for a navigation task; and (4) recorded behaviors rarely include hard constraint violation such as collision, and these rare events need to be engineered/trained based on domain knowledge.

There are two datasets widely used in both social robot navigation and human motion prediction literature: ETH [198] and UCY [150]. The ETH dataset contains trajectory data from two scenes, namely, *ETH* and *Hotel*, while the UCY dataset contains *Univ*, *Zara1*, and *Zara2*. *ETH* and *Hotel* have more linear trajectories and lower pedestrian density when compared to others. *Zara1* and *Zara2* have relatively more complex trajectories and higher pedestrian density. *Univ* has the most non-linear trajectories and the highest density out of these five scenes.

Using pedestrian datasets for evaluation is favored due to two reasons. First, although the recorded behaviors are not fully representative of human pedestrians, the datasets still contain the trajectories of humans navigating in a real environment. Second, it is practically much more convenient to evaluate on a dataset when compared to testing in a real human environment. The dataset-based evaluation, however, presents several limitations. These metrics are only applicable when the evaluation is done by comparing navigation algorithm outputs to dataset recordings. This way of evaluating navigation algorithms carries strong prior assumption about the optimality of human behavior recorded in dataset, while in reality paths chosen by human pedestrians might not be optimal or there is likely more than one optimal path for the same scene. Forcing navigation algorithms to give the exact same path as recorded ground truth may not be wise.

4.2.2 Using Human Datasets for Evaluation. One possible way to derive a navigation protocol can be as the one proposed by Reference [249]. This approach requires identifying a pedestrian within a frame, extracting the start and end position of that pedestrian, removing that pedestrian from the observation dataset of the navigation the algorithm, and finally to provide the start and

end positions of the removed pedestrian and the current and previous positions of the remaining agents to the navigation algorithm. In this way, it is possible to assure that one path through the crowd exists. This testing protocol provides us with actual human performance in the same situation as encountered by the navigation algorithm. To determine the safety threshold, Reference [249] computes the shortest distance that any two humans in the ETH dataset ever came to each other.

4.2.3 Simulation. Real-life experiments are usually expensive in both time and resources. Crowd simulators can be easily customized to support new tasks and environments in a way that would not be otherwise possible in the real world. They are essential tools to evaluate social navigation algorithms and train data-driven or reinforcement learning models [37, 39]. As Fraichard and Levesy [76] point out, existing crowd simulators rely on unrealistic assumptions that make them far from being a high-fidelity representation of the real world. They assume that all agents obey the same law (*homogeneity*) and are aware of the state of all other neighboring agents (*omniscience*). Additionally, in simulations, virtual agents are allowed to collide as long as these collisions do not impact visual realism (*safety*). Nevertheless, once these assumptions are relaxed, we run into safety issues. It is impossible to guarantee that a policy trained in such simulated environments and then transferred into the real world will not cause the robot to collide with other human agents.

4.2.3.1 Models. The literature proposes several techniques to model the crowd's dynamics, where pedestrian simulation engines are drawn from work in crowd analysis dynamics [103] or computer graphics [59, 123, 261]. However, current state-of-the-art does not offer a comprehensive evaluation approach for each model. Existing works mainly focus on simulating crowd behaviors in a critical context, such as the circular aggregation or flow of pedestrians around a narrow passage and lanes' formation between two or more pedestrians' flows [92, 104, 105, 190, 194] or on designing more artificial scenarios to test interactions between agents [91, 190, 263].

Although it is not common to observe these kinds of scenarios in real life, they provide the desired situations to evaluate collision avoidance algorithms in progressively changing density conditions. A crowd simulator must achieve sufficient generality and reach high complexity and randomness to avoid overfitting and guarantee the policy's convergence during training. There are three main approaches to model the crowd dynamics, i.e., microscopic, macroscopic, and mesoscopic models. See Table 3 for a taxonomy of widely employed methods.

The microscopic model attempts to study crowd behavior by modeling each individual's action through force-based or velocity-based models. The **Social Force Model (SFM)** [106] is one of the most prevalent models in this category. It is a force-based method, simulating crowd dynamics through interactions between agents and obstacles, where both socio-psychological and physical forces influence the individual decision-making process. It was first introduced to study the crowd dynamics in a situation of escape panic. Other works build upon the SFM algorithms by introducing a mobile grid to make each agent adjust its preferred velocity [215], resulting in more smooth trajectories, or adding evasive forces to predict future collisions [120] and interactions, and reduce particle oscillations [112]. Under the velocity-based modeling approaches, we can find techniques such as **Velocity Obstacle (VO)** [71], **Reciprocal Velocity Obstacle (RVO)** [263], **Optimal Reciprocal Collision Avoidance (ORCA)** [261], and **Hybrid Reciprocal Velocity Obstacle (HRVO)** [230]. These algorithms generate collision-free trajectories by considering combining local interactions and the neighbor's information into the decision-making process.

The macroscopic approach models the crowd dynamics at a large-scale. The crowd is represented as a continuous entity, where potential fields and fluid dynamics models govern the pedestrian motion. This modeling approach works very well to describe high-density situations, where it is possible to recognize the groups' appropriate behavior. However, it does not consider

Table 3. A Taxonomy of Crowd Models

Model	Microscopic	Macroscopic	Mesoscopic
Helbing and Molnár [105]	✓		
Saboia and Goldenstein [215]	✓		
Karamouzas et al. [121]	✓		
Fiorini and Shiller [71]	✓		
Van Den Berg et al. [264]	✓		
Snape et al. [229]	✓		
Hughes [108]	✓		
Treuille et al. [251]		✓	
Jiang et al. [113]		✓	
Narain et al. [185]		✓	
Patil et al. [197]		✓	
Dion [54]			✓
Lewin [151]			✓
Karamouzas et al. [121]			✓
Ulicny et al. [258]			✓
Yersin et al. [282]			✓
Guy et al. [93]			✓
Sakuma et al. [218]			✓
Wiggins [278]			✓
Eysenck [62]			✓

individual-level interactions between agents and the environment. Under this category, there are mainly three ways to design pedestrian flow dynamics, i.e., continuum theory, aggregate dynamics, and potential field. The continuum model [108, 113] describes the pedestrian flow dynamics through the use of the continuum theory. The crowd is a flow characterized by density and speed. This way of modeling crowd behavior, however, is more designed for large groups with common goals. The aggregate dynamics model [185] uses a fluid dynamical model to represent the crowd flow and discrete agents or continuous systems to describe the crowd density. The potential-field method [197] aims to model one or more agent actions by integrating guidance fields designed by a user to generate smooth, collision-free navigation fields.

Recent approaches investigate mesoscopic models to design group simulations rather than single or large-scale models. This method includes dynamic group behaviors [54, 122, 151] (i.e., the social relationships among individuals), which uses dynamic analysis to obtain general rules of group phenomena, interactive group formation [258, 282], and social-psychological crowds [62, 93, 218, 278], used for emergency evacuation and parades. The latter model considers personality traits and emotion contagion theories and how they affect human behaviors into the modeling framework.

These methods describe different ways to model crowd behaviors. The choice between the three categories is based on personal purpose, e.g., interest in modelling the crowd as a single entity or as a group. This section aims to give an overview of the current state-of-the-art. However, there is no way, as of today, to evaluate how each of these approaches compares to the real-world crowd behavior. The lack of a good amount of real-world pedestrian dataset does not make it possible to run a sufficient statistical analysis for their evaluation.

4.2.3.2. Software. This section reviews some relevant crowd simulator software available as research platforms for crowd behavioral studies.

PedSim [84], *OpenSteer* [205], *Menge* [51], *Continuum* [251], both written in C++, are the most common tools used to simulate crowd scenarios. *PedSim* [84] is a C++ library based on the social-force-based steering model. One of this method's main drawbacks is that it has no global planning ability, neither high-level behaviors and goals. *OpenSteer* [205] is a C++ application used to explore steering behaviors for vehicles but not specifically designed to simulate human pedestrians (i.e., the human agent derives from a "vehicle" class). Both these tools, however, have limited capabilities. They do not account for different human behaviors and have no global planning capability. *Menge* [51] is a C++ modular framework with the ability to use runtime plug-ins to change the pedestrian model, global navigation algorithms, and the way we can visualize the simulation. Fluid dynamical models have inspired other approaches to simulate crowds. The *Continuum Crowd* [251] is a tool based on continuum dynamics and dynamic potential field. This software is based on the continuum model and represents pedestrians as a continuous density field where partial differential equations describe their dynamics. As a drawback, none of this software considers the robot as part of the crowd scenario.

PedSim-Ros builds upon Reference [189] to include a robot agent in the simulation environment. *Webots*, **ROS (Robot Operating System)** [202], *Gazebo* [137], and *Stage* [83] are other frameworks used to represent virtual environments, where each agent is defined by its geometry and dynamics and navigation strategy. These tools, however, lack in both graphical and physical realism. *MengeROS* [6] is an open-source simulator that can simulate a variety of human crowd behavior (force-based, velocity-based) and integrate a robot in the scene.

CrowdSim [29] is a rule-based crowd simulation software developed to simulate pedestrian behaviors during evacuations, along with human comfort and safety. It is an add-on for *Blender* [73], and it is used primarily for visual effects. This software offers the possibility to generate random test data with high variability and controlled complexity. *UCrowds* [257] is a simulator available as 3D Plugin for *Unity* [95]. Its primary objective is to simulate the movements of large masses of people under a variety of circumstances. It offers faster and more accurate simulations and the possibility of producing dynamic adjustments, making it very appealing for crowd behavioral studies.

Golaem Crowd [88], *Massive* [164], and *Miarmy* [15] are photo-realistic simulators available as a plug-in for *AutoDesk Maya* [7]. They are explicitly used for generating special effects and visualization, but they are not suitable for crowd behavioral studies.

Last, an experimental platform, namely, *SEAN-EP* [255], has been introduced as a high visual fidelity and open-source system social navigation simulation platform, built on top of *Unity*, that evaluates and gathers feedback on navigation algorithms.

4.2.4 Experimental Evaluation. A crucial part of evaluation in social robot navigation research involves testing a navigation framework in close proximity with humans. In this section, we classify existing work into categories, depending on the experimental methodology followed for validation. Broadly, we observe that researchers have followed three main approaches: (a) experimental demonstrations presenting proof-of-concept prototypes; (b) lab studies under controlled settings; and (c) field studies in public environments. Table 4 presents a taxonomy of existing works based on the validation methodology followed.

4.2.4.1 Experimental Demonstrations. An experimental demonstration in the presence of humans often represents the first real-world step following a simulated case study for social robot navigation research. This type of evaluation often features limited control over experimental variables and limited reporting of qualitative or quantitative performance measures. For example, Bennewitz et al. [18] report a series of 10 experiments involving robot navigation in a hallway of limited crowd density. Sisbot et al. [228] describe a series of navigation interactions between a robot running their

Table 4. A Taxonomy of Studies in Algorithmic Social Robot Navigation

Study	Demo	Lab	Field
Burgard et al. [26]			✓
Thrun et al. [243]			✓
Bennewitz et al. [18]			✓
Pacchierotti et al. [191]		✓	
Sisbot et al. [228]	✓		
Ziebart et al. [288]		✓	
Foka and Trahanias [72]			✓
Kirby [132]		✓	
Park et al. [196]	✓		
Kruse et al. [141]	✓		
Bandyopadhyay et al. [11]	✓		
Luber et al. [161]		✓	
Shiomi et al. [226]	✓		
Kato et al. [124]			✓
Kim and Pineau [131]			✓
Trautman et al. [249]			✓
Kretzschmar et al. [140]	✓		
Chen et al. [39]	✓		
Truong and Ngo [252]		✓	
Mavrogiannis et al. [165]		✓	
Lo et al. [158]		✓	
Tsai and Oh [254]		✓	

proposed navigation algorithm and a human in a lab environment. Park et al. [196] test their control framework on a robotic wheelchair inside a corridor of an academic building and report a set of successful collision-avoidance encounters. Kretzschmar et al. [140] also deploy their model on a robotic wheelchair documenting a set of experiments in a narrow hallway under controlled settings. Chen et al. [39] document an experimental demo, involving a robot navigation experiment in a crowded area of an academic building.

