Considerations about Social Norms Compliance in a Shared Elevator Scenario

Danilo Gallo, Shreepriya Shreepriya, Tommaso Colombino, Maria Antonietta Grasso, and Cecile Boulard

Abstract—

In this paper, we present our ongoing research on socially acceptable robot navigation for an indoor elevator sharing scenario. We highlight the current challenge of designing interactions for a robot behavior, both effective in accomplishing tasks but not intrusive or at risk of breakdown. We discuss the advantages and limitations of modeling these behaviors based on a full human-like approach. In particular, we discuss the risk that a full human-like approach presents of creating the illusion of social competence. It has been observed that this illusion often leads to breakdowns when the technology is faced with complex and potentially ambiguous social situations. We propose the principle of "machine-like yet human-friendly" behavior to address the risks of the completely human mimicking approach. We believe that this approach can provide more understandable and less disruptive behaviors for routine integration into human spaces. We conclude by discussing the need for a multi-layer experiment set up to evaluate and validate this approach.

Keywords: Social norms; Robot navigation; Robot legibility.

I. INTRODUCTION

Questions around what constitutes socially acceptable behavior for autonomous agents are not new to the HRI community [1]. The safe and harmonious deployment of robots in public spaces requires their social behavior to be understandable and predictable [2]. The social norms that constitute the fabric of human interactions can be very informative to model the robot behaviors and blend them into a social setting. These norms are potentially useful for modeling machine behavior in an understandable way and adapting it to the context [3].

In this paper, we present how our current work in the area of social navigation has benefited from a fine-grained understanding of the social norms that are present during the activity of taking a shared elevator. In this work, we have focused on indoor scenarios using a robot platform developed by our organization. These robots are capable of navigating autonomously and carrying objects (Figure 1). These scenarios include delivering parcels and food orders in the context of a large office building. The robots we design have some dedicated infrastructure (such as dedicated elevators in a specific office building). However, they also have to be able to utilize shared infrastructure and spaces with humans at certain times. The robots are controlled by a centralized processing robot brain to keep each robot relatively simple and modular, reducing unit costs. In these office scenarios, we focus our investigation on aspects of social robot navigation involving the use of shared elevators with humans.

Previous research has focused on the technical development of robots able to autonomously operate and ride elevators with humans, such as [4] and [5]. However, we are not aware of any prior art that explicitly investigates the appropriate navigation behavior a robot should display considering social norms and human preferences for that specific context. Indeed, this context presents a range of challenges that go beyond the interaction with the infrastructure, such as negotiation of priorities or movement and coordination in reduced spaces.

Figure 1. The NAVER robotic platform.

When investigating the broader domain of social robot navigation, we have come across several approaches and specific contributions to enable a socially acceptable interaction. However, we encountered the lack of a common approach that characterizes the aimed-for human-friendliness that we could adopt in our work. This led to the definition of our own approach, which is design-driven and grounded in a fine-grained understanding of the social interactions at play in the context of interest. In doing so, we also drew inspiration from the body of work that has highlighted the challenges involved in making complex AI systems and decision making transparent to lay users, when encountering these systems in shared spaces and infrastructure [6][7], e.g., in the context of Autonomous Vehicles.

