Rethinking Legibility in Social Robot Navigation: Impact of Intent Representation and Human Attention

Pranav Goyal
University of Michigan at Ann Arbor
Ann Arbor USA
prgoyal@umich.edu

Andrew Stratton
University of Michigan at Ann Arbor
Ann Arbor USA
arstr@umich.edu

Christoforos Mavrogiannis
University of Michigan at Ann Arbor
Ann Arbor USA
cmavro@umich.edu

Abstract

Legible motion enables humans to anticipate robot behavior during social navigation, but existing approaches largely assume open spaces, static interactions, and fully attentive pedestrians. We study legibility in the ubiquitous and realistic setting of hallway navigation through two user studies. Study 1 ($N = 45$) evaluates how intent should be represented for legible navigation within a model predictive control framework. We find that expressing intent at the interaction level (i.e., passing side) and dynamically adapting it to human motion leads to smoother human trajectories and higher perceived competence than destination-based or non-legible baselines. Study 2 ($N = 45$) examines whether legibility remains beneficial when pedestrians are cognitively distracted. While legible motion still reduced abrupt human motion relative to the non-legible baseline, subjective impressions were less sensitive under distraction. Together, these results demonstrate that legibility is most effective when grounded in immediate interaction objectives and highlight the need to account for attentional variability.

CCS Concepts

• **Computer systems organization → Robotic autonomy**; *Robotic control*; • **Human-centered computing → Laboratory experiments**; *Collaborative interaction*; • **Computing methodologies** → *Cognitive robotics*; *Robotic planning*.

Keywords

Implicit communication, Social robot navigation

ACM Reference Format:
Pranav Goyal, Andrew Stratton, and Christoforos Mavrogiannis. 2026. Rethinking Legibility in Social Robot Navigation: Impact of Intent Representation and Human Attention. In *Companion Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction (HRI Companion '26), March 16–19, 2026, Edinburgh, Scotland, UK*. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3776734.3794499

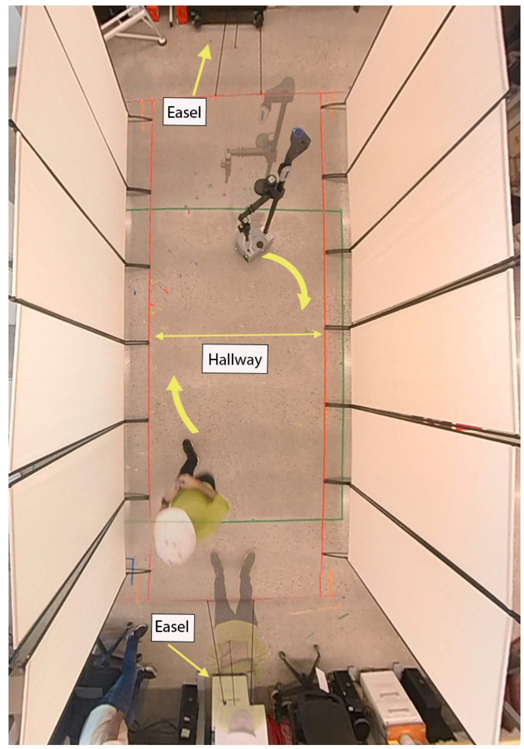

Figure 1: Instance from our study investigating legible social robot navigation in hallways. Footage from our experiments can be found at https://youtu.be/P5O5BwfcUCo.

1 Introduction

Robots moving in pedestrian spaces like the workplace, and public spaces must do more than simply avoid collisions: they must move in ways that people can intuitively interpret. In human–human encounters, pedestrians rely on subtle cues such as gaze, orientation, and small path adjustments to infer each other's intentions and coordinate smooth passages [37]. Robots, however, lack these ingrained social signals. When their motion is difficult to interpret, pedestrians may hesitate, yield unnecessarily, or experience discomfort [23]. This motivates the integration of legibility [7, 15, 22, 33] into the social robot navigation (SRN) control stack.