4.2.4.2. Lab Studies. Lab studies provide the advantage of the ability to control experimental variables that play a significant role during execution. For this reason, lab studies represent an important step for design and evaluation of social robot navigation research. Pacchierotti et al. [191] test their control framework with a study involving an autonomous robot navigating next to human subjects in a corridor under controlled settings. They present their findings from the interactions of 10 participants with a robot exhibiting different navigation strategies corresponding to different passing distance. Kirby et al. [133] and Kirby [132] present a user study involving 27 human subjects navigating alongside a robot in an academic hallway. Kruse et al. [141] document a series of interactions between a mobile robot and a human subject in a lab study involving 10 participants. Truong and Ngo [252] document a series of interactions between a robot running their planning algorithm and human participants in a lab environment. Our past work [168] featured an experimental validation of our planning framework [170] in a lab experiment involving interactions between a navigating robot and three human subjects at a time under dense navigation settings, yielding a total sample of 105 participants. Lo et al. [158] evaluate a series of

robot collision-avoidance strategies on a self-balancing mobile robot in a lab study with 98 human subjects. Tsai and Oh [254] evaluated their approach on a physical robot in a controlled outdoor environment where they tested robot behavior against various settings of pedestrians. They used two types of human evaluations based on the feedback from the evaluators observing the scene from a third-person view and the participants (pedestrians).

4.2.4.3. Field Studies. Another common approach involves deploying robots in the wild in public environments. Burgard et al. [26] deployed the RHINO guide robot in the Museum of Bonn, in Germany, in 1997 and documented thousands of in-person and virtual interactions between the robot and visitors over 47 hours of runtime spanning 6 days. Thrun et al. [243] deployed MINERVA, the second-generation of robot tour guides at a Smithsonian Museum in Washington, D.C., in 1998 and also documented thousands of interactions with visitors for two weeks. Both studies documented statistics related to performance, collision avoidance, and visitors' impressions of the robot.

Foka and Trahanias [72] report logs and performance aspects upon running their robot for 70 hours in an indoor academic building. Shiomi et al. [226] test their navigation framework through a 4-hour field study in a shopping mall. Trautman et al. [249] test their navigation framework on a robot in a field study comprising 488 robot runs in a crowded cafeteria. Kato et al. [124] test their approach on a humanlike robot employee in a crowded mall and record interactions with 130 people. Kim and Pineau [131] evaluate their framework on a robotic wheelchair in a crowded hallway over 10 field runs.

4.3 Open Problems and Directions for Future Work

In this section, we highlight a series of problems related to existing practices on the evaluation of social robot navigation frameworks and offer suggestions for improvement.

4.3.1 Limitations of Existing Simulation Practices. Simulation is a valuable tool for prototyping and testing social robot navigation frameworks in controlled settings. However, simulating crowd navigation is a challenging task on its own, as it requires adopting assumptions that inevitably abstract away too much out of the rich interactions that naturally arise in real-world situations. Therefore, it is crucial to carefully construct an experimental setup that can provide meaningful insights. This requires determining the models governing the behavior of simulated humans, the specific scenarios in which we evaluate an algorithm, the metrics with respect to which we evaluate them, and more.

In the absence of a uniformly agreed upon evaluation standard, existing evaluation practices found in the literature often involve strong and unrealistic assumptions that prohibit the extraction of meaningful insights. For instance, the majority of recent literature [36, 37, 61, 153] relies on ORCA [262] as the main behavioral engine governing simulated humans for both training and testing of learning algorithms. Reference [37], for example, uses ORCA simulation environment as demonstrator to initialize the Value Network for the navigation algorithms. Then, the network keeps being trained on self-play distributions generated from a network initialized with ORCA agents behavior. While this choice is motivated, it comes with serious shortcomings. We uncover these shortcomings with a simulation study. Our experimental setup session aims to provide an additional overview of what happens when the Value Network in Reference [37] is being initialized with different agents model. The objective of our study aims to show how different pedestrian models can change the outcome of current-state-of-the art algorithms and to reason about the necessity of gathering more pedestrian datasets to develop realistic simulation models for better training and testing evaluations.

4.3.1.1. Experimental Setup. We consider an experimental setup in which a robot navigates within a crowd of simulated human agents in a circular workspace with a radius of $3m$. All agents (the robot and humans), represented as disks of $0.3m$ radius, are initially placed at random starting locations along the boundary of the workspace. Each agent is tasked with navigating from its starting location to the antipodal point on the circle. We consider two main environments¹—an ORCA [261] environment in which all human agents navigate by running the ORCA algorithm, and an SFM environment [106] in which all human agents navigate by running the SFM algorithm. In parallel, the robot navigates by running one of the following algorithms:

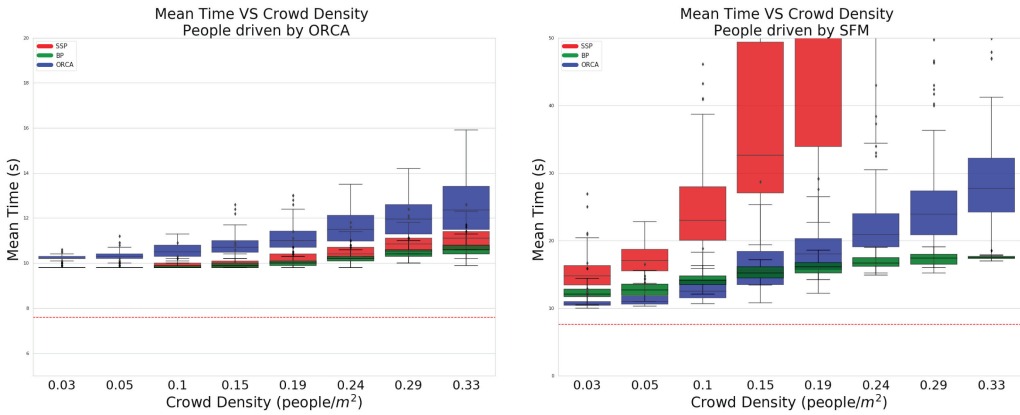
- **Simple Social Planner (SSP):** A robot moving straight toward the goal and stopping when getting closer than $0.2m$ to people.
- **Blind Planner (BP):** A robot with no access to any sensory data.
- **Optimal Reciprocal Collision Avoidance (ORCA) [262]:** A robot implementing the reciprocal collision avoidance algorithm given by Reference [263].
- **Socially Aware Reinforcement Learning (SARL) [39]:** A robot that learns human-robot interactions with respect to their future states through an attentive pooling mechanism over pairwise interaction features.
- **Collision Avoidance Deep Reinforcement Learning (CADRL) [40]:** A robot that moves following a policy that tries to minimize the expected time to the goal subject to collision avoidance constraints.
- **Relational Graph Learning (RGL) [36]:** A robot that models the crowd as a spatial graph where agents are the nodes and edges represents human-robot interactions and uses a constant velocity model algorithm to predict agents position at the next time step.
- **Model Predictive Relational Graph Learning (MPRGLP) [36]:** A robot that models the crowd as a spatial graph where the nodes and the edges represent the agents and the human-robot interactions, respectively, and that uses a constant velocity model algorithm to predict the positions of the agents at the next time step.

SARL, CADRL, RGL, and MPRGL have been trained using first imitation learning with collected experience from a demonstrator ORCA or SFM policy to initialize the model, and then RL to refine the policy. We use the subscript “SOCIAL” to differentiate the algorithms trained using a demonstrator SFM from the ones trained using demonstrator ORCA. These algorithms have been trained for 10,000 epochs using Adam Optimizer, a learning rate of 0.001, and a discount factor of 0.9. The robot moves following holonomic kinematics, and the action space consists of 80 discrete actions: 5 speeds exponentially spaced between $(0, v_{pref}]$ (where the preferred velocity is set to $1m/s$) and 16 headings spaced between $[0, 2\pi)$.

We consider eight different density levels, corresponding to scenarios involving 5, 7, 14, 21, 28, 35, 42, and 49 simulated human agents, respectively. For each density level, we generate 100 random scenarios corresponding to distinct initial arrangements of agents along the boundary of the workspace. We examine the performance of the robot measured with respect to: (a) average time to reach the goal and (b) average number of collisions over the 100 trials. We then present a series of plots that make the case for (a) the insufficiency of simulators commonly used for evaluating social navigation algorithms and (b) the suboptimality of state-of-the-art social navigation frameworks.

4.3.1.2. Insufficiency of Simulation Environments. In particular, Figure 1 reports the average time to reach the goal for SSP, ORCA, and BP as density increases in the ORCA and SFM environments,

¹We use the implementation and parameters of Chen et al. [37], hosted at <https://github.com/vita-epfl/CrowdNav>. Note that in contrast to their default implementation, as soon as the agent reaches its goal, it is assigned a new target, randomly picked on the opposite center-symmetric side. We added this level of complexity, such that the simulated environment keeps remaining dynamic for the whole time the robot needs to reach his destination.



(a) Average time to reach the goal on ORCA. The SSP or BP significantly outperforms an ORCA-driven robot in an ORCA environment. ORCA algorithm chooses velocity over paths, and so lacks a mechanism to represent the general navigation task; meanwhile, a naive planner such as SSP preserves this geometry information for path planning, achieving greater performance.

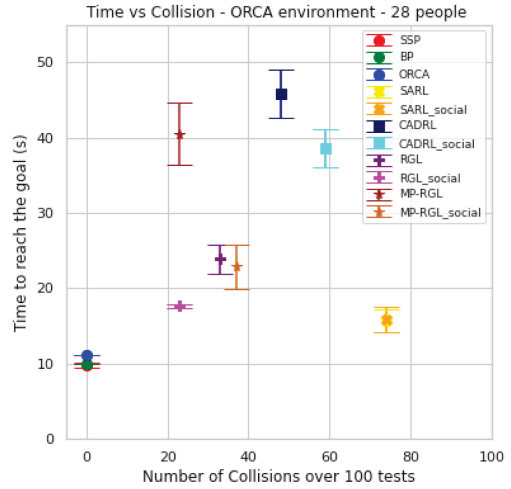
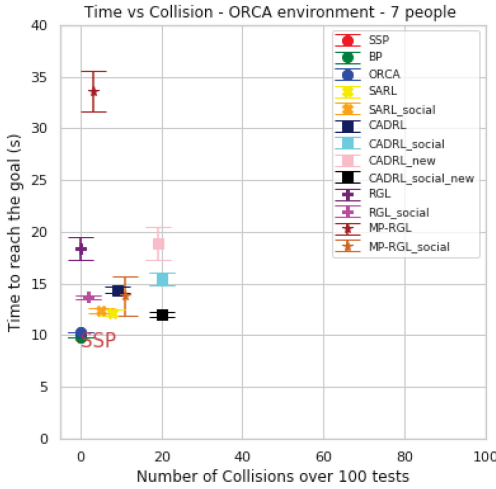
(b) Average time to reach the goal on SFM simulator. In this scenario, the SSP underperforms compared to the other methods. It fails to reach the goal when crowd density over 0.2 people/m^2 while the BP presents more than 70% of collision cases. This sort of behavior is the one that is more likely to appear in a real-world setting.

Fig. 1. Average time to reach the goal as crowd density increases. An ORCA-based simulator is not necessarily informative for an algorithm's performance, meanwhile, SFM adds higher levels of complexity to the testing scenarios. Navigation algorithms should seek training and testing scenarios that include a different level of complexity, heterogeneity, and cooperations from the human side, such as the one provided by the SFM.

respectively. Surprisingly, in Figure 1(a), we see that no matter how dense the crowd is, even a simple social planner such as SSP or BP significantly outperforms an ORCA-driven robot in an ORCA environment. ORCA agents are designed to be cooperative, allowing each other agent to take half of the responsibility of avoiding pairwise collisions. Besides, ORCA is only guaranteed not to collide and choose the maximum velocity in the next step, but it does not consider the shortest path overall. It lacks a mechanism to represent the general navigation task, while a naive planner such as SSP preserves this geometry information for path planning, achieving greater performance. For reference, in a more complex and heterogeneous scenario, such as the one represented by the SFM environment, the SSP underperforms compared to the other methods. SSP fails to reach the goal when the density surpasses a certain level as seen in Figure 1(b), while the BP ends in a great number of collisions. This sort of behavior for basic planners is the one that is more likely to appear in a real-world setting.