Authors are with Naver Labs Europe, France.
Name.surname@naverlabs.com

II. SOCIAL NAVIGATION WHILE TAKING A SHARED ELEVATOR

As mentioned in the introduction, our research focuses on indoor office scenarios. These scenarios involve robots sharing elevators with employees and visitors at our corporate headquarters. The robots deliver parcels and food orders in the context of a large office building, where they display autonomous behaviors and are controlled by a dedicated infrastructure. This dedicated infrastructure foresees the use of robot-only elevators and moments in which robots would need to be able to use shared elevators, e.g., when the workload is high. While we have seen commercial robots navigate in and out of elevators, there is little focus on the issues of proxemics and cultural preferences [8] in this context. Designing for such issues requires a distinction between what might be characterized and reduced to rules easily built into the robots' navigation behaviors and more nuanced matters of elevator social norms that would require the robots to identify things like human gaze and posture and *read* their interactional meaning and valence. While in our corporate headquarters, the centralized robot "brain" can operate the elevators for the robots, the scenario of robots using the elevators also points to the need for more complex interactions designed to enhance the robots' flexibility. Such interactions include robots requesting assistance operating elevators that might not be explicitly designed for their use and not integrated into the robotic platform's infrastructure. This need is similar to those investigated in the context of collaborative robotics [4] aimed at creating robots capable of compensating for their physical (ability to manipulate the environment) or perceptual limitations by eliciting human assistance [9].

We are aware of existing commercial robots using elevators, which rely mainly on speech interfaces, i.e., the robot announcing its intention to enter the elevator and declaring where they will position themselves [10]. While this can be an effective strategy, it clearly places the burden of making the interaction work on the people sharing the elevator with the robot. It may even be socially acceptable if the interactions are occasional (as might be the case, for example, with delivery robots in a hotel where any given guest might encounter the robot once during their stay). However, in an office environment with service robots performing routine tasks, encounters with robots in elevators are likely to be a daily occurrence for people working in the building, which means that negotiating the use of elevators through loud verbal announcements could quickly become tiresome.

A different and far more ambitious technology and interaction design paradigm might be to develop a platform capable of reading non-verbal behavior and the social context and its norms. This understanding could enable more subtle interactions, with robots that treat people as agents occupying a shared space with rules to follow for things like order of service and priority, rather than just obstacles to be avoided. This, of course, presents a substantial challenge as there is a semantic gap to be bridged between detecting things like posture or predicting movement (intentionality) and making dynamic and contextually appropriate decisions in what is (most of the time, *but* not always) for us a straightforward, but quite nuanced social interaction.

A. *Understanding the activity*

In our approach to the elevator scenario, we first studied what we might call the *practice* of taking the elevator. While it is in many ways a straightforward accomplishment (arguably more than driving a car in traffic), it is also constituted of practices that are methodical and accountable, with normative components and nuanced, often non-verbal use of space and resources. We analyzed approximately 16 hours of video data gathered by placing a video camera in the elevator lobby of one of our company's research labs. We adopted an ethnomethodological analytic orientation [11] to understand the specific practices of waiting for, entering, and exiting an elevator. We do not go into all the findings in detail here, but for the purpose of this discussion, we focus on how the order of service is managed (who enters the elevator first when multiple people are waiting). Our observations reveal that there is a general first-come-first-served principle that is applied, but it is a *weak* one. People do not form proper queues, especially where multiple elevators are linked to a single call function. In this common scenario, people often drift towards the lobby's center and only position themselves clearly in front of a door when the elevator lights indicate it will be the next available elevator. However, the elevator lights are not entirely reliable as the elevator status may change, and the next available elevator may, in fact, be at the other end of the lobby. In such cases, people moving back and forth across the lobby (*chasing* the next available elevator) may lose or have to renegotiate their priority in the order of service. Additionally, groups of people wait in the lobby with the intention of taking the elevator together. These groups also have *weak* form attributes regarding closeness and body orientation, making it ambiguous sometimes even to a human observer. These groups might stand in front of the elevator, even calling for it while waiting for another member (Figure 2). The difference between a fully formed group and one under composition is related to elements like distance and body orientation. However, also in this case, several ambiguous situations have been observed in our videos.

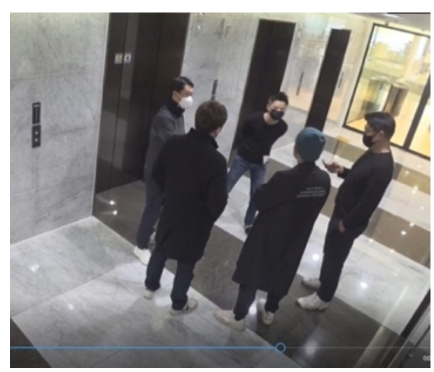

Figure 2. Waiting for another member to join in front of the elevator.