Legibility, the property of motion that enables an observer to quickly and confidently infer the robot's intent has been extensively studied in human-robot interaction (HRI) [3, 6–9, 17, 22, 31, 33]. Early work established that legibility can improve collaboration in manipulation tasks by reducing ambiguity and building trust. For example, Dragan and Srinivasa [8] showed that legible trajectory planning can make a robot's intended goal easier to infer, and subsequent studies demonstrated that intent-expressive motion enhances perceived safety and efficiency in shared tasks [6, 17, 31]. Later work recognized the value of legibility for any joint HRI activity [15], discussed the impact of important parameters like the observer's viewpoint [24, 33], and proposed frameworks for transferring the benefits of legibility in other domains [17, 22, 35].

ACM ISBN 979-8-4007-2321-6/2026/03
https://doi.org/10.1145/3776734.3794499

Figure 2: Hardware and materials used across our studies.

Despite rich activity in legibility for HRI, several gaps remain. Most prior formulations assume that the observer's inference is over the robot's destinations. This abstraction suits manipulation tasks [3, 7, 31], but it breaks down in navigation: pedestrians do not need to know where others are ultimately headed; they just need to know how to resolve a possible conflict [22, 37]. Additionally, prior work often evaluates legibility via studies involving passive observation of robot motion [7, 14, 32, 33], overlooking that in many HRI contexts, people are both observers and actors who dynamically adapt to the robot's movements. While prior explorations of legibility in navigation domains often emphasize open-space navigation [17, 22, 33], constrained spaces are more common in the real world and more challenging to navigate [29]. Earlier studies on hallway navigation emphasize human-robot proxemics and side-passing conventions [20, 25, 26], but lack a systematic investigation of legibility representations. Finally, prior work on legible motion assumes that humans are fully attentive observers [3, 6–9, 17, 22, 31, 33], yet in many real-world environments, including SRN domains, human attention is often scattered across multiple tasks like texting, speaking on the phone, carrying objects, etc. These tasks act as distractors that may limit the ability of the robot to influence human behavior.

Motivated by these gaps, we revisit legible motion in SRN and formulate two research questions: (Q1) *How does intent representation shape legible SRN?*, (Q2) *How does human distraction impact benefits of legible SRN?* To approach these questions, we embed alternative formulations of legible motion into model predictive control and evaluate them across two studies in a hallway setting. Study 1 ($N = 45$) explores the impact of intent representation on human navigation performance and impressions. Study 2 ($N = 45$) examines the value of legible motion under distraction. Across both studies, we analyze subjective perceptions of competence, comfort, and effort, alongside objective measures of human motion.

2 How Intent Representation Shapes Legible SRN

We present an IRB-approved lab study (HUM00268645) examining how intent representation in legible motion generation affects navigation performance and user impressions (Q1).

2.1 Experiment Design

Experimental setup. The study took place in a hallway of size 2 m×5.8 m, constructed with room dividers inside the lab (see Fig. 1). An easel pad was placed at each hallway end. We used the Hello Robot Stretch 2 mobile platform. An overhead camera (Insta360) recorded RGB video if the user gave consent. A motion capture system (OptiTrack) was used to localize the robot and the users. The full experimental apparatus is shown in Fig. 2.

Procedure. Participants provided informed consent, received task instructions, and completed one practice encounter, followed by five trials (~80 s each) and a brief questionnaire. Sessions concluded with demographics and debriefing, lasted under 40 minutes, and participants were compensated $20.

Task description. Participants performed a mock "factory inspection" task by walking between easel pads at hallway ends to place stickers while the robot navigated in the opposite direction, creating repeated head-on encounters. Each trial used a different navigation algorithm. Turnarounds were synchronized via a gong, and the robot's slower speed (0.33 m/s) was accommodated by placing its turnaround points closer to ensure consistent encounters.

2.2 Conditions

Each participant completed five trials in a within-subjects design, with the robot executing a different navigation strategy in each trial. All strategies used the same state-of-the-art MPC controller [36], with identical safety and efficiency objectives and a Constant Velocity (CV) human-motion predictor from prior work [16, 19, 27, 28, 30, 34]. The only manipulated factors were the robot's expressed intent and whether it was adaptively updated during the encounter.

Goal-based legibility (GL). Here the robot's intent is modeled as the set of the two hallway endpoints, corresponding to the robot's intended destination, similar to prior legibility implementations [7, 17]. Intuitively, this algorithm is favoring actions that consistently communicate the robot's intended destination.