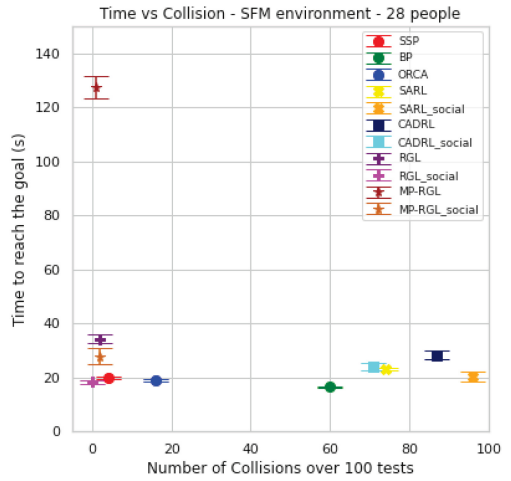
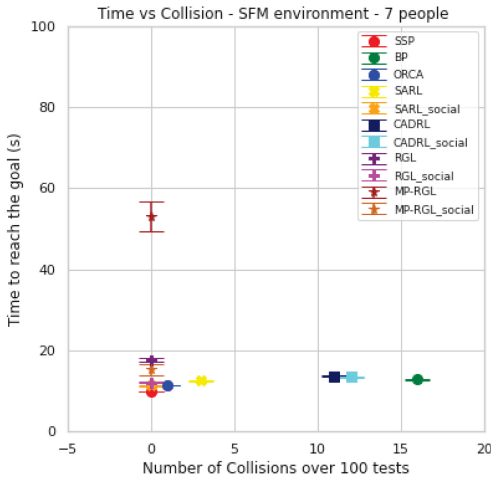
Our main takeaway from this example study is that testing *only* on the ORCA simulator environment is not necessarily informative for an algorithm's performance. Additionally, the SFM adds higher levels of complexity to the testing scenarios. Hence, social navigation algorithms that can outperform these simple planners in such an environment, or an even more complex and heterogeneous setup, are most likely to show good performance in a real-world setting.

4.3.1.3. Suboptimality of State-of-the-art Approaches. Figure 2 reports the performance of all algorithms in the ORCA and SFM environments, respectively, for two density levels of interest: low density— 0.05 people/m^2 (7 human agents)—and high density— 0.19 people/m^2 (28 human agents). Overall, we see the performance of the learning-based baselines does not transfer in domains



(a) Fully cooperative simulation environment, ≈ 0.05 people/ m^2 . The subscript *social* indicates that the algorithms have been trained in a SFM environment.

(b) Fully cooperative simulation environment, ≈ 0.19 people/ m^2 . The subscript *social* indicates that the algorithms have been trained in a SFM environment.



(c) Fully cooperative simulation environment, ≈ 0.05 people/ m^2 . The subscript *social* indicates that the algorithms have been trained in a SFM environment.

(d) Fully cooperative simulation environment, ≈ 0.19 people/ m^2 . The subscript *social* indicates that the algorithms have been trained in a SFM environment.

Fig. 2. Average time and standard error to reach the goal vs. number of collisions for different RL algorithms as crowd density increases on ORCA (2(a), 2(b)) and SFM (2(c), 2(d)) environment. The SSP outperforms all state-of-the-art RL methods and is nearly identical to the best-performing RL method. Training in the SFM environment rather than in the ORCA environment strongly affects the policy’s performance. A robot trained in the ORCA environment results in a more conservative policy. However, a robot trained in the SFM environment moves more efficiently (faster) and less safely (more collisions). Additionally, it shows better performances when tested out-of-distribution, represented by the ORCA environment in this specific case.

falling outside of the test set distribution. Training in the SFM environment rather than in the ORCA environment strongly affects the policy’s performance. A robot trained in the ORCA environment results in a more conservative policy. However, a robot trained in the SFM environment moves more efficiently (faster) and less safely (more collisions). Additionally, an algorithm

trained in a more complex and heterogeneous environment, such as the one represented by the SFM, shows better performances when tested out-of-distribution, represented by the ORCA environment in this specific case. These findings illustrate the importance of having a good high-level demonstrator, which, in our case, is given by the crowd's model. The algorithm chosen to represent the crowd's dynamics highly affects the performance of the learning algorithm. They also demonstrate that existing social navigation algorithms tend to struggle to outperform naively simple baselines in existing simulation environments, highlighting the lack of more realistic simulation suites and simulation benchmarking standards.

4.3.1.4. Behavioral Model Assumptions. Another limitation, as pointed out by Fraichard and Levesy [76], is that typical crowd simulators assume that all agents obey the same law, and all agents know the state of all other neighboring agents. However, individual behavior can affect the whole crowd's motion. Current crowd models fail to describe the relationships between human behaviors systematically and to enable heterogeneity. A central challenge for a crowd simulator is building models capable of integrating different crowd behaviors and interpreting how such actions affect the individuals' movement under a unified mechanism representing different scenarios. Hence, one way to improve existing simulator suites would be to consider combining different modeling approaches that increase the crowd heterogeneity while modeling and representing complex individual behavior in different scenarios. Additionally, it is possible to improve existing methods by varying the degrees of "attention" of the human agents. For example, a fraction of the agents could be considered inattentive, e.g., some people in crowds are looking at their phones. However, the agent should not be considered "fully" inattentive, since the robot in the periphery could still trigger attentiveness. A crowd simulator should also account for the agent's "flexibility," i.e., how much humans are willing to compromise on their desired trajectory. People change their speed, desired directions, and they stop walking for several reasons. These rapid changes in the human pattern are not just due to the nearby people or obstacles, but to anything that caught their attention.

4.3.2 Towards a Benchmarking Protocol. Despite the limitations of existing simulation practices, simulation remains an important tool for social navigation research. Real-world studies are often costly, require extensive IRB reviews, and are challenging to design properly. Therefore, before proceeding to real-world testing procedures, it is crucial to acquire insights from repeated testing in simulation. Overall, as argued by Steinfeld et al. [233], methodologically rigorous human modeling is of high value role to human-robot interaction research. However, this necessitates addressing the human modeling issues raised above. Core research on the design of models for human locomotion and organization may offer significant insights into the design of models for human simulation [208, 275].

The complete validation of a social navigation framework requires extensive, statistically rigorous experimental testing in close proximity with humans in real-world environments. Existing literature in social navigation research is dominated by simulation studies with the limitations described in the previous sections, whereas experimental work is often lacking statistical significance (e.g., see Demos in Table 4). In fact, much of the literature reports experimental demonstrations that inevitably capture a small, biased sample of rich behaviors that often emerge in real-world environments. However, even studies with more statistical rigor often suffer from design considerations. Specifically, existing lab studies often capture unrealistically simplistic or overly specific interactions with users. In contrast, the findings of field studies conducted in the wild may suffer from uncontrolled variables and the noise underlying open-form interaction with users. Furthermore, there is no consensus over objective or subjective measures for evaluating social navigation algorithms, neither a standardized repertoire of scenarios to test or platforms to experiment with.

Overall, these issues derive from the lack of a standardized experimental benchmark. Given the rich interest in the field, it is important as a community to converge towards a benchmark definition, clearly specifying all required experimental attributes. At the minimum, these should include: the exact definition of the task that the robot is performing; the types of experimental the implications of using nonholonomically constrained systems a set of recommended platforms with their suitability and limitations for the task, e.g., positive and negative implications of the systems with nonholonomic kinematics; the number and role of humans in the experiment; the context and the role of the environment; the metrics we should be measuring performance with; the baselines we should be comparing against.

4.3.3 Formal Verification. It should be noted that the existence of established real-world benchmarks is not sufficient on its own to guarantee desirable performance. Any benchmark is still inevitably at best capturing a small sample of representative real-world interactions. Further, experiments are costly in terms of time and capital, prohibiting the possibility of exhaustive real-world testing. However, navigation in close proximity with humans is a safety-critical application, and it is important that systems deployed in real-world settings are provably fail-safe. This highlights the important role that simulation still plays and the need for addressing the issues raised in Section 4.3.1. It also motivates the possible employment of methods for formal verification as practices to ensure robust performance [139]. This is especially important, considering the poor performance of existing algorithms on out-of-distribution test cases. While there is a wide body of work on the formal verification of robotic systems, including work specifically on collision avoidance [177], relatively limited attention has been placed to the more complex problem of verifying navigation systems operating in pedestrian spaces [117]. Given the richness of pedestrian domains and the importance of delivering safe performance even under the highly likely emergence of system failures, more research is necessary to ensure graceful operation in realistic settings.

5 CONCLUSION

We presented an overview of the field of social robot navigation with a survey spanning three decades of exciting developments. Acknowledging the important progress to date, we presented our perspective on a series of core challenges that prohibit mobile robots from seamlessly navigating in human-populated environments today. We classified these challenges into three broad categories: planning, behavior design, and evaluation.

First, we highlighted the importance of modeling interaction for social navigation research. We argued for the tight coupling between prediction and planning and traced its benefits in the literature. We also shared our thoughts on fundamental challenges underlying the problem of planning robot motion in crowded human environments. Specifically, planning in such spaces is NP-hard, and the non-convexity of balancing standard efficiency/safety tradeoffs prohibits guaranteeing practical performance in real-world applications. Additional complications arise from the difficulty of modeling human agents and from the sensitivity to different types of contexts.

In terms of behavior design, we called for a need to organize and establish principled knowledge that governs social navigation behavior. We focused on spatial and motion behavior of pedestrians and mobile robots. And we reached the conclusion that literature in this domain is vast and varied, but no single theory stands out. It is generally accepted that considerations of certain behavior elements such as proxemics are important, but it is often unclear what is the accepted way to incorporate such elements. We argued that to establish such principled knowledge, both extensive testing in diverse large-scale environments and practicality testing by incorporating into social navigation models are needed for the proposed behavior models in the literature.

Finally, we pointed out issues with existing evaluation practices. We highlighted the shortcomings of commonly adopted simulation conventions. Surprisingly, we demonstrated that evaluating

on the widely employed ORCA environment is not necessarily informative, since a naive agent can leverage the cooperativeness of ORCA agents and outperform ORCA agents. Further, we saw that learning-based state-of-the-art approaches are extremely sensitive to out-of-distribution testing (equivalently, learning-based methods do not generalize well). Motivated by these observations, we highlighted other strong assumptions used by existing simulation environments (homogeneity, fixed attention, and fixed flexibility) and argued for the importance of those parameters as design variables. Finally, we highlighted the lack of standardized benchmarks for social navigation research and argued for their importance as the community grows.

5.1 Other Challenges

While our categorization of open challenges captures a wide range of problems, the performance of social robot navigation frameworks is practically constrained by a number of broad challenges that were not directly within the scope of this review.

5.1.1 Robot Design. Since the early days, roboticists have experimented with a wide variety of robot platforms, featuring a wide range of kinematic models, aesthetics, and anthropomorphic features.

In terms of kinematics, the prevalent paradigm involves wheeled mobile robots with a differential drive. While well-studied and universally available, such robots tend to produce motion that is not humanlike, often confusing humans. Recently, there have been notable alternatives, including the *Ballbot* paradigm [64, 158, 181], a robot resembling an inverted pendulum that is dynamically stabilizing on a single point of contact with the floor. Ballbots feature a humanlike form factor and agility that opens up the possibility of smooth merging in densely crowded environments.

In terms of appearance, many works tend to ignore the aesthetics aspect, focusing on the development of customized functional systems or directly using existing robot platforms. Notably, RHINO [26] did not feature any anthropomorphic features but MINERVA [244], the second-generation museum guide robot, featured a motorized face. Over the following years, many works employed a wide range of platforms with varying levels of humanlikeness, including anthropomorphic features such as a hat [249] or a virtual face drawn on a screen [168]. There is a long and extensive body of work in the HRI literature focusing on the tradeoffs of face design [55, 118]. Further, it has been shown that the form factor itself is an important aspect of design when it comes to HRI applications [203].

5.1.2 Perception Challenges. The planning and behavior prediction literature in social navigation often assumes that the robot is equipped with perfect on-board perception capabilities enabling precise robot localization and precise tracking of surrounding humans. However, outside of the lab, on-board perception will inevitably be limited and often prone to failures. Besides the mobile navigation stack required for robot localization and scene segmentation, a social navigation system must be equipped with a people-tracking module that provides robust metric information for surrounding humans. While existing approaches [27] may provide such functionalities, their performance is highly dependent on the sensing hardware of the robot and the conditions in the robot's environment. Further research is required to enable social robot navigation systems to carefully account for the inherent perception uncertainty at the prediction and planning cycle. It is also important that these systems employ principled mechanisms for handling perception failures during runtime.

5.1.3 Novelty Effects. An important feature that is often not explicitly accounted for in the existing literature is the *novelty effect*. While the widespread introduction of robot vacuum cleaners in households has contributed to the familiarization of humans to robots, it is expected—and in

many cases anecdotally confirmed [249]—that the spontaneous, unscripted encounter of humans with navigating robots in the wild features a great amount of unmodeled phenomena. The robot often becomes the pole of attraction, spiking the curiosity of bystanders who approach it, touch it, and try to play with it or challenge it [25]. These types of interactions may complicate the robot’s reaction and yield unpredictable and even unsafe responses. Finally, it is expected that these novelty effects will gradually dampen out as humans in a particular context become more familiar with a system. Therefore, it is important that the design follows principles of lifelong learning [245], accounting for the gradual adaptation of humans.