This seemingly obvious and easily accomplished behavior would present serious challenges if we wanted a socially competent robot to fully understand and mimic it. For example, what [12] describes as *the order of waiting* is constituted of both ordered and disordered formations, with demarcations and affiliations that are constantly produced and renegotiated. Even if we had computer vision technology that was capable of reliably and dynamically detecting things like posture, orientation, distance and displacement with respect to other people and elevator doors [13], gaze and facial expression, and computing the semantically and situationally appropriate reading of the situation, the nuances of the context would still be hard to capture, as they are for a human being at times. Additionally, we question how comfortable people would be with robots that fully mimic human-like behaviors and how these behaviors might contribute to an *illusion* of social competence.

One of the problems that concern us here and that we think is of interest to the HRI community is that, as [14] observed many years ago, the breakdowns in the interactions between technology and its users were often instantiated by what could be described as the *illusion* of social competence. Therefore, the design challenge we are exposed to is how an agent can effectively navigate shared spaces with people, focusing on safety and minimal disruption, but without necessarily being burdened with the normative expectations of being perceived as a fully human-like socially competent agent. This challenge resonates with discussions that have been done in a broader sense on the use (and extent of use) of anthropomorphic elements for robotic visual and behavioral elements' design [15].

With these considerations, we are experimenting with an intermediate option in which robots both exploit an understanding of the social situation and retain its representation as a tool to exhibit "machine-like yet human-friendly behaviors." Previous research [16] has defined machine-like behavior as the behavior that exploits machines' characteristics like sensors that humans do not have. For instance, autonomous vehicles (AVs) can know the position of other AVs without seeing them and act accordingly. Human-like behaviors are defined as the typical behaviors exhibited by humans based on their social understanding of other actors and the context. In our machine-like yet human-friendly behavior paradigm, we mix machine and human-like behaviors. We adopt human behavior elements that allow the robot to demonstrate the necessary level of social understanding that avoids disrupting the activity while ensuring task completion (e.g., position and direction with respect to the elevator door). We further incorporate machine-specific elements that convey the robot's role as subordinate entities with limited social understanding (e.g., unidirectional communication of intent and priority to humans as much as possible). We hypothesize that this approach would prevent the illusion problem introduced by taking a full human-like approach.

Our approach is reflected in a number of design choices aimed at creating robots with elegant, clear, and direct interaction mechanisms that encourage users to, for example, limit their needs of engaging with the service robots to what is part of their tasks and within their scope, and not beyond. This approach can limit potential breakdowns while ensuring that the robot does not disrupt the routine activity, i.e., taking the elevator. Indeed, this activity should be designed to become a non-experience, something people do without thinking and without consciously experiencing the interactions [17].

B. *The reality of social competence*

The notions of social competence come into play when technology designers use generic interaction metaphors like *human-like* or *pet-like* that give connotations that burden the agent individual and culturally dependent expectations. Moreover, as mentioned, the complicated technology computations required to model these socially nuanced and potentially ambiguous situations contextually are still a limiting factor to consider.

TABLE I. DEFINING OUR PROPOSED APPROACH WITHIN THE SPECTRUM OF MACHINE-LIKE AND HUMAN-LIKE BEHAVIORS IN THE SHARED ELEVATOR CONTEXT

	Machine-like	Machine-like yet Human-friendly	Human-like
Social awareness	No. The robot is not able to differentiate humans from other obstacles.	Some The robot can detect humans and is **aware of people entering and exiting** to decide its actions.	Yes The robot detects humans and their intentions. It adopts queuing as done by humans and moves of position according to an understanding of situations.
Communication of intent	Non-verbal, e.g., sounds. The robot only gives information for consumption.	The robot uses subtle **non-verbal interface elements** to broadcast intent, which are deliberate design choices for information consumption only.	Verbal and non-verbal like gaze, body posture, etc. The robot and the humans acknowledge and exchange information.
Movement and position	The robot positions itself in front of the door and always enters first.	The robot **takes a fixed waiting position and gives priority**, except in urgency.	Mimics humans with queuing and position adjustments.