Passing-side legibility (PL). Here the robot's intent consists of artificial subgoals placed to the left and right of the hallway endpoints, effectively creating the notion of passing-side interactions. By steering toward one of these offsets, the robot conveys its intended side of passage. The robot's true intent is assigned at random before the trial and held fixed for each encounter.

Dynamic passing side legibility (DPL). Here, robot's intent is also defined as the notion of passing-side interaction like PL, but the true intent is dynamically adapted at run time. Specifically, the robot selects the passing side associated with the lower predicted probability of being chosen by the human, based on CV predictions of the human's motion.

Social Momentum (SM) [22]. Here robot's intent explicitly consists of passing from the left or right side. Unlike PL and DPL, which encode passing-side intent through artificial offset goals, SM reasons about passing directly in the joint human–robot interaction space. The passing side is determined using the angular momentum of the system: its sign encodes the passing side, while its magnitude reflects confidence [22]. The true robot's intent is dynamically adapted based on the passing preference of the human.

No legibility (NL). The baseline is a standard MPC formulation in which the legibility term is omitted, resulting in purely functional robot navigation with only safety and efficiency objectives.

2.3 Measures

To analyze human motion, we measure human path inefficiency using the Path Irregularity (Human PI) measure [11], defined as the amount of unnecessary turning per unit path length (rad/m). We also measure human path jerkiness as the average human path acceleration over a trial (Human AA).

We also study user impressions collected via questionnaires. We use the *Discomfort* and *Competence* subscales of RoSAS [5], presented in randomized order on 9-point Likert scales. We also use the *Mental, Physical, Temporal, Frustration, Performance*, and *Effort* demand scales from NASA-TLX [12], presented in a 21-point format. To capture perceived goal clarity, we asked users to rate: (*L1*) – "The robot will bump into me in the future" (perceived collision risk); (*L2*) – "I was quickly and accurately able to tell where the robot wants to go" (perceived legibility, based on Dragan et al. [7]), both presented as 7-point scales. Finally, we collected open-form responses to capture insights not covered by the structured scales.

2.4 Hypotheses

We study how different legibility implementations impact navigation performance and impressions by investigating the following:

H1: "Legible algorithms will be more positively perceived and enable higher user performance." We hypothesize that legible algorithms (GL, SM, PL, DPL) will lead to lower acceleration and more regular paths for the users, compared to non-legible algorithms (NL). We further expect legible algorithms to be rated as more competent and comfortable on the RoSAS scale, and to impose lower workload on users as measured by the NASA TLX, in contrast to non-legible algorithms.

H2: "Legibility over the robot's *passing side* will be more positively perceived and enable higher user performance compared to legibility over the robot's goal." We hypothesize that passing side legibility (PL, DPL, SM) will enable faster and more accurate inference of the robot's intent than goal-based legibility (GL), as reflected in responses. Passing side legibility is also expected to yield smoother trajectories (lower acceleration and more regular paths), and to be perceived as more competent, comfortable, and less effortful for users compared to goal-based legibility.

H3: "Dynamically adapting the robot's legibility intent based on user reaction will be more positively perceived and enable higher user performance compared to legibility over a fixed intent." We hypothesize that dynamic adaptation (SM, DPL) will outperform fixed intent legibility (PL) under the same passing side representation. Specifically, dynamic algorithms are expected to produce smoother human motion (lower acceleration and more regular paths), be perceived as more competent and comfortable, and require less effort and workload than fixed intent approaches.

2.5 Analysis

Data from 45 participants recruited from University of Michigan were analyzed ($M_{\text{age}} = 22.42$, $SD = 2.75$; self-reported robotics familiarity $M = 3.68$, $SD = 1.03$) using linear mixed-effects model with random intercepts for participants, fixed effects for algorithm and order, and Benjamini–Hochberg (BH) corrected pairwise contrasts. Discomfort and Competence ratings are summarized in Table 1; workload-related measures, L1 and Human AA, are shown in Fig. 3. No significant effects were observed for Human PI, L2, or the NASA–TLX Performance, Effort, and Temporal Demand.