5.1.4 Robust Autonomy. Despite any advances in the development of socially aware navigation systems, the complexity of real-world environments can always pose a threat to robot autonomy. For instance, a robot may get physically stuck, run out of battery, get delocalized, fall, and so on. These situations could not only make the robot fail at its primary task but also possibly discomfort or put in danger surrounding humans. Thus, it is crucial to build mechanisms and systems that could enable robots to recover effectively.

One way is via robust navigation strategies like wandering [23, 184]. Notably, such strategies have empowered robot vacuum cleaners to operate effectively without localization [46] in human households. This strategy may work in simpler environments, when localization is not required by the primary task or when it is not important to account for the type of motion that the robot generates. While on many occasions this is indeed the case, environments featuring complex crowd motion and randomly placed obstacles with various shapes and properties could limit the effectiveness of techniques like wandering.

A more general way to enhance the robustness of a navigation system could be via human help [211]. Minimal human help has enabled even robots with limited sensing and compute capabilities to effectively navigate challenging spaces like long academic corridors [184]. Such studies elevate the use of human interventions and feedback for mobile robots into an important medium for robot recovery and lifelong learning. We envision that robots would eventually be able to leverage bystander help to recover via or learn from human help towards improving their capabilities and requiring less intervention in future deployments. Doing so is challenging in itself, as human help can be contingent on a person’s internal state [182]. When a person is rushing, it is less likely to help a robot in need for help. When a robot asks for help too often, it might predispose people into not helping it again in the future. Extracting models of relevant human internal states and dynamics could enable robots to infer when and how to ask for help. This requires more research on modeling human interruptibility [12] and integrating it into robot decision-making.

5.2 Moving Forward

The rich interest of the community in the area of social robot navigation and fast-paced neighboring fields, such as human behavior prediction [214] and autonomous driving [24], are promising indicators for exciting developments over the next decade. To this end, we aspire that this survey will constitute an important asset for researchers as they plan their research agenda in the years to come.

REFERENCES

- [1] Henny Admoni and Brian Scassellati. 2017. Social Eye gaze in human-robot interaction: A review. *J. Hum.-Robot Interact.* 6, 1 (2017), 25–63.
- [2] Alexandre Alahi, Kratharth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. 2016. Social LSTM: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR’16)*. 961–971.

- [3] Xavier Alameda-Pineda, Yan Yan, Elisa Ricci, Oswald Lanz, and Nicu Sebe. 2015. Analyzing free-standing conversational groups: A multimodal approach. In *Proceedings of the ACM International Conference on Multimedia (MM'15)*. 5–14.
- [4] P. Althaus, H. Ishiguro, T. Kanda, T. Miyashita, and H. I. Christensen. 2004. Navigation for human-robot interaction tasks. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'04)*, Vol. 2. 1894–1900.
- [5] Toshitaka Amaoka, Hamid Laga, Suguru Saito, and Masayuki Nakajima. 2009. Personal space modeling for human-computer interaction. In *Entertainment Computing (ICEC'09)*. Springer Berlin, 60–72.
- [6] Anoop Aroor, Susan L. Epstein, and Raj Korpan. 2018. Mengers: A crowd simulation tool for autonomous robot navigation. *arXiv preprint arXiv:1801.08823* (2018).
- [7] AUTODESK. 2019. Maya. Retrieved from <https://www.autodesk.com/products/maya>.
- [8] Mohammad Babaiezadeh, Iuri Frosio, Stephen Tyree, Jason Clemons, and Jan Kautz. 2016. Reinforcement learning through asynchronous advantage actor-critic on a GPU. *arXiv e-prints*, Article arXiv:1611.06256.
- [9] H. Bai, S. Cai, N. Ye, D. Hsu, and W. S. Lee. 2015. Intention-aware online POMDP planning for autonomous driving in a crowd. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'15)*. 454–460.
- [10] C. L. Baker, J. Tenenbaum, and R. Saxe. 2009. Action understanding as inverse planning. *Cognition* 113, 3 (2009), 329–349.
- [11] Tirthankar Bandyopadhyay, Kok Sung Won, Emilio Frazzoli, David Hsu, Wee Sun Lee, and Daniela Rus. 2013. Intention-aware motion planning. In *Algorithmic Foundations of Robotics X*, Emilio Frazzoli, Tomas Lozano-Perez, Nicholas Roy, and Daniela Rus (Eds.). Springer Berlin, 475–491.
- [12] Siddhartha Banerjee, Andrew Silva, and Sonia Chernova. 2018. Robot classification of human interruptibility and a study of its effects. *Trans. Hum.-robot Interact.* 7, 2 (Oct. 2018).
- [13] Kim Baraka, Stephanie Rosenthal, and Manuela Veloso. 2016. Enhancing human understanding of a mobile robot's state and actions using expressive lights. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 652–657.
- [14] Marie-Lou Barnaud, Nicolas Morgado, Richard Palluel-Germain, Julien Diard, and Anne Spalanzani. 2014. Proxemics models for human-aware navigation in robotics: Grounding interaction and personal space models in experimental data from psychology. In *Proceedings of the "Assistance and Service Robotics in a Human Environment" Workshop, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- [15] Basefount. 2020. Miarmy. Retrieved from <http://www.basefount.com/miarmy.html>.
- [16] L. Bazzani, M. Cristani, and V. Murino. 2012. Decentralized particle filter for joint individual-group tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 1886–1893.
- [17] N. Bellotto and H. Hu. 2010. A bank of unscented Kalman filters for multimodal human perception with mobile service robots. *Int. J. Soc. Robot.* 2 (2010), 121–136.
- [18] Maren Bennewitz, Wolfram Burgard, Grzegorz Cielniak, and Sebastian Thrun. 2005. Learning motion patterns of people for compliant robot motion. *Int. J. Robot. Res.* 24, 1 (2005), 31–48.
- [19] Aniket Bera, Tanmay Randhavana, Rohan Prinja, and Dinesh Manocha. 2017. SocioSense: Robot navigation amongst pedestrians with social and psychological constraints. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 7018–7025.
- [20] Mitchell A. Berger. 2001. Topological invariants in braid theory. *Lett. Math. Phys.* 55, 3 (2001), 181–192.
- [21] Joan S. Birman. 1975. *Braids Links and Mapping Class Groups*. Princeton University Press.
- [22] Bruno Brito, Boaz Floor, Laura Ferranti, and Javier Alonso-Mora. 2019. Model predictive contouring control for collision avoidance in unstructured dynamic environments. *IEEE Robot. Autom. Lett.* 4, 4 (2019), 4459–4466. DOI: <https://doi.org/10.1109/LRA.2019.2929976>
- [23] R. Brooks. 1986. A robust layered control system for a mobile robot. *IEEE J. Robot. Autom.* 2, 1 (1986), 14–23.
- [24] Kyle Brown, Katherine Driggs-Campbell, and Mykel J. Kochenderfer. 2020. A taxonomy and review of algorithms for modeling and predicting human driver behavior. *arXiv e-prints*, Article arXiv:2006.08832.
- [25] Dražen Brščić, Hiroyuki Kidokoro, Yoshitaka Suehiro, and Takayuki Kanda. 2015. Escaping from children's abuse of social robots. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI'15)*. 59–66.
- [26] Wolfram Burgard, Armin B. Cremers, Dieter Fox, Dirk Hähnel, Gerhard Lakemeyer, Dirk Schulz, Walter Steiner, and Sebastian Thrun. 1998. The interactive museum tour-guide robot. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI'98)*. 11–18.
- [27] Z. Cao, G. Hidalgo Martinez, T. Simon, S. Wei, and Y. A. Sheikh. 2021. OpenPose: Realtime multi-person 2D pose estimation using part affinity fields. *IEEE Trans. Pattern Anal. Mach. Intell.* 43, 1 (2021), 172–186.
- [28] Daniel Carton, Wiktor Olszowy, and Dirk Wollherr. 2016. Measuring the effectiveness of readability for mobile robot locomotion. *Int. J. Soc. Robot.* 8, 5 (2016), 721–741.
- [29] Vinícius Jurinic Cassol. 2016. *Crowdsim: A Framework to Estimate Safety of Egress Performance in Real Life Scenarios*. Ph.D. Dissertation. Pontificia Universidade Católica do Rio Grande do Sul.

- [30] Isarun Chamveha, Yusuke Sugano, Yoichi Sato, and Akihiro Sugimoto. 2013. Social group discovery from surveillance videos: A data-driven approach with attention-based cues. In *Proceedings of the the British Machine Vision Association (BMVC)*.
- [31] Rohan Chandra, Tianrui Guan, Srujan Panuganti, Trisha Mittal, Uttaran Bhattacharya, Aniket Bera, and Dinesh Manocha. 2019. Forecasting trajectory and behavior of road-agents using spectral clustering in graph-LSTMs. *arXiv e-prints*, Article arXiv:1912.01118.
- [32] M. Chang, N. Krahnstoever, and W. Ge. 2011. Probabilistic group-level motion analysis and scenario recognition. In *Proceedings of the International Conference on Computer Vision (ICCV'11)*. 747–754.
- [33] Konstantinos Charalampous, Ioannis Kostavelis, and Antonios Gasteratos. 2017. Recent trends in social aware robot navigation: A survey. *Robot. Auton. Syst.* 93 (2017), 85–104.
- [34] I. Chatterjee and A. Steinfeld. 2016. Performance of a low-cost, human-inspired perception approach for dense moving crowd navigation. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 578–585.
- [35] Antonio Chella and Irene Macaluso. 2009. The perception loop in CiceRobot, a museum guide robot. *Neurocomputing* 72, 4–6 (2009), 760–766.
- [36] Changan Chen, Sha Hu, Payam Nikdel, Greg Mori, and Manolis Savva. 2019. Relational graph learning for crowd navigation. *arXiv preprint arXiv:1909.13165* (2019).
- [37] Changan Chen, Yuejiang Liu, Sven Kreiss, and Alexandre Alahi. 2018. Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning. *arXiv preprint arXiv:1809.08835* (2018).
- [38] Weihua Chen, Tie Zhang, and Yanbiao Zou. 2018. Mobile robot path planning based on social interaction space in social environment. *Int. J. Adv. Robot. Syst.* 15, 3 (2018), 1729881418776183.
- [39] Y. F. Chen, M. Everett, M. Liu, and J. P. How. 2017. Socially aware motion planning with deep reinforcement learning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 1343–1350.
- [40] Yu Fan Chen, Miao Liu, Michael Everett, and Jonathan P. How. 2017. Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 285–292.
- [41] A. M. Cheriyyadat and R. J. Radke. 2008. Detecting dominant motions in dense crowds. *IEEE J. Select. Topics Sig. Process.* 2, 4 (2008), 568–581.
- [42] Chiho Choi, Srikanth Malla, Abhishek Patil, and Joon Hee Choi. 2019. DROGON: A trajectory prediction model based on intention-conditioned behavior reasoning. *arXiv e-prints*, Article arXiv:1908.00024.
- [43] Wongun Choi, Yu-Wei Chao, Caroline Pantofaru, and Silvio Savarese. 2014. Discovering groups of people in images. In *Proceedings of the European Conference on Computer Vision (ECCV)*. Springer International Publishing, Cham, 417–433.
- [44] Shu-Yun Chung and Han-Pang Huang. 2010. A mobile robot that understands pedestrian spatial behaviors. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 5861–5866.
- [45] A. Clodic, S. Fleury, R. Alami, R. Chatila, G. Bailly, L. Brethes, M. Cottret, P. Danes, X. Dollat, F. Elisei, I. Ferrane, M. Herrb, G. Infantes, C. Lemaire, F. Lerasle, J. Manhes, P. Marcoul, P. Menezes, and V. Montreuil. 2006. Rackham: An interactive robot-guide. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 502–509.
- [46] CNET. 2021. This Is Why Your Roomba’s Random Patterns Actually Make Perfect Sense. Retrieved from <https://www.cnet.com/home/kitchen-and-household/this-is-why-your-roombas-random-patterns-actually-make-perfect-sense/>.
- [47] Gregory F. Cooper. 1990. The computational complexity of probabilistic inference using Bayesian belief networks. *Artif. Intell.* 42, 2 (1990), 393–405.
- [48] CROWDBOT. 2021. CROWDBOT. Retrieved from <http://www.crowdbot.eu>.
- [49] Gergely Csibra and György Gergely. 2007. “Obsessed with goals”: Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychol.* 124, 1 (Jan. 2007), 60–78.
- [50] N. P. Cuntoor, R. Collins, and A. J. Hoogs. 2012. Human-robot teamwork using activity recognition and human instruction. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 459–465.
- [51] S. Curtis, A. Best, and D. Manocha. 2016. *Menge: A Modular Framework for Simulating Crowd Movement*. Technical Report. University of North Carolina at Chapel Hill.
- [52] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. 2009. The complexity of computing a Nash equilibrium. *SIAM J. Comput.* 39, 1 (2009), 195–259.
- [53] F. Dellaert, D. Fox, W. Burgard, and S. Thrun. 1999. Monte Carlo localization for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Vol. 2. 1322–1328.
- [54] Kenneth L. Dion. 2000. Group cohesion: From “field of forces” to multidimensional construct. *Group Dynam. Theor. Res. Pract.* 4, 1 (2000), 7.

