1 Introduction

This chapter provides an overview of the research domain. We begin by outlining the context of autonomous mobile robot deployment, followed by the specific problem statement, the research questions addressed in this thesis, and the main scientific contributions.

(a)

(b)

(c)

(d)

(e)

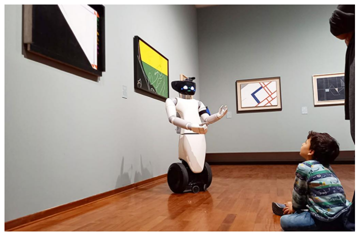

(f)

Figure 1.1: Deployment of autonomous mobile robots (AMRs) in various domains. (a) Spot, developed by Boston Dynamics, is used for inspection tasks at a BASF chemical plant. (b) BellaBot, developed by PUDU Robotics, autonomously delivers food to designated tables and returns empty dishes. (c) Ona, one of the first autonomous delivery robots being tested in a real urban environment in Europe. (d) ABB integrates advanced technologies into its mobile platforms enable safe navigation in dynamic environments. (e) Cruzr, developed by UBTECH, was deployed during the COVID-19 pandemic in Rwanda to assist with fever screening at airports. (f) R1, developed by the Istituto Italiano di Tecnologia (IIT), serves as a tour guide in an art museum in Turin. Image sources and credits: see Appendix A.6

1.1 Research Context

With rapid advancements in sensing, embedded computation, and robotic autonomy, autonomous mobile robots (AMRs) are no longer confined to factories or labs. Today, they are actively deployed in diverse domains such as industrial inspection (Fig. 1.1a), hospitality (Fig. 1.1b), logistics and delivery (Fig. 1.1c), lab automation (Fig. 1.1d), public service (Fig. 1.1e) and cultural institutions (Fig. 1.1f). Statistics show a growing trend in the global AMR market. For example, the International Federation of Robotics Report 2024 shows that service robots have seen continuous growth, with 205,000 new units shipped, 30% of which are AMRs [Int24]. Additionally, Statista shows that service robotics accounted for the largest share of the global robotics market in 2023, generating US$35.45 billion, or 77% of total revenue. The market is projected to grow at a compound annual growth rate of 10% until 2029 [Sta25].

Today, AMRs are equipped with well-established motion planning and control methods that enable them to navigate environments and execute assigned tasks efficiently. For example, warehouse AGVs and food-serving robots typically plan time-efficient paths in static or semi-static environments. In addition, they use onboard sensors, such as Lidars and bumpers, to detect unexpected obstacles or humans, and respond by stopping or altering their trajectory. While these reactive behaviors ensure safety, the overly conservative behaviors, such as unnecessary stops, can compromise operational efficiency. Moreover, humans may feel uncomfortable if the robot is getting too close or moving unnaturally, even in the absence of collisions.

Consequently, the seamless deployment of AMRs in dynamical, socially interactive environments remains both an active area of research and a key focus of real-world robotics deployments. This need has driven growing interest in *social robots* and the field of *Socially Aware Navigation* ***(SAN)***, which shifts the paradigm from treating humans as mere obstacles to integrating human-centric factors into autonomous robot motion planning. The goal of SAN research is to develop motion planning strategies that allow AMRs to navigate safely, respectfully, and effectively in shared human environments, thereby enhancing public trust and broader acceptance of these technologies.

1.2 Socially Aware Navigation

What is socially aware navigation? Based on discussions from the 2022 Social Navigation Symposium, researchers proposed a crisp definition of social robots [FPDL+25]: **A socially navigating robot acts and interacts with humans or other robots, achieving its navigation goals while modifying its behavior so**

the experience of agents around the robot is not degraded or is even enhanced. To fulfill this, a social robot must be capable of understanding and adapting to the behavior of other agents, including humans and other robots. In the context of navigation, this involves adjusting its trajectory in accordance with social norms or by interpreting and responding to the needs of surrounding agents. SAN is generally defined by adherence to key principles such as **safety, comfort**, and **sociability** [KPAK13]. Recently, these principles were expanded into a broader set of eight, capturing more detailed aspects of socially aware behavior of AMRs [FPDL+25]. This reflects the growing depth and maturity of SAN research in recent years. In this work, the principles of SAN are broadly summarized as **Safety, Comfort, Legibility** and **Sociability**, from the perspective of motion planning:

Safety: The most fundamental requirement. A robot must reliably avoid collisions with nearby humans, other robots, and the environment.

Comfort: Humans should feel comfortable when interacting with robots. As noted in [KPAK13], comfort is not only about maintaining a physical distance, but also about minimizing perceived annoyance or stress. For example, inspired by the Proxemics Theory [HH66], a robot should respect personal zones rather than simply avoiding collision at a safety distance.