H1. Human AA was significantly higher under NL than under SM ($p < .01$) and DPL/PL ($p < .05$), indicating jerkier human motion during encounters with a non-legible robot. NL was rated less competent than SM and DPL ($p < .05$) and more likely to collide (L1; $p < .05$). *Interpretation:* Legibility enhances both perceived and objective interaction quality; however, improvements depend on the specific formulation. **H1 is partially supported.**

Table 1: RoSAS ratings (EMM [95% CI]). Different letters indicate significant pairwise differences ($p < .05$, BH corrected). No letters denote no significant differences.

Algorithm	Competence ↑	Discomfort ↓
SM	5.75 [5.19, 6.31]^a	2.52 [2.04, 2.99]
DPL	5.66 [5.11, 6.22]^a	2.30 [1.83, 2.78]^b
PL	5.05 [4.49, 5.61]^b	2.95 [2.48, 3.43]^a
NL	5.05 [4.50, 5.61]^b	2.53 [2.06, 3.01]
GL	4.62 [4.06, 5.17]^b	3.09 [2.62, 3.57]^a

H2. GL consistently underperformed the passing-side strategies: lower competence ($p < .001$ vs. SM/DPL), higher discomfort ($p < .05$ vs. DPL), and the highest Human AA ($p < .01$ vs. SM/DPL/PL). *Interpretation:* Destination-based legibility miscommunicates intent in tight encounters, whereas passing-side intent provides clearer conflict-resolution cues. **H2 is partially supported.**

H3. SM and DPL were rated more competent and less collision-prone than PL ($p < .05$). Human AA differences were mixed and not statistically significant. *Interpretation:* With a passing-side intent representation, dynamic adaptation improves perceived competence and comfort, though objective measures showed no significant differences. **H3 is partially supported.**

Summary. Across all three hypotheses, results indicate that legibility is most effective when (i) intent is expressed at the interaction level (passing side) and (ii) legibility signals adapt to evolving human motion. SM and DPL best satisfied these criteria, yielding the smoothest trajectories and highest perceived competence, whereas GL consistently degraded both performance and impressions.

3 How Human Distraction Impacts Benefits of Legible SRN

In Study 2, we investigate the role of legible motion under distraction (Q2). The distraction factor was implemented between subjects: the no-distraction data came from Study 1, while a new participant cohort performed the same task under distraction.

3.1 Study Design

We used the same task setup as Study 1, where participants performed a mock factory inspection task in hallway scenario.

Distraction task. Participants listened to a narrated passage via earbuds and answered multiple-choice questions on a mobile device (Fig. 4), inducing divided attention consistent with prior work on distracted walking [2, 21]. Passages were standardized and validated via pilot testing to ensure consistent cognitive load.

Procedure. Participants completed three~80 s trials while performing the distraction task and completed the same subjective measures as in Study 1 after each trial. Sessions lasted up to 30 min, and participants were compensated $15.

Conditions. Three navigation algorithms were tested in a within-subjects design: Dynamic Passing-side Legibility (DPL), Social Momentum (SM), and No Legibility (NL). DPL and SM were selected as the strongest legible performers from Study 1, with NL as the non-legible reference.

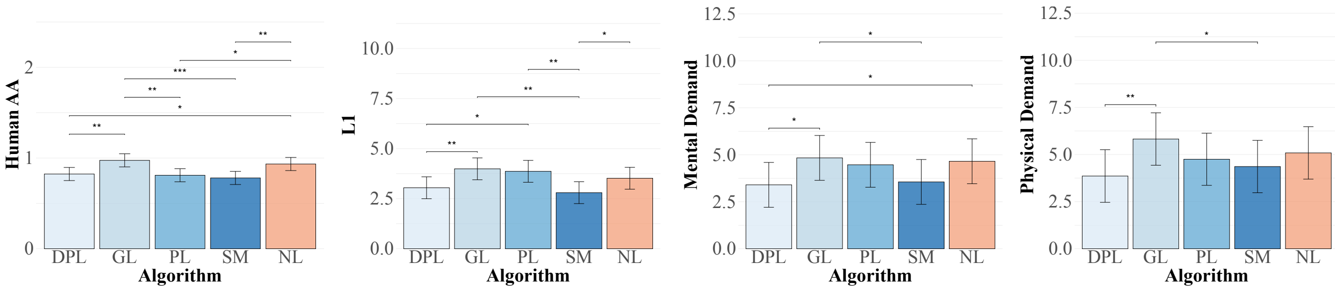

(a) **(b)** **(c)** **(d)**

Figure 3: **Estimated marginal means (95% CI) for Human Average Acceleration** (m/s^2)**,** ***L1*** **(1–7 Likert), Mental Demand (1–21), and Physical Demand (1–21). Bars with asterisks indicate significant pairwise differences:** $p < .001$, $p < .01$, $p < .05$.