- [55] Carl F. DiSalvo, Francine Gemperle, Jodi Forlizzi, and Sara Kiesler. 2002. All robots are not created equal: The design and perception of humanoid robot heads. In *Proceedings of the Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques (DIS)*. 321–326.
- [56] A. D. Dragan, K. C. T. Lee, and S. S. Srinivasa. 2013. Legibility and predictability of robot motion. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 301–308.
- [57] Anca D. Dragan and Siddhartha Srinivasa. 2014. Integrating human observer inferences into robot motion planning. *Auton. Robots* 37, 4 (2014), 351–368.
- [58] N. Du Toit and J. Burdick. 2012. Robot motion planning in dynamic, uncertain environments. *IEEE Trans. Robot.* 28, 1 (2012), 101–115.
- [59] T. Dutra, R. Marques, J. Cavalcante-Neto, C. Vidal, and J. Pettre. 2017. Gradient-based steering for vision-based crowd simulation algorithms. *Computat. Graph. Forum* 36, 2 (2017), 337–348.
- [60] Ivan Dynnikov and Bert Wiest. 2007. On the complexity of braids. *J. Eur. Math. Soc.* 009, 4 (2007), 801–840.
- [61] M. Everett, Y. F. Chen, and J. P. How. 2018. Motion planning among dynamic, decision-making agents with deep reinforcement learning. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 3052–3059.
- [62] Hans J. Eysenck. 1950. *Dimensions of Personality*. Vol. 5. Transaction Publishers.
- [63] T. Fan, X. Cheng, J. Pan, P. Long, W. Liu, R. Yang, and D. Manocha. 2019. Getting robots unfrozen and unlost in dense pedestrian crowds. *IEEE Robot. Autom. Lett.* 4, 2 (2019), 1178–1185.
- [64] Péter Fankhauser. 2010. *Modeling and Control of a Ballbot*. Master’s Thesis. ETH Zurich.
- [65] Francesco Farina, Daniele Fontanelli, Andrea Garulli, Antonio Giannitrapani, and Domenico Prattichizzo. 2017. Walking ahead: The headed social force model. *PLoS One* 12, 1 (2017), 1–23.
- [66] D. Feil-Seifer and M. Mataric. 2011. People-aware navigation for goal-oriented behavior involving a human partner. In *Proceedings of the IEEE International Conference on Development and Learning (ICDL)*, Vol. 2. 1–6.
- [67] L. Feng and B. Bhanu. 2015. Tracking people by evolving social groups: An approach with social network perspective. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*. 109–116.
- [68] Rolando Fernandez, Nathan John, Sean Kirmani, Justin Hart, Jivko Sinapov, and Peter Stone. 2018. Passive demonstrations of light-based robot signals for improved human interpretability. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 234–239.
- [69] Franck Feurtey. 2000. *Simulating the Collision Avoidance Behavior of Pedestrians*. Master’s Thesis. University of Tokyo, Tokyo, Japan.
- [70] Michelangelo Fiore, Harmish Khambhaita, Grégoire Milliez, and Rachid Alami. 2015. An adaptive and proactive human-aware robot guide. In *Social Robotics*. Springer International Publishing, Cham, 194–203.
- [71] Paolo Fiorini and Zvi Shiller. 1998. Motion planning in dynamic environments using velocity obstacles. *Int. J. Robot. Res.* 17, 7 (1998), 760–772.
- [72] A. Foka and P. Trahanias. 2010. Probabilistic autonomous robot navigation in dynamic environments with human motion prediction. *Int. J. Soc. Robot.* 2 (2010), 79–94.
- [73] Blender Foundation. 2018. Blender—A 3D Modelling and Rendering Package. Retrieved from <http://www.blender.org>.
- [74] Dieter Fox, Wolfram Burgard, and Sebastian Thrun. 1998. A hybrid collision avoidance method for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Vol. 2. 1238–1243.
- [75] T. Fraichard. 2007. A short paper about motion safety. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 1140–1145.
- [76] T. Fraichard and V. Levesy. 2020. From crowd simulation to robot navigation in crowds. *IEEE Robot. Autom. Lett.* 5, 2 (2020), 729–735.
- [77] D. Fridovich-Keil, E. Ratner, L. Peters, A. D. Dragan, and C. J. Tomlin. 2020. Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 1475–1481.
- [78] A. Garrell and A. Sanfeliu. 2010. Local optimization of cooperative robot movements for guiding and regrouping people in a guiding mission. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 3294–3299.
- [79] Anais Garrell and Alberto Sanfeliu. 2012. Cooperative social robots to accompany groups of people. *Int. J. Robot. Res.* 31, 13 (2012), 1675–1701.
- [80] W. Ge, R. T. Collins, and R. B. Ruback. 2012. Vision-based analysis of small groups in pedestrian crowds. *IEEE Trans. Pattern Anal. Mach. Intell.* 34, 5 (2012), 1003–1016.
- [81] G. Gennari and G. D. Hager. 2004. Probabilistic data association methods in visual tracking of groups. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, Vol. 2.
- [82] Martin Gérin-Lajoie, Carol L. Richards, Joyce Fung, and Bradford J. McFadyen. 2008. Characteristics of personal space during obstacle circumvention in physical and virtual environments. *Gait Post.* 27, 2 (2008), 239–247.

- [83] Brian Gerkey, Richard T. Vaughan, and Andrew Howard. 2003. The player/stage project: Tools for multi-robot and distributed sensor systems. In *Proceedings of the International Conference on Advanced Robotics (ICAR)*, Vol. 1. 317–323.
- [84] C. Gloor. 2016. PedSim: Pedestrian Crowd Simulation. Retrieved from <http://pedsim.silmaril.org>.
- [85] Rachel Gockley, Jodi Forlizzi, and Reid Simmons. 2007. Natural person-following behavior for social robots. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 17–24.
- [86] Erving Goffman. 1966. *Behavior in Public Places: Notes on the Social Organization of Gatherings*. Free Press.
- [87] Erving Goffman. 1971. *Relations in Public: Microstudies of the Public Order*. Basic Books, New York.
- [88] Golaem. 2020. Golaem Crowd. Retrieved from <http://golaem.com/>.
- [89] C. Granata and P. Bidaud. 2012. A framework for the design of person following behaviors for social mobile robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 4652–4659.
- [90] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. 2018. Social GAN: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2255–2264.
- [91] Stephen J. Guy, Jatin Chhugani, Changkyu Kim, Nadathur Satish, Ming Lin, Dinesh Manocha, and Pradeep Dubey. 2009. ClearPath: Highly parallel collision avoidance for multi-agent simulation. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. 177–187.
- [92] Stephen J. Guy, Sean Curtis, Ming C. Lin, and Dinesh Manocha. 2012. Least-effort trajectories lead to emergent crowd behaviors. *Phys. Rev. E* 85, 1 (2012), 016110.
- [93] Stephen J. Guy, Sujeong Kim, Ming C. Lin, and Dinesh Manocha. 2011. Simulating heterogeneous crowd behaviors using personality trait theory. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. 43–52.
- [94] Jérôme Guzzi, Alessandro Giusti, Luca M. Gambardella, Guy Theraulaz, and Gianni A. Di Caro. 2013. Human-friendly robot navigation in dynamic environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 423–430.
- [95] John K. Haas. 2014. *A History of the Unity Game Engine*. Master’s Thesis. Worcester Polytechnic Institute.
- [96] E. T. Hall. 1963. A system for notation of proxemic behavior. *Amer. Anthropol.* 65, 5 (1963), 1003–1026.
- [97] E. T. Hall. 1966. *The Hidden Dimension*. Anchor Books.
- [98] Mahmoud Hamandi, Mike D’Arcy, and Pooyan Fazli. 2019. DeepMoLon: Learning to navigate like humans. In *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. 1–7.
- [99] Yanlin Han and Piotr Gmytrasiewicz. 2018. Learning others’ intentional models in multi-agent settings using interactive POMDPs. In *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS)*, Vol. 31. 5634–5642.
- [100] Justin W. Hart, Brian Gleeson, Matthew Pan, Ajung Moon, Karon MacLean, and Elizabeth Croft. 2014. Gesture, gaze, touch, and hesitation: Timing cues for collaborative work. In *Proceedings of the HRI Workshop on Timing in Human-Robot Interaction*.
- [101] L. A. Hayduk. 1981. The shape of personal space: An experimental investigation. *Canad. J. Behav. Sci.* 13, 1 (1981), 87–93.
- [102] L. Hayduk. 1994. Personal space: Understanding the simplex model. *J. Nonverb. Behav.* 18 (1994), 245–260.
- [103] D. Helbing. 1998. A fluid dynamic model for the movement of pedestrians. *Complex Syst.* 6, 5 (1998), 391–415.
- [104] Dirk Helbing, Illés Farkas, and Tamas Vicsek. 2000. Simulating dynamical features of escape panic. *Nature* 407, 6803 (2000), 487–490.
- [105] Dirk Helbing and Péter Molnár. 1995. Social force model for pedestrian dynamics. *Phys. Rev. E* 51, 5 (1995), 4282–4286.
- [106] D. Helbing, P. Molnar, I. Farkas, and K. Bolay. 2001. Self-organizing pedestrian movement. *Environ. Plan.ning B: Plan. Des.* 28, 3 (2001), 361–383.
- [107] P. Henry, C. Vollmer, B. Ferris, and D. Fox. 2010. Learning to navigate through crowded environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 981–986.
- [108] Roger L. Hughes. 2002. A continuum theory for the flow of pedestrians. *Transport. Res. Part B: Methodol.* 36, 6 (2002), 507–535.
- [109] Hayley Hung and Ben Kröse. 2011. Detecting F-Formations as dominant sets. In *Proceedings of the 13th International Conference on Multimodal Interfaces (ICMI)*. 231–238.
- [110] T. Iqbal, S. Rack, and L. D. Riek. 2016. Movement coordination in human-robot teams: A dynamical systems approach. *IEEE Trans. Robot.* 32, 4 (2016), 909–919.
- [111] T. Iqbal and L. D. Riek. 2017. Coordination dynamics in multihuman multirobot teams. *IEEE Robot. Autom. Lett.* 2, 3 (2017), 1712–1717.
- [112] Qingge Ji, Fuchuan Wang, and Ting Zhu. 2016. VPBS: A velocity-perception-based SFM approach for crowd simulation. In *Proceedings of the IEEE International Conference on Virtual Reality and Visualization (ICVRV)*. 317–324.