Legibility: At the low-level motion planning stage, the robot's movement should be easily interpretable by nearby agents. This involves making its goal and intention predictable, such as maintaining a consistent direction of movement or using gazing cues, if applicable. Smooth and non-erratic trajectories also help avoid confusion and distraction.

Sociability: At the high-level decision-making stage, humans typically resolve spatial conflicts by following shared social conventions. Therefore, robots should likewise follow these norms to enable cooperative interactions without relying on explicit communication methods such as sounds or lights. Since social norms can vary across contexts, robots must also be capable of inferring others' intentions and respond proactively.

1.3 Research Challenges in Socially Aware Navigation

While the principles of Safety, Comfort, and Sociability (defined above) provide a clear theoretical goal, implementing them in real-world environments presents significant technical challenges, requiring careful design of control frameworks that can account for both physical constraints and social context. This section outlines the specific control and communication challenges that currently prevent standard AMRs from achieving true social awareness.

One possible approach to socially aware motion planning is to formulate it as a distributed control problem in a multi-agent system, where both robots and humans pursue their individual goals while interacting through implicit communication [C+17, MAOTK22]. In conventional distributed control systems, the sub-systems (or agents) compute their control policies based on their local measurements and the information exchange between them. However, implementing explicit communication between humans and robots, such as through wearable devices [COS20], remains impractical in many real-world scenarios. Instead, in SAN, robots typically rely on onboard sensors to detect and track human motion, from which they predict human trajectories or infer intent, and often leverage social conventions to implicitly model interactions.

A comprehensive review of the state of the art is detailed in Chapter 2. However, despite the significant progress highlighted in that review, substantial gaps remain between current SAN capabilities and the requirements for real-world deployment.

1.4 Contributions and Outline

In response to the challenges outlined in Section 1.3, the primary goal of this thesis is to develop a novel motion planning framework that bridges the gap between theoretical social compliance and real-world deployment. To achieve this, we propose extending the reactive Dynamical System Modulation (DSM) framework with proactive social inference modules, specifically utilizing the Social Force Model (SFM) and opinion dynamics (OD) to enable robots to navigate safely, comfortably and legibly without explicit communication.

Consequently, the main contributions of this thesis are as follows:

- **Explicit Modeling of Social Norms:** We leverage the analytical nature of the DSM to explicitly encode social factors, such as the shape of personal zones and the robot speed profile, into motion planning. This allows the robot to respect spatial and speed-based social norms without relying on opaque, learning-based "black box" policies.
- **Framework 1: Social Inference Modulated DSM (SIM-DSM, Chapter 4):** We introduce an SFM-inspired approach that enables the robot to infer appropriate avoidance directions by analyzing the spatial relationship between agents and the human's passing intent.
- **Framework 2: Opinion Dynamics-based DSM (OD-DSM, Chapter 6):** We propose a novel formulation that embeds OD directly into the DSM eigenvalue design. This mechanism enables the robot to proactively adjust its motion,

resulting in smoother and more socially compatible navigation compared to purely reactive methods.

- **Comprehensive Evaluation of Social Awareness**: We validate the proposed frameworks through extensive simulations and real-world user studies. These experiments quantify the impact of proactive behaviors on human-perceived comfort, offering new empirical insights into the correlation between robot motion dynamics and social acceptance.

The thesis is organized as follows:

Chapter 2 reviews the state of the art in SAN, with a focus on motion planning methodologies, the integration of social norms, and evaluation metrics. It concludes with a comparative analysis that highlights the specific distinctions between this thesis and existing works.

Chapter 3 outlines the foundational concepts of this thesis. First, it analyzes the fundamental requirements of SAN and reformulates them within the context of a robot motion planning framework. Second, it introduces the preliminaries of DSM, which serves as the foundational motion planning method throughout the thesis.

Chapter 4 presents a novel SAN motion planner, SIM-DSM, which integrates an SFM-inspired module to infer suitable avoidance directions in human-robot interactions. Inspired by the idea of SFM, we propose the Social Inference Model (SIM), which incorporates social cues and human intent estimation. This model is combined additively with the robot's original dynamics to generate socially compliant trajectories. The approach is validated in simulation by comparing its performance against baseline collision avoidance methods within a pedestrian simulation environment.

Chapter 5 details the real-world validation of the planner introduced in Chapter 4. The evaluation includes two key aspects. First, the planner's behavior is empirically validated under varying controller parameters and intent estimation settings. Second, a user study assesses the impact of controller variations on human-perceived comfort, providing initial insights into factors influencing social awareness. The limitations identified in this experiment motivate the alternative design introduced in the subsequent chapter.