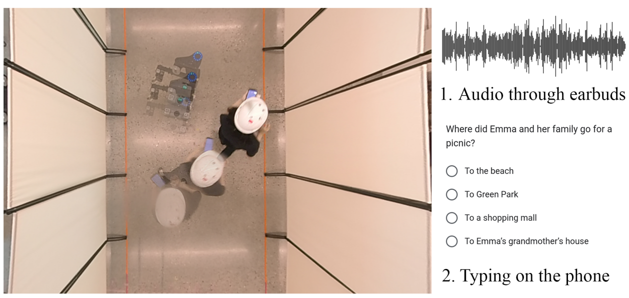

Figure 4: **Experimental setup for the distraction condition. Participants walked while listening to audio passages and answering comprehension questions on a phone, inducing divided attention during head-on encounters with the robot.**

Hypothesis H4: 'Legible algorithms will be more positively perceived and enable higher user performance in the presence of cognitive distraction, compared to respective ratings and performance under no distraction.' We hypothesize that divided attention increases navigational uncertainty, increasing the value of clear robot intent for conflict resolution.

3.2 Analysis

Data from 45 new participants were analyzed ($M_{\text{age}} = 23.55$, $SD = 4.05$; self-reported robotics familiarity $M = 3.40$, $SD = 0.91$). To validate the distraction manipulation, we compared cognitive- and motion-related outcomes across conditions. Participants reported significantly higher mental demand under distraction than no-distraction ($t(241.9) = -8.87$, $p < .001$). Distraction also reduced walking speed (0.96 vs. 1.11 m/s; $t(244.23) = 8.71$, $p < .001$) and increased robot acceleration during encounters (1.70 vs. 1.34 m/s^2), confirming increased navigational difficulty.

Algorithm effects under distraction were modeled using linear mixed-effects regression as in Study 1. Subjective impressions did not differ significantly between algorithms. However, legible motion reduced Human AA: SM yielded significantly lower acceleration than NL (0.803 vs. 0.947, $p < .05$), and DPL (0.887) showed a similar trend. These improvements were comparable in magnitude to those observed in the no-distraction condition.

H4. We tested whether distraction amplified the benefits of legibility by including fixed effects of Algorithm, Distraction, and their interaction. Likelihood-ratio tests showed that the interaction term did not improve model fit, indicating that distraction did not significantly alter algorithm performance on any metric. Thus, **H4 was not supported**: legible algorithms did not produce stronger gains under distraction than in the no-distraction setting.

3.3 Exploratory Insights

Residual variance in competence and discomfort ratings was lower under distraction (competence: 1.24 vs. 1.01; discomfort: 0.71 vs. 0.46), indicating more uniform judgments when attention was divided and reducing the ability to detect algorithm-level differences. In contrast, residual variance in Human AA increased (0.041 vs. 0.072), consistent with more variable walking behavior. Speed-moderation analyses showed that legibility effects intensified at higher walking speeds. Under distraction, SM produced significantly lower Human AA than NL across the 25th, 50th, and 75th percentiles ($p < .05$), with the gap widening as speed increased.

Summary. Legible motion retained its objective benefits under distraction, particularly at higher and typical walking speeds, but subjective impressions were less sensitive due to reduced attentional capacity. These findings suggest that legibility remains valuable when users are inattentive.

4 Discussion

Across two studies, we examined how intent should be represented for legible SRN and whether legibility remains effective under pedestrian distraction. Study 1 showed that widely adopted destination-based formulations [1, 7, 17, 18, 33] increased workload and produced less smooth human motion, whereas dynamically adaptive passing-side intent provided clearer cues for resolving head-on encounters. Study 2 confirmed that these objective benefits persisted under distraction, with stronger effects at typical and higher walking speeds. Overall, the results suggest that legibility in navigation should prioritize real-time interaction coordination over predictability of the robot's final goal.