- [113] Hao Jiang, Wenbin Xu, Tianlu Mao, Chunpeng Li, Shihong Xia, and Zhaoqi Wang. 2010. Continuum crowd simulation in complex environments. *Comput. Graph.* 34, 5 (2010), 537–544.
- [114] Michiel Jooisse, Manja Lohse, Niels Van Berkel, Aziez Sardar, and Vanessa Evers. 2021. Making appearances: How robots should approach people. *ACM Trans. Hum.-robot Interact.* 10, 1 (2021).
- [115] Joshua Joseph, Finale Doshi-Velez, Albert S. Huang, and Nicholas Roy. 2011. A Bayesian nonparametric approach to modeling motion patterns. *Auton. Robots* 31, 4 (2011), 383–400.
- [116] E. Jung, B. Yi, and S. Yuta. 2012. Control algorithms for a mobile robot tracking a human in front. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 2411–2416.
- [117] Sebastian Junges, Nils Jansen, Joost-Pieter Katoen, Ufuk Topcu, Ruohan Zhang, and Mary Hayhoe. 2018. Model checking for safe navigation among humans. In *Quantitative Evaluation of Systems*, Annabelle McIver and Andras Horvath (Eds.). Springer International Publishing, Cham, 207–222.
- [118] Alisa Kalegina, Grace Schroeder, Aidan Allchin, Keara Berlin, and Maya Cakmak. 2018. Characterizing the design space of rendered robot faces. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 96–104.
- [119] S. Karaman and E. Frazzoli. 2011. Sampling-based algorithms for optimal motion planning. *Int. J. Robot. Res.* 30, 7 (2011), 846–894.
- [120] Ioannis Karamouzas, Peter Heil, Pascal Van Beek, and Mark H. Overmars. 2009. A predictive collision avoidance model for pedestrian simulation. In *Proceedings of the International Workshop on Motion in Games*. Springer, 41–52.
- [121] Ioannis Karamouzas, Peter Heil, Pascal van Beek, and Mark H. Overmars. 2009. A predictive collision avoidance model for pedestrian simulation. In *Motion in Games*. Springer Berlin, 41–52.
- [122] Ioannis Karamouzas and Mark Overmars. 2011. Simulating and evaluating the local behavior of small pedestrian groups. *IEEE Trans. Visualiz. Comput. Graph.* 18, 3 (2011), 394–406.
- [123] I. Karamouzas, B. Skinner, and S. J. Guy. 2014. Universal power law governing pedestrian interactions. *Phys. Rev. Lett.* (2014).
- [124] Yusuke Kato, Takayuki Kanda, and Hiroshi Ishiguro. 2015. May I help you?: Design of human-like polite approaching behavior. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 35–42.
- [125] Kapil Katyal, Yuxiang Gao, Jared Markowitz, Sara Pohland, Corban Rivera, I-Jeng Wang, and Chien-Ming Huang. 2020. Learning a group-aware policy for robot navigation. *arXiv e-prints*, Article arXiv:2012.12291.
- [126] L. E. Kavraki, P. Svestka, J. Latombe, and M. H. Overmars. 1996. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robot. Autom.* 12, 4 (1996), 566–580.
- [127] A. Kendon. 1990. Conducting interaction: Patterns of behavior in Focused Encounters. *Stud. Int. Socioling.* 7 (1990).
- [128] Harmish Khambhaita and Rachid Alami. 2020. Viewing robot navigation in human environment as a cooperative activity. In *Robotics Research*, Nancy M. Amato, Greg Hager, Shawna Thomas, and Miguel Torres-Torriti (Eds.). Springer International Publishing, Cham, 285–300.
- [129] Sultan Daud Khan, Giuseppe Vizzari, Stefania Bandini, and Saleh Basalamah. 2015. Detection of social groups in pedestrian crowds using computer vision. In *Advanced Concepts for Intelligent Vision Systems*, Sebastiano Battiato, Jacques Blanc-Talon, Giovanni Gallo, Wilfried Philips, Dan Popescu, and Paul Scheunders (Eds.). Springer International Publishing, Cham, 249–260.
- [130] O. Khatib. 1985. Real-time obstacle avoidance for manipulators and mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Vol. 2. 500–505.
- [131] Beomjoon Kim and Joelle Pineau. 2016. Socially adaptive path planning in human environments using inverse reinforcement learning. *Int. J. Soc. Robot.* 8, 1 (2016), 51–66.
- [132] Rachel Kirby. 2010. *Social Robot Navigation*. Ph.D. Dissertation. Carnegie Mellon University, Pittsburgh, PA.
- [133] Rachel Kirby, Reid Simmons, and Jodi Forlizzi. 2009. COMPANION: A constraint-optimizing method for person-acceptable navigation. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 607–612.
- [134] Kris M. Kitani, Brian D. Ziebart, James Andrew Bagnell, and Martial Hebert. 2012. Activity forecasting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, Andrew Fitzgibbon, Svetlana Lazebnik, Pietro Perona, Yoichi Sato, and Cordelia Schmid (Eds.). Springer Berlin, 201–214.
- [135] Ross A. Knepper and Daniela Rus. 2012. Pedestrian-inspired sampling-based multi-robot collision avoidance. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 94–100.
- [136] H. Knight and R. Simmons. 2014. Expressive motion with x, y and theta: Laban Effort Features for mobile robots. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 267–273.
- [137] Nathan Koenig and Andrew Howard. 2004. Design and use paradigms for Gazebo, an open-source multi-robot simulator. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Vol. 3. 2149–2154.
- [138] Y. Koren and J. Borenstein. 1991. Potential field methods and their inherent limitations for mobile robot navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Vol. 2. 1398–1404.

- [139] Hadas Kress-Gazit, Kerstin Eder, Guy Hoffman, Henny Admoni, Brenna Argall, Rüdiger Ehlers, Christoffer Heckman, Nils Jansen, Ross Knepper, Jan Křetínský, Shelly Levy-Tzedek, Jamy Li, Todd Murphey, Laurel Riek, and Dorsa Sadigh. 2021. Formalizing and guaranteeing human-robot interaction. *Commun. ACM* 64, 9 (Aug. 2021), 78–84. DOI: <https://doi.org/10.1145/3433637>
- [140] Henrik Kretzschmar, Markus Spies, Christoph Sprunk, and Wolfram Burgard. 2016. Socially compliant mobile robot navigation via inverse reinforcement learning. *Int. J. Robot. Res.* 35, 11 (2016), 1289–1307.
- [141] Thibault Kruse, Patrizia Basili, Stefan Glasauer, and Alexandra Kirsch. 2012. Legible robot navigation in the proximity of moving humans. In *Proceedings of the IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO)*. 83–88.
- [142] Thibault Kruse, Amit Kumar Pandey, Rachid Alami, and Alexandra Kirsch. 2013. Human-aware robot navigation: A survey. *Robot. Auton. Syst.* 61, 12 (2013), 1726–1743.
- [143] H. Kuzuoka, Y. Suzuki, J. Yamashita, and K. Yamazaki. 2010. Reconfiguring spatial formation arrangement by robot body orientation. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 285–292.
- [144] Minae Kwon, Malte F. Jung, and Ross A. Knepper. 2016. Human expectations of social robots. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 463–464.
- [145] C. Lam, C. Chou, K. Chiang, and L. Fu. 2011. Human-centered robot navigation—Towards a harmoniously human-robot coexisting environment. *IEEE Trans. Robot.* 27, 1 (2011), 99–112.
- [146] Przemysław A. Lasota, Terrence Fong, Julie A. Shah, et al. 2017. A survey of methods for safe human-robot interaction. *Found. Trends® Robot.* 5, 4 (2017), 261–349.
- [147] Boris Lau, Kai O. Arras, and Wolfram Burgard. 2010. Multi-model hypothesis group tracking and group size estimation. *Int. J. Soc. Robot.* 2, 1 (2010), 19–30.
- [148] Steven M. LaValle and James J. Kuffner Jr. 2001. Randomized kinodynamic planning. *Int. J. Robot. Res.* 20, 5 (2001), 378–400.
- [149] L. Leal-Taixé, G. Pons-Moll, and B. Rosenhahn. 2011. Everybody needs somebody: Modeling social and grouping behavior on a linear programming multiple people tracker. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops*. 120–127.
- [150] Alon Lerner, Yiorgos Chrysanthou, and Dani Lischinski. 2007. Crowds by example. *Comput. Graph. Forum* 26, 3 (2007), 655–664.
- [151] Kurt Lewin. 1945. The research center for group dynamics at Massachusetts Institute of Technology. *Sociometry* 8, 2 (1945), 126–136.
- [152] Jiachen Li, Fan Yang, Masayoshi Tomizuka, and Chiho Choi. 2020. EvolveGraph: Multi-agent trajectory prediction with dynamic relational reasoning. *arXiv e-prints*, Article arXiv:2003.13924.
- [153] Kunming Li, Mao Shan, Karan Narula, Stewart Worrall, and Eduardo Nebot. 2020. Socially aware crowd navigation with multimodal pedestrian trajectory prediction for autonomous vehicles. In *Proceedings of the IEEE International Conference on Intelligent Transportation Systems (ITSC)*. 1–8.
- [154] X. Li, X. Ying, and M. C. Chuah. 2019. GRIP: Graph-based interaction-aware trajectory prediction. In *Proceedings of the IEEE Intelligent Transportation Systems Conference (ITSC)*. 3960–3966.
- [155] Junwei Liang, Lu Jiang, Kevin Murphy, Ting Yu, and Alexander Hauptmann. 2020. The garden of forking paths: Towards multi-future trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 10505–10515.
- [156] Junwei Liang, Lu Jiang, Juan Carlos Nieves, Alexander G. Hauptmann, and Li Fei-Fei. 2019. Peeking into the future: Predicting future person activities and locations in videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 5718–5727.
- [157] C. Lichtenthäler, T. Lorenzy, and A. Kirsch. 2012. Influence of legibility on perceived safety in a virtual human-robot path crossing task. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 676–681.
- [158] S. Lo, K. Yamane, and K. Sugiyama. 2019. Perception of pedestrian avoidance strategies of a self-balancing mobile robot. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 1243–1250.
- [159] T. Lorenz, A. Mörtl, B. Vlaskamp, A. Schubö, and S. Hirche. 2011. Synchronization in a goal-directed task: Human movement coordination with each other and robotic partners. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 198–203.
- [160] Matthias Luber and Kai Oliver Arras. 2013. Multi-hypothesis social grouping and tracking for mobile robots. In *Proceedings of Robotics: Science and Systems (RSS)*.
- [161] M. Luber, L. Spinello, J. Silva, and K. O. Arras. 2012. Socially-aware robot navigation: A learning approach. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 902–907.
- [162] M. Luber, G. Tipaldi, and K. Arras. 2010. People tracking with human motion predictions from social forces. In *Proceedings of the IEEE International Conference on Robotics and Automation*.

- [163] E. A. Martinez-Garcia, Ohya Akihisa, and Shin'ichi Yuta. 2005. Crowding and guiding groups of humans by teams of mobile robots. In *Proceedings of the IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO)*. 91–96.
- [164] Massive. 2017. MASSIVE. Retrieved from <http://www.massivesoftware.com/>.
- [165] Christoforos Mavrogiannis, Patrícia Alves-Oliveira, Wil Thomason, and Ross A. Knepper. 2022. Social momentum: Design and evaluation of a framework for socially competent robot navigation. *ACM Trans. Hum.-robot Interact.* 11, 2 (2022).
- [166] Christoforos Mavrogiannis, Valts Blukis, and Ross A. Knepper. 2017. Socially competent navigation planning by deep learning of multi-agent path topologies. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 6817–6824.
- [167] Christoforos Mavrogiannis, Jonathan A. DeCastro, and Siddhartha S. Srinivasa. 2020. Implicit multiagent coordination on unsignalized intersections via multimodal inference enabled by topological braids. *arXiv e-prints*, Article arXiv:2004.05205.
- [168] Christoforos Mavrogiannis, Alena M. Hutchinson, John Macdonald, Patrícia Alves-Oliveira, and Ross A. Knepper. 2019. Effects of distinct robotic navigation strategies on human behavior in a crowded environment. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 421–430.
- [169] Christoforos Mavrogiannis and Ross A. Knepper. 2019. Multi-agent path topology in support of socially competent navigation planning. *Int. J. Robot. Res.* 38, 2-3 (2019), 338–356.
- [170] Christoforos Mavrogiannis, Wil B. Thomason, and Ross A. Knepper. 2018. Social momentum: A framework for legible navigation in dynamic multi-agent environments. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 361–369.
- [171] Christoforos I. Mavrogiannis and Ross A. Knepper. 2020. Multi-agent trajectory prediction and generation with topological invariants enforced by Hamiltonian dynamics. In *Algorithmic Foundations of Robotics XIII*, Marco Morales, Lydia Tapia, Gildardo Sánchez-Ante, and Seth Hutchinson (Eds.). Springer International Publishing, Cham, 744–761.
- [172] Bruce A. Maxwell. 2007. Building robot systems to interact with people in real environments. *Auton. Robots* 22, 4 (May 2007), 353–367.
- [173] R. Mazzon, F. Poiesi, and A. Cavallaro. 2013. Detection and tracking of groups in crowd. In *Proceedings of the IEEE International Conference on Advanced Video and Signal-based Surveillance (AVSS)*. 202–207.
- [174] R. Mead and M. J. Matarić. 2012. A probabilistic framework for autonomous proxemic control in situated and mobile human-robot interaction. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 193–194.
- [175] Ross Mead and Maja J. Matarić. 2016. Perceptual models of human-robot proxemics. In *Experimental Robotics: The 14th International Symposium on Experimental Robotics*, M. Ani Hsieh, Oussama Khatib, and Vijay Kumar (Eds.). Springer International Publishing, Cham, 261–276.
- [176] F. Michaud, C. Côté, D. Létourneau, Y. Brosseau, J. M. Valin, É. Beaudry, C. Raïevsky, A. Ponchon, P. Moisan, P. Lepage, Y. Morin, F. Gagnon, P. Giguère, M. A. Roux, S. Caron, P. Frenette, and F. Kabanza. 2007. Spartacus attending the 2005 AAAI Conference. *Auton. Robots* 22, 4 (2007).
- [177] Stefan Mitsch, Khalil Ghorbal, David Vogelbacher, and André Platzer. 2017. Formal verification of obstacle avoidance and navigation of ground robots. *Int. J. Robot. Res.* 36, 12 (2017), 1312–1340.
- [178] Abdulllah Mohamed, Kun Qian, Mohamed Elhoseiny, and Christian Claudel. 2020. Social-STGCNN: A social spatio-temporal graph convolutional neural network for human trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 14424–14432.
- [179] M. Montemerlo, N. Roy, and S. Thrun. 2003. Perspectives on standardization in mobile robot programming: The Carnegie Mellon Navigation (CARMEN) Toolkit. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Vol. 3. 2436–2441.
- [180] M. Mucientes and W. Burgard. 2006. Multiple hypothesis tracking of clusters of people. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 692–697.
- [181] U. Nagarajan, G. Kantor, and R. L. Hollis. 2009. Human-robot physical interaction with dynamically stable mobile robots. In *Proceedings of the 4th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 205–205.
- [182] Amal Nanavati, Christoforos Mavrogiannis, Kevin Weatherwax, Leila Takayama, Maya Cakmak, and Siddhartha S. Srinivasa. 2021. Modeling human helpfulness with individual and contextual factors for robot planning. In *Proceedings of Robotics: Science and Systems (RSS)*.
- [183] A. Nanavati, X. Z. Tan, J. Connolly, and A. Steinfeld. 2019. Follow the robot: Modeling coupled human-robot dyads during navigation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 3836–3843.
- [184] Amal Nanavati, Nick Walker, Lee Taber, Christoforos Mavrogiannis, Leila Takayama, Maya Cakmak, and Siddhartha Srinivasa. 2022. Not all who wander are lost: A localization-free system for in-the-wild mobile robot deployments. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 422–431.