Based on the findings in previous chapters, Chapter 6 introduces OD-DSM, a novel DSM-based planning framework that incorporates OD for human intent estimation. Compared to the additive extension design in Chapter 4, the novelty in this approach lies in the integration of proactive behavior into the DSM's eigenvalue formulation, resulting in smoother and more adaptive robot trajectories.

Chapter 7 extends the experimental design of Chapter 5 to validate the controller proposed in Chapter 6. In addition to assessing its standalone performance and social compliance, we perform a comparative evaluation between the two developed controllers, analyzing both objective and subjective outcomes.

Chapter 8 summarizes the main contributions and findings of the thesis. It also discusses limitations of the current setup and outlines directions for future research to further advance SAN frameworks.

In Appendix A, we present the statistics of the participant survey and the mobile platform design.

2 Related Work

The goal of this thesis is to develop a socially aware motion planner capable of encoding social norms as well as to explicitly evaluate the impact of these social factors for future SAN studies. Guided by the research challenges and objectives outlined in Section 1.3, we organize the literature review around the following aspects:

- Section 2.1 analyzes current research trends and identifies open questions in the field by synthesizing insights from survey works over the last decade.
- Section 2.2 reviews motion planning methods in SAN from the perspective of planner frameworks and algorithm design.
- Section 2.3 discusses the modeling of human-centric factors in motion planning.
- Section 2.4 examines recent approaches to the experimental evaluation of SAN.

Finally, in Section 2.5, we highlight the novelty of this work through a comparative analysis with state-of-the-art methods.

2.1 Survey Works on SAN

Early deployment of AMRs in human environments relied on classical motion planning methods and modeled humans as dynamic obstacles for collision avoidance. We refer the reader to [KPAK13, Sin22, MBW+23] for an overview of SAN research prior to 2010. Since the 2010s, the field has progressed to incorporate human-centric factors, such as modeling of human-robot interaction, human trajectory prediction and intent estimation.

The evolution of SAN in the last decade has been extensively documented in several influential survey papers [KPAK13, RMSL15, CKG17, MBW+23, SBBM+24, FPDL+25]. Among these, the survey [KPAK13] is widely regarded as a foundational work in SAN research over the last decade, which introduced three core features that have since influenced the field's development: **comfort, naturalness** and **sociability**. The survey [RMSL15] extended the scope of SAN beyond distance-based proxemics, and proposed that the robot should consider social cues and social conventions, such as gaze direction, human pose, or group behavior. The survey [CKG17] introduced the concept of 'social mapping', which focuses on interpreting sensor-derived social signals and encoding them into spatial representations that support socially aware motion planning.

More recently, [MBW+23] highlighted the importance of integrating human-robot interaction into SAN. They categorized planning strategies into decoupled and coupled prediction–planning paradigms, depending on whether the robot considers mutual influence during navigation. This paper further reviewed robot behavior design in SAN through three key dimensions: proxemics (spatial preference and norms), intention communication (explicit or implicit signaling of goals), and social space context (environmental or situational cues that influence behavior).

The recent surveys [SBBM+24, FPDL+25] provided a comprehensive review of methodologies developed in SAN research over the past decade. Both works built upon and expanded the three core features in [KPAK13] to a broader spectrum. In particular, [SBBM+24] defined four minimum requirements for SAN: (1) treating humans as special entities and prioritizing safety, (2) minimizing disturbance caused by robot behavior, (3) expressing intent through interpretable actions, and (4) resolving conflicts in a socially compliant manner. Similarly, [FPDL+25] refined and elaborated on existing concepts, introducing dimensions such as legibility, politeness, social competency and proactivity, while maintaining safety and comfort as core concerns.

Despite these advances, SAN remains an active area of research. [SBBM+24] concluded with several key proposals for future work, emphasizing the need for improved human behavior modeling, adaptive intent communication strategies, and comfort-oriented metrics that extend beyond traditional safety considerations. Moreover, all major surveys recognized evaluation as a persistent challenge in SAN. Objective metrics commonly used in the evaluation of traditional robot motion planning, such as minimum distance to obstacles or path efficiency, struggle to capture the subjective nature of human experience. On the other hand, subjective evaluations, while valuable, are heavily influenced by the experimental design and context, limiting their generalizability, reproducibility, and comparability across studies. In addition, evaluations based on participant questionnaires are resource-intensive, requiring human subjects and careful experimental protocols. To address these challenges, [FPDL+25] offers a set of comprehensive guidelines for the design and evaluation of SAN systems.

Building upon the overview of research in SAN, the remainder of this chapter provides a comprehensive overview and categorization of the motion planning methods used in social navigation.