Both studies were conducted in a controlled hallway with scripted one-on-one encounters, enabling precise comparisons while limiting ecological complexity. Robot speed was capped at 0.33 m/s for safety, consistent with prior constrained-space HRI studies [4, 10, 13, 20], likely reducing collision risk. Future work will evaluate faster platforms and multi-agent interactions in richer crowd settings, and explore online estimation of human attentional state to balance efficiency and legibility alongside contextual factors such as environmental complexity, physical effort, and task urgency.

References

[1] Javad Amirian, Mouad Abrini, and Mohamed Chetouani. 2024. Legibot: Generating Legible Motions for Service Robots Using Cost-Based Local Planners. In *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (ROMAN)*. 461–468.

[2] Hassan Bazzi and Anthony T. Cacace. 2023. Altered gait parameters in distracted walking: a bio-evolutionary and prognostic health perspective on passive listening and active responding during cell phone use. *Frontiers in Integrative Neuroscience* Volume 17 - 2023 (2023).

[3] Christopher Bodden, Daniel Rakita, Bilge Mutlu, and Michael Gleicher. 2018. A flexible optimization-based method for synthesizing intent-expressive robot arm motion. *The International Journal of Robotics Research* 37, 11 (2018), 1376–1394.

[4] John Travis Butler and Arvin Agah. 2001. Psychological Effects of Behavior Patterns of a Mobile Personal Robot. *Autonomous Robots* 10, 2 (2001), 185–202.

[5] Colleen M. Carpinella, Alisa B. Wyman, Michael A. Perez, and Steven J. Stroessner. 2017. The Robotic Social Attributes Scale (RoSAS): Development and Validation. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 254–262.

[6] Anca D. Dragan, Shira Bauman, Jodi Forlizzi, and Siddhartha S. Srinivasa. 2015. Effects of Robot Motion on Human-Robot Collaboration. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 51–58.

[7] Anca D. Dragan, Kenton C.T. Lee, and Siddhartha S. Srinivasa. 2013. Legibility and predictability of robot motion. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 301–308.

[8] Anca D. Dragan and Siddhartha S. Srinivasa. 2013. Generating Legible Motion. In *Robotics: Science and Systems (RSS)*.

[9] Miguel Faria, Francisco S. Melo, and Ana Paiva. 2024. "Guess what I'm doing": Extending legibility to sequential decision tasks. *Artificial Intelligence* 330 (2024), 104107.

[10] Jonatan Ginés, Francisco Martín, David Vargas, Francisco J. Rodríguez, and Vicente Matellán. 2019. Social Navigation in a Cognitive Architecture Using Dynamic Proxemic Zones. *Sensors* 19, 23 (2019).

[11] Jerome Guzzi, Alessandro Giusti, and Luca Maria Gambardella. 2013. Human-friendly robot navigation in dynamic environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 423–430.

[12] S. G. Hart and L. E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. *Human Mental Workload* 1, 3 (1988), 139–183.

[13] Mitsuhiro Kamezaki, Ayano Kobayashi, Yuta Yokoyama, Hayato Yanagawa, Moondeep Shrestha, and Shigeki Sugano. 2020. A Preliminary Study of Interactive Navigation Framework with Situation-Adaptive Multimodal Inducement: Pass-By Scenario. *International Journal of Social Robotics* 12, 2 (2020), 567–588.

[14] Lawrence H. Kim and Sean Follmer. 2021. Generating Legible and Glanceable Swarm Robot Motion through Trajectory, Collective Behavior, and Pre-attentive Processing Features. *ACM Transactions on Human-Robot Interaction* 10, 3, Article 21 (July 2021), 25 pages.

[15] Ross A. Knepper, Christoforos I. Mavrogiannis, Julia Proft, and Claire Liang. 2017. Implicit Communication in a Joint Action. In *Proceedings of the ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 283–292.

[16] Marina Kollmitz, Kaijen Hsiao, Johannes Gaa, and Wolfram Burgard. 2015. Time dependent planning on a layered social cost map for human-aware robot navigation. *Proceedings of the European Conference on Mobile Robots (ECMR)* (2015), 1–6.

[17] 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.

[18] Christina Lichtenthäler and Alexandra Kirsch. 2013. Towards Legible Robot Navigation - How to Increase the Intend Expressiveness of Robot Navigation Behavior. In *Workshop on Embodied Communication of Goals and Intentions, International Conference on Social Robotics (ICSR)*.