- [185] Rahul Narain, Abhinav Golas, Sean Curtis, and Ming C. Lin. 2009. Aggregate dynamics for dense crowd simulation. *ACM Trans. Graph. (Proc. SIGGRAPH Asia)* 28, 5 (2009), 122:1–122:8.
- [186] Nicoletta Noceti and Francesca Odone. 2014. Humans in groups: The importance of contextual information for understanding collective activities. *Pattern Recog.* 47, 11 (2014), 3535–3551.
- [187] I. R. Nourbakhsh, C. Kunz, and T. Willeke. 2003. The mobot museum robot installations: A five year experiment. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Vol. 4. 3636–3641.
- [188] Jean Oh, Thomas M. Howard, Matthew R. Walter, Daniel Barber, Menglong Zhu, Sangdon Park, Arne Suppe, Luis Navarro-Serment, Felix Duvallat, Abdeslam Boularias, Oscar Romero, Jerry Vinokurov, Terence Keegan, Robert Dean, Craig Lennon, Barry Bodt, Marshal Childers, Jianbo Shi, Kostas Daniilidis, Nicholas Roy, Christian Lebiere, Martial Hebert, and Anthony Stentz. 2017. Integrated intelligence for human-robot teams. In *Proceedings of the International Symposium on Experimental Robotics (ISER)*, Dana Kulić, Yoshihiko Nakamura, Oussama Khatib, and Gentiane Venture (Eds.). Springer International Publishing, Cham, 309–322.
- [189] Billy Okal and Timm Linder. 2014. PedSim-ROS. Retrieved from https://github.com/srl-freiburg/pedsim_ros.
- [190] Jan Ondřej, Julien Pettré, Anne-Hélène Olivier, and Stéphane Donikian. 2010. A synthetic-vision based steering approach for crowd simulation. *ACM Trans. Graph.* 29, 4 (2010), 1–9.
- [191] Elena Pacchierotti, Henrik I. Christensen, and Patric Jensfelt. 2006. Evaluation of passing distance for social robots. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 315–320.
- [192] A. K. Pandey and R. Alami. 2009. A step towards a sociable robot guide which monitors and adapts to the person's activities. In *Proceedings of the International Conference on Advanced Robotics*. 1–8.
- [193] A. K. Pandey and R. Alami. 2010. A framework towards a socially aware mobile robot motion in human-centered dynamic environment. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. 5855–5860. DOI: <https://doi.org/10.1109/IROS.2010.5649688>
- [194] Sébastien Paris, Julien Pettré, and Stéphane Donikian. 2007. Pedestrian reactive navigation for crowd simulation: A predictive approach. *Comput. Graph. Forum* 26, 3 (2007), 665–674.
- [195] Chonhyon Park, Jan Ondřej, Max Gilbert, Kyle Freeman, and Carol O'Sullivan. 2016. HI Robot: Human intention-aware robot planning for safe and efficient navigation in crowds. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 3320–3326.
- [196] Jong Jin Park, Collin Johnson, and Benjamin Kuipers. 2012. Robot navigation with model predictive equilibrium point control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 4945–4952.
- [197] Sachin Patil, Jur Van Den Berg, Sean Curtis, Ming C. Lin, and Dinesh Manocha. 2010. Directing crowd simulations using navigation fields. *IEEE Trans. Visualiz. Comput. Graph.* 17, 2 (2010), 244–254.
- [198] S. Pellegrini, A. Ess, K. Schindler, and L. Van Gool. 2009. You'll never walk alone: Modeling social behavior for multi-target tracking. In *Proceedings of the IEEE International Conference on Computer Vision*. 261–268.
- [199] Stefano Pellegrini, Andreas Ess, and Luc Van Gool. 2010. Improving data association by joint modeling of pedestrian trajectories and groupings. In *Computer Vision—ECCV 2010*, Kostas Daniilidis, Petros Maragos, and Nikos Paragios (Eds.). Springer Berlin, 452–465.
- [200] C. Pérez-D'Arpino and J. A. Shah. 2015. Fast target prediction for human-robot manipulation using time series classification. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 6175–6182.
- [201] N. Pradhan, T. Burg, and S. Birchfield. 2011. Robot crowd navigation using predictive position fields in the potential function framework. In *Proceedings of the American Control Conference (ACC)*. 4628–4633.
- [202] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Eric Berger, Rob Wheeler, and Andrew Ng. 2009. ROS: An open-source robot operating system. In *Proceedings of the Workshop on Open-Source Robotics, IEEE International Conference on Robotics and Automation (ICRA)*.
- [203] I. Rae, L. Takayama, and B. Mutlu. 2013. The influence of height in robot-mediated communication. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 1–8.
- [204] John H. Reif and Stephen R. Tate. 1993. The complexity of N-body simulation. In *Automata, Languages and Programming*, Andrzej Lingas, Rolf Karlsson, and Svante Carlsson (Eds.). Springer Berlin, 162–176.
- [205] Craig W., Reynolds. 1999. Steering behaviors for autonomous characters. In *Proceedings of the Game Developers Conference*, Vol. 1999. Citeseer, 763–782.
- [206] Nicholas Rhinehart, Rowan McAllister, Kris Kitani, and Sergey Levine. 2019. PRECOG: PREdiction conditioned on goals in visual multi-agent settings. In *Proceedings of the International Conference on Computer Vision*. 2821–2830.
- [207] Nicholas Rhinehart, Rowan McAllister, and Sergey Levine. 2018. Deep imitative models for flexible inference, planning, and control. *arXiv e-prints*, Article arXiv:1810.06544.
- [208] Kevin W. Rio, Gregory C. Dachner, and William H. Warren. 2018. Local interactions underlying collective motion in human crowds. *Proc. Roy. Societ. B: Biolog. Sci.* 285, 1878 (2018), 20180611.

- [209] Jorge Rios-Martinez, Anne Spalanzani, and Christian Laugier. 2015. From proxemics theory to socially-aware navigation: A survey. *Int. J. Soc. Robot.* (2015).
- [210] Junha Roh, Christoforos Mavrogiannis, Rishabh Madan, Dieter Fox, and Siddhartha S. Srinivasa. 2020. Multimodal trajectory prediction via topological invariance for navigation at uncontrolled intersections. In *Proceedings of the Conference on Robot Learning*.
- [211] Stephanie Rosenthal, Manuela Veloso, and Anind K. Dey. 2011. Learning accuracy and availability of humans who help mobile robots. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 1501–1506.
- [212] N. Roy, W. Burgard, D. Fox, and S. Thrun. 1999. Coastal navigation-mobile robot navigation with uncertainty in dynamic environments. *IEEE Intell. Syst.* (1999).
- [213] N. Roy and S. Thrun. 1999. Online self-calibration for mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Vol. 3. 2292–2297.
- [214] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M. Kitani, Dariu M. Gavrila, and Kai O. Arras. 2019. Human motion trajectory prediction: A survey. *arXiv preprint arXiv:1905.06113* (2019).
- [215] Priscila Saboia and Siome Goldenstein. 2012. Crowd simulation: Applying mobile grids to the social force model. *Vis. Comput.* 28, 10 (2012), 1039–1048.
- [216] Amir Sadeghian, Vineet Kosaraju, Ali Sadeghian, Noriaki Hirose, Hamid Rezaatofghi, and Silvio Savarese. 2019. Sophie: An attentive GAN for predicting paths compliant to social and physical constraints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1349–1358.
- [217] M. Saerbeck and C. Bartneck. 2010. Perception of affect elicited by robot motion. In *Proceedings of the 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 53–60. DOI : <https://doi.org/10.1109/HRI.2010.5453269>
- [218] Takeshi Sakuma, Tomohiko Mukai, and Shigeru Kuriyama. 2005. Psychological model for animating crowded pedestrians. *Comput. Anim. Virt. Worlds* 16, 3-4 (2005), 343–351.
- [219] Satoru Satake, Takayuki Kanda, Dylan F. Glas, Michita Imai, Hiroshi Ishiguro, and Norihiro Hagita. 2009. How to approach humans? Strategies for social robots to initiate interaction. In *Proceedings of the ACM/IEEE International Conference on Human Robot Interaction (HRI)*. 109–116.
- [220] S. Satake, T. Kanda, D. F. Glas, M. Imai, H. Ishiguro, and N. Hagita. 2013. A robot that approaches pedestrians. *IEEE Trans. Robot.* 29, 2 (2013), 508–524.
- [221] Adarsh Jagan Sathyamoorthy, Utsav Patel, Moumita Paul, Nithish K. Sanjeev Kumar, Yash Savle, and Dinesh Manocha. 2022. CoMet: Modeling group cohesion for socially compliant robot navigation in crowded scenes. *IEEE Robot. Autom. Lett.* 7, 2 (2022), 1008–1015. DOI : <https://doi.org/10.1109/LRA.2021.3135560>
- [222] Kristin E. Schaefer, Jean Oh, Derya Aksaray, and Daniel Barber. 2019. Integrating context into artificial intelligence: Research from the robotics collaborative technology alliance. *AI Mag.* 40, 3 (2019), 28–40.
- [223] J. Schulte, C. Rosenberg, and S. Thrun. 1999. Spontaneous, short-term interaction with mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Vol. 1. 658–663.
- [224] Francesco Setti, Chris Russell, Chiara Bassetti, and Marco Cristani. 2015. F-Formation detection: Individuating free-standing conversational groups in images. *PLoS One* 10, 5 (05 2015), 1–26.
- [225] M. Shiomi, T. Kanda, S. Koizumi, H. Ishiguro, and N. Hagita. 2007. Group attention control for communication robots with Wizard of OZ approach. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 121–128.
- [226] Masahiro Shiomi, Francesco Zanlungo, Kotaro Hayashi, and Takayuki Kanda. 2014. Towards a socially acceptable collision avoidance for a mobile robot navigating among pedestrians using a pedestrian model. *Int. J. Soc. Robot.* 6, 3 (2014), 443–455.
- [227] R. Siegwart, K. Arras, S. Bouabdallah, D. Burnier, G. Froidevaux, B. Jensen, A. Lorotte, L. Mayor, M. Meisser, R. Pigué, G. Ramel, G. Terrien, and N. Tomatis. 2003. Robox at expo.02: A large scale installation of personal robots. *Robot. Auton. Syst.* (2003).
- [228] Emrah Akin Sisbot, Luis Felipe Marin-Urias, Rachid Alami, and Thierry Siméon. 2007. A human aware mobile robot motion planner. *IEEE Trans. Robot.* 23, 5 (2007), 874–883.
- [229] J. Snape, J. van den Berg, S. Guy, and D. Manocha. 2011. The hybrid reciprocal velocity obstacle. *IEEE Trans. Robot.* (2011).
- [230] Jamie Snape, Jur Van Den Berg, Stephen J. Guy, and Dinesh Manocha. 2010. Smooth and collision-free navigation for multiple robots under differential-drive constraints. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 4584–4589.
- [231] F. Solera, S. Calderara, and R. Cucchiara. 2016. Socially constrained structural learning for groups detection in crowd. *IEEE Trans. Pattern Anal. Mach. Intell.* 38, 5 (2016), 995–1008. DOI : <https://doi.org/10.1109/TPAMI.2015.2470658>
- [232] SPENCER. 2016. SPENCER. Retrieved from <http://www.spencer.eu>.
- [233] Aaron Steinfeld, Odest Chadwicke Jenkins, and Brian Scassellati. 2009. The Oz of Wizard: Simulating the human for interaction research. In *Proceedings of the ACM/IEEE International Conference on Human Robot Interaction (HRI)*. 101–108.