2.2 Motion Planning Methods for Social Navigation

In this work, we focus on the motion planning aspect of SAN. First, we narrow our scope to local motion planning, where we detail analytical methods and compare them with learning-based approaches. Second, we investigate the integration of social awareness into motion planning frameworks. Finally, we discuss evaluation methods in SAN, with a specific emphasis on human-centric factors.

Global and Local Motion Planning

SAN builds upon classical robot motion planning methods, which can be generally categorized into global path planning and local motion planning. Global planners (e.g., Dijkstra's algorithm [Dij59], A* [HNR68]) compute collision-free paths from start to goal positions using a static or semi-static map. The main objectives of global planners include task efficiency (e.g., finding the shortest path) and avoiding dead ends. For example, Dijkstra's algorithm guarantees global optimality by exhaustively exploring all possible routes, while A* improves search efficiency using heuristics, and still yields optimal paths when heuristics are admissible and consistent. However, global planning methods generally assume a static environment, making them less suitable for dynamic settings with moving obstacles such as humans. Therefore, global planners typically work in tandem with a local motion planner, which is responsible for adjusting the global path or trajectory in real time based on sensor data, ensuring safe navigation.

Therefore, a robot navigation framework typically combines both global and local planning: the global planner computes a task-efficient, dead-end-free path, while the local planner handles dynamic obstacle avoidance. In this context, SAN planners are generally developed at the local planning level, where they adapt to human detection and tracking in real-time, and integrate socially aware behaviors into motion generation. We broadly categorize existing SAN approaches at the local planner level into analytical and data-driven methods.

Analytical Methods for Local Motion Planning

Analytical approaches explicitly model and reason about the robot's behavior and its environment, and can be broadly categorized into rule-based, sampling-based, optimization-based, and hybrid techniques.

Rule-based methods for SAN explicitly encode social norms and behavioral rules into a robot's motion planning framework. Notable examples include the Artificial Potential Field (APF) [Kha86] and the Velocity Obstacle (VO) [FS98]. In APF, a robot is typically modeled as a point mass subject to virtual forces, such as attractive forces pulling it toward the target position and repulsive forces pushing it away from obstacles. A well-known application of APF in SAN is the Social Force Model [HM95], which models pedestrians as particles influenced by social forces. Recent works have extended SFM to incorporate human-robot interactions, environmental social cues, and data-driven calibration of force parameters. For example, [FZCS17, AW20] extended SFM to include human-robot interaction forces and used data-driven approaches to calibrate the force parameters. In particular, [FZCS17] further refined the parameters based on user feedback. [LJSP22] extended SFM by considering the non-holonomic kinematic constraint of an E-Scooter, and a velocity-dependent force model to include the influence of relative velocities. The SFM parameters are calibrated using experimental data. While traditional SFM treats pedestrians as position-based agents, [YZCF19] emphasizes the importance of pedestrians' orientation within groups and extended SFM to handle interaction with human groups. [RMK21] introduced context-aware social factors into SFM, such as passing on the left or allowing overtaking on the right, enabling robots to conform to localized social rules.

Although SFM captures human-robot interactions in a clear and interpretable manner, it relies on a one-step sense-perceive-react cycle and does not proactively anticipate humans' future movements, which makes it less suitable for robot motion planning. Besides, due to the force-based modeling, the proxemics theory was reflected only implicitly in robot motion. As a result, it fails to provide theoretical guarantees for collision avoidance, nor does it maintain explicitly defined personal spaces. Hence, some works focus on modeling human-robot interaction using SFM and integrating it into a broader planning framework. For instance, [KTK+22] proposed the inducible SFM (iSFM), which uses extended SFM for human motion prediction, and in cases where a potential deadlock is predicted, the robot applies a real physical force via a robot arm to actively induce human movement and resolve the conflict. The parameters of SFM were determined by manual tuning in simulation. [KCKY22] addressed the challenge of local minima in SFM-based navigation by integrating it with A*, which provides waypoints that guide the robot out of these areas. [PZCBN24] modeled the human personal zone using asymmetric Gaussian functions, and introduced social forces to consider the human-object interaction spaces.

Velocity Obstacle (VO) is another widely used foundation in SAN for dynamic obstacle avoidance. It represents obstacles within the robot's velocity space to enable predictive collision avoidance. Traditional VO-based approaches assume that other agents maintain constant velocities, limiting their effectiveness in dynamic, interactive environments. To address this limitation, the Reciprocal VO (RVO) method [VDBGLM11] extended VO by assuming mutual cooperation among agents, where all agents share responsibility for collision avoidance. This cooperative behavior enables RVO to proactively anticipate future collisions,