[19] Shuijing Liu, Peixin Chang, Zhe Huang, Neeloy Chakraborty, Kaiwen Hong, Weihang Liang, D. Livingston McPherson, Junyi Geng, and Katherine Driggs-Campbell. 2023. Intention Aware Robot Crowd Navigation with Attention-Based Interaction Graph. In *IEEE International Conference on Robotics and Automation (ICRA)*. 12015–12021.

[20] David V. Lu and William D. Smart. 2013. Towards more efficient navigation for robots and humans. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 1707–1713.

[21] Yue Luo, Nicolas Grimaldi, Haolan Zheng, Wayne C.W. Giang, and Boyi Hu. 2023. Distraction From Smartphones Changed Pedestrians' Walking Behaviors in Open Areas. *Motor Control* 27, 2 (2023), 275 – 292.

[22] 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 Transactions on Human-Robot Interaction (THRI)* 11, 2 (2022), 1–37.

[23] Christoforos Mavrogiannis, Francesca Baldini, Allan Wang, Dapeng Zhao, Pete Trautman, Aaron Steinfeld, and Jean Oh. 2023. Core Challenges of Social Robot Navigation: A Survey. *Transactions on Human-Robot Interaction* 12, 3 (2023).

[24] Stefanos Nikolaidis, Anca Dragan, and Siddhartha Srinivasa. 2016. Viewpoint-based legibility optimization. In *Proceedings of ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. 271–278.

[25] Elena Pacchierotti, Henrik 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.

[26] Elena Pacchierotti, Henrik I. Christensen, and Patric Jensfelt. 2006. Embodied Social Interaction for Service Robots in Hallway Environments. In *Field and Service Robotics*. Springer Berlin Heidelberg, 293–304.

[27] Sriyash Poddar, Christoforos Mavrogiannis, and Siddhartha S. Srinivasa. 2023. From Crowd Motion Prediction to Robot Navigation in Crowds. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 6765–6772.

[28] Christoph Schöller, Vincent Aravantinos, Florian Lay, and Alois Knoll. 2020. What the Constant Velocity Model Can Teach Us About Pedestrian Motion Prediction. *IEEE Robotics and Automation Letters* 5, 2 (2020), 1696–1703.

[29] Andrew Stratton, Kris Hauser, and Christoforos Mavrogiannis. 2025. Characterizing the Complexity of Social Robot Navigation Scenarios. *IEEE Robotics and Automation Letters* 10, 1 (2025), 184–191.

[30] Andrew Stratton, Phani Teja Singamaneni, Pranav Goyal, Rachid Alami, and Christoforos Mavrogiannis. 2026. How Human Motion Prediction Quality Shapes Social Robot Navigation Performance in Constrained Spaces. In *Proceedings of the IEEE/ACM International Conference on Human-Robot Interaction (HRI)*.

[31] Freek Stulp, Jonathan Grizou, Baptiste Busch, and Manuel Lopes. 2015. Facilitating intention prediction for humans by optimizing robot motions. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 1249–1255.

[32] Daniel Szafir, 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.

[33] Ada V. Taylor, Ellie Mamantov, and Henny Admoni. 2022. Observer-Aware Legibility for Social Navigation. In *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. 1115–1122.

[34] Phani Teja Singamaneni, Anthony Favier, and Rachid Alami. 2021. Human-Aware Navigation Planner for Diverse Human-Robot Interaction Contexts. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 5817–5824.

[35] Nick Walker, Christoforos Mavrogiannis, Siddhartha S. Srinivasa, and Maya Cakmak. 2021. Influencing Behavioral Attributions to Robot Motion During Task Execution. In *Proceedings of the Conference on Robot Learning (CoRL)*.

[36] Grady Williams, Nolan Wagener, Brian Goldfain, Paul Drews, James M. Rehg, Byron Boots, and Evangelos A. Theodorou. 2017. Information theoretic MPC for model-based reinforcement learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 1714–1721.

[37] Nicholas H. Wolfinger. 1995. Passing Moments: Some Social Dynamics of Pedestrian Interaction. *Journal of Contemporary Ethnography* 24, 3 (1995), 323–340.

Received 2025-12-08; accepted 2026-01-12