- [234] L. Sturlaugson and J. W. Sheppard. 2014. Inference complexity in continuous time Bayesian networks. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*. 772–779.
- [235] M. Svenstrup, T. Bak, and J. Andersen. 2010. Trajectory planning for robots in dynamic human environments. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 4293–4298.
- [236] Daniel Szafrir, Bilge Mutlu, and Terrence Fong. 2014. Communication of intent in assistive free flyers. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 358–365.
- [237] L. Tai, J. Zhang, X. Lin, and W. Burgard. 2018. Socially compliant navigation through raw depth inputs with generative adversarial imitation learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 1111–1117.
- [238] L. Takayama and C. Pantofaru. 2009. Influences on proxemic behaviors in human-robot interaction. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 5495–5502.
- [239] Charlie Tang and Russ R. Salakhutdinov. 2019. Multiple futures prediction. In *Proceedings of the International Conference on Neural Information Processing Systems (NeurIPS)*, Vol. 32. 15424–15434.
- [240] Angelique Taylor, Darren M. Chan, and Laurel D. Riek. 2020. Robot-centric perception of human groups. *ACM Trans. Hum.-robot Interact.* 9, 3 (2020), 1–21.
- [241] Andrea Thomaz, Guy Hoffman, and Maya Cakmak. 2016. Computational human-robot interaction. *Found. Trends Robot.* 4, 2-3 (2016), 105–223.
- [242] S. Thompson, T. Horiuchi, and S. Kagami. 2009. A probabilistic model of human motion and navigation intent for mobile robot path planning. In *Proceedings of the International Conference on Autonomous Robots and Agents (ICARA)*. 663–668.
- [243] Sebastian Thrun, Michael Beetz, Maren Bennewitz, Wolfram Burgard, Armin B. Cremers, Frank Dellaert, Dieter Fox, Dirk Hähnel, Chuck Rosenberg, Nicholas Roy, Jamieson Schulte, and Dirk Schulz. 2000. Probabilistic algorithms and the interactive museum tour-guide robot MINERVA. *Int. J. Robot. Res.* 19, 11 (2000), 972–999.
- [244] S. Thrun, M. Bennewitz, W. Burgard, A. B. Cremers, F. Dellaert, D. Fox, D. Hähnel, C. Rosenberg, N. Roy, J. Schulte, and D. Schulz. 1999. MINERVA: A second-generation museum tour-guide robot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Vol. 3. 1999–2005.
- [245] Sebastian Thrun and Tom M. Mitchell. 1995. Lifelong robot learning. *Robot. Auton. Syst.* 15, 1 (1995), 25–46.
- [246] S. Thrun, J. Schulte, and C. Rosenberg. 2000. Interaction with mobile robots in public places. *IEEE Intell. Syst.* (2000), 7–11.
- [247] Elena Torta, Raymond H. Cuijpers, James F. Juola, and David van der Pol. 2011. Design of robust robotic proxemic behaviour. In *Proceedings of the International Conference on Social Robotics (ICSR)*. Berlin, 21–30.
- [248] P. Trautman and A. Krause. 2010. Unfreezing the robot: Navigation in dense interacting crowds. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 797–803.
- [249] P. Trautman, J. Ma, R. M. Murray, and A. Krause. 2015. Robot navigation in dense human crowds: Statistical models and experimental studies of human robot cooperation. *Int. J. Robot. Res.* 34, 3 (2015), 335–356.
- [250] Peter Trautman and Karankumar Patel. 2020. Real time crowd navigation from first principles of probability theory. *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*. 459–467.
- [251] Adrien Treuille, Seth Cooper, and Zoran Popović. 2006. Continuum crowds. *ACM Trans. Graph.* 25, 3 (2006), 1160–1168.
- [252] Xuan-Tung Truong and Dung Ngo Ngo. 2017. Toward socially aware robot navigation in dynamic and crowded environments: A proactive social motion model. *IEEE Trans. Autom. Sci. Eng.* 14, 4 (2017), 1743–1760.
- [253] Xuan-Tung Truong and Trung-Dung Ngo. 2016. Dynamic social zone based mobile robot navigation for human comfortable safety in social environments. *Int. J. Soc. Robot.* 8, 5 (5 2016), 663–684.
- [254] Chieh-En Tsai and Jean Oh. 2020. A generative approach for socially compliant navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 2160–2166.
- [255] Nathan Tsoi, Mohamed Hussein, Olivia Fugikawa, J. D. Zhao, and Marynel Vázquez. 2020. SEAN-EP: A platform for collecting human feedback for social robot navigation at scale. *arXiv preprint arXiv:2012.12336* (2020).
- [256] Annemarie Turnwald and Dirk Wollherr. 2019. Human-like motion planning based on game theoretic decision-making. *Int. J. Soc. Robot.* 11, 1 (2019), 151–170.
- [257] UCrowd. 2021. UCrowd. Retrieved from <https://www.ucrowds.com/>.
- [258] Branislav Ulicny, Pablo de Heras Ciechowski, and Daniel Thalmann. 2004. Crowdbush: Interactive authoring of real-time crowd scenes. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. 243–252.
- [259] V. V. Unhelkar, C. Pérez-D’Arpino, L. Stirling, and J. A. Shah. 2015. Human-robot co-navigation using anticipatory indicators of human walking motion. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 6183–6190.
- [260] J. Šochman and D. C. Hogg. 2011. Who knows who—Inverting the Social Force Model for finding groups. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops*. 830–837.

- [261] Jur van den Berg, Stephen J. Guy, Ming C. Lin, and Dinesh Manocha. 2009. Reciprocal n -body collision avoidance. In *Proceedings of the International Symposium on Robotics Research (ISRR)*. 3–19.
- [262] Jur van den Berg, Stephen J. Guy, Jamie Snape, Ming C. Lin, and Dinesh Manocha. 2016. RVO2 Library: Reciprocal Collision Avoidance for Real-time Multi-agent Simulation. Retrieved from <https://gamma.cs.unc.edu/RVO2/>.
- [263] Jur van den Berg, Ming Lin, and Dinesh Manocha. 2008. Reciprocal velocity obstacles for real-time multi-agent navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 1928–1935.
- [264] Jur van den Berg, Sachin Patil, Jason Sewall, Dinesh Manocha, and Ming Lin. 2008. Interactive navigation of multiple agents in crowded environments. In *Proceedings of the ACM/SIGGRAPH Symposium on Interactive 3D Graphics and Games (i3D)*. 139–147.
- [265] Dizan Vasquez, Billy Okal, and Kai O. Arras. 2014. Inverse reinforcement learning algorithms and features for robot navigation in crowds: An experimental comparison. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 1341–1346.
- [266] Marynel Vázquez, Elizabeth J. Carter, Jo Ana Vaz, Jodi Forlizzi, Aaron Steinfeld, and Scott E. Hudson. 2015. Social group interactions in a role-playing game. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI) Extended Abstracts*. 9–10.
- [267] M. Vazquez, A. Steinfeld, and S. E. Hudson. 2015. Parallel detection of conversational groups of free-standing people and tracking of their lower-body orientation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 3010–3017.
- [268] Araceli Vega, Ramón Cintas, Luis J. Manso, Pablo Bustos, and Pedro Núñez. 2020. Socially-accepted path planning for robot navigation based on social interaction spaces. In *Robot 2019: Fourth Iberian Robotics Conference*, Manuel F. Silva, José Luís Lima, Luís Paulo Reis, Alberto Sanfeliu, and Danilo Tardioli (Eds.). Springer International Publishing, Cham, 644–655.
- [269] Anirudh Vemula, Katharina Muelling, and Jean Oh. 2017. Modeling cooperative navigation in dense human crowds. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 1685–1692.
- [270] A. Vemula, K. Muelling, and J. Oh. 2018. Social attention: Modeling attention in human crowds. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 4601–4607.
- [271] J. Vroon, M. Jooze, M. Lohse, J. Kolkneier, J. Kim, K. Truong, G. Englebienne, D. Heylen, and V. Evers. 2015. Dynamics of social positioning patterns in group-robot interactions. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 394–399.
- [272] Nick Walker, Christoforos Mavrogiannis, Siddhartha Srinivasa, and Maya Cakmak. 2022. Influencing behavioral attributions to robot motion during task execution. In *Proceedings of the 5th Conference on Robot Learning (Proceedings of Machine Learning Research)*, Vol. 164. 169–179.
- [273] Allan Wang, Christoforos Mavrogiannis, and Aaron Steinfeld. 2021. Group-based motion prediction for navigation in crowded environments. In *Proceedings of the Conference on Robot Learning (CoRL)*.
- [274] A. Wang and A. Steinfeld. 2020. Group split and merge prediction with 3D convolutional networks. *IEEE Robot. Autom. Lett.* 5, 2 (2020), 1923–1930.
- [275] William H. Warren. 2006. The dynamics of perception and action. *Psychol. Rev.* 113, 2 (2006), 358–389.
- [276] Jarosław Waś, Bartłomiej Gudowski, and Paweł J. Matuszyk. 2006. Social distances model of pedestrian dynamics. In *Cellular Automata*. Springer Berlin, 492–501.
- [277] Eva Wiese, Agnieszka Wykowska, Jan Zwickel, and Hermann J. Müller. 2012. I see what you mean: How attentional selection is shaped by ascribing intentions to others. *PLoS One* 7, 9 (2012), 1–7.
- [278] Jerry S. Wiggins. 1996. *The Five-factor Model of Personality: Theoretical Perspectives*. Guilford Press.
- [279] Nicholas H. Wolfinger. 1995. Passing moments: Some social dynamics of pedestrian interaction. *J. Contemp. Ethnogr.* 24, 3 (1995), 323–340.
- [280] Takuma Yagi, Karttikeya Mangalam, Ryo Yonetani, and Yoichi Sato. 2018. Future person localization in first-person videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 7593–7602.
- [281] Xinjie Yao, Ji Zhang, and Jean Oh. 2019. Autonomous human-aware navigation in dense crowds. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- [282] Barbara Yersin, Jonathan Maim, Julien Pettré, and Daniel Thalmann. 2009. Crowd patches: Populating large-scale virtual environments for real-time applications. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (i3D)*. 207–214.
- [283] Zeynep Üçel, Francesco Zanlungo, Tetsushi Ikeda, Takahiro Miyashita, and Norihiro Hagita. 2013. Deciphering the crowd: Modeling and identification of pedestrian group motion. *Sensors* 13, 1 (2013), 875–897.
- [284] Matteo Zanotto, Loris Bazzani, Marco Cristani, and Vittorio Murino. 2012. Online Bayesian nonparametrics for group detection. In *Proceedings of the British Machine Vision Conference (BMVC)*. 111.1–111.12.
- [285] H. Zender, P. Jensfelt, and G. M. Kruijff. 2007. Human- and situation-aware people following. In *Proceedings of the IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*. 1131–1136.

- [286] Pu Zhang, Wanli Ouyang, Pengfei Zhang, Jianru Xue, and Nanning Zheng. 2019. SR-LSTM: State refinement for LSTM towards pedestrian trajectory prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 12085–12094.
- [287] Dapeng Zhao and Jean Oh. 2021. Noticing motion patterns: Temporal CNN with a novel convolution operator for human trajectory prediction. *IEEE Robot. Autom. Lett.* 6, 2 (2021), 628–634.
- [288] Brian D. Ziebart, Nathan Ratliff, Garratt Gallagher, Christoph Mertz, Kevin Peterson, J. Andrew Bagnell, Martial Hebert, Anind K. Dey, and Siddhartha S. Srinivasa. 2009. Planning-based prediction for pedestrians. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 3931–3936.

Received 26 March 2021; revised 17 October 2022; accepted 5 January 2023