To explain, we can take an example of waiting for an elevator scenario in which the robot is designed to have *human-like* behaviors. In this scenario, the robot should detect if a crowd of people is actually waiting for the elevator to act accordingly. This would require understanding if the group is waiting for the elevator or another group member to join them. In our video observations, we identify human pose and distance to the elevator as indicators of intention. However, we recognize that these elements are not enough to identify with confidence the group intention, even for a human evaluator. In order to disambiguate the situation, the robot would then need to initiate interactions with the group or take guesses to queue behind the unstructured group, potentially (and unnecessarily) delaying the service accomplishment. We further hypothesize that people would not feel comfortable with a robot that moves like humans, making continuous adjustments rather than taking a designated waiting position. Hence, we distance ourselves from following the strict approach of fully exploiting understanding and adherence to the elevator social norms in our work and establish the notion of *machine-like yet human-friendly* interaction behaviors. We define such behaviors as ones that respect human and social considerations particularly relevant to this activity and social context without explicitly mimicking human behaviors. As previously described, one of our findings exposed the fluidity of queuing and taking an elevator in a multi-elevator setting linked to common call buttons. In this case, a *human-like* behavior of changing queues and moving from elevator to elevator, as a human would do, would be relatable yet unpredictable, annoying, and potentially hazardous. Considering this, we designed specific actions to be taken by the robot based on only some elements of the elevator human etiquette while also putting in place specific and clear robot behaviors (TABLE I.).

C. Designing the robot behavior

The design of the activity is broken down into the following stages: waiting, entering, riding, and exiting the elevator. At the same time, each stage is divided into several smaller actions performed by the robot, during which the robot will display a series of communication states. Each state aims to convey a specific intent by using different communication modalities (i.e., sound, light, displays, projection, and anticipatory movements) and combinations of them. For this paper, we only describe the waiting and entry stages of the activity, along with the specific actions involved.

Waiting for the elevator:

a) Calls elevator

- The robot navigates to the elevator hall (it might encounter people waiting for the elevator in an unstructured queue or be the first to arrive).
- The robot calls the elevator and receives the information about the elevator that will arrive next. Elevator indicators should reflect this information.

b) Navigate to Waiting Position

- The robot navigates to the fixed waiting position (Figure 3). It commits to the corresponding elevator, even if the situation changes and another elevator arrives first (unlike a human-like behavior).
- If it detects humans within 46 cm [18], it tries to go around them or starts a Request state to ask permission to pass. We should point out that this distance might need to change in different cultures or contexts (e.g., within the elevator, where there is limited space). If the robot detects obstacles within 10 cm is tries to go around them.

c) Robot waits

- Once it reaches the Waiting Position, the robot displays it is in Waiting state.
- Regardless of people's movements around it, the robot remains in that place to avoid disrupting them with small position adjustments (unlike a human-like behavior).

Entering the elevator:

d) Elevator arrives

- The robot detects people exiting the elevator and waits for them to exit.

e) The robot lets people enter first

- After everyone left the elevator, the robot detects people are entering.
- If people are entering, it remains in place and triggers a Yielding state and lets people waiting to enter, regardless of the order of arrival (unlike a human-like behavior).
- Once everyone has entered the elevator, the robot triggers an In-Motion state.

f) Robot enters first

- In certain cases, if the robot's task (such as delivering hot coffee or food) is to be completed within a certain timeframe, the robot triggers an Urgent state right after people stop exiting the elevator. The Urgent state communicates the intention to enter first (even if it detects people are entering), and the robot enters the elevator. This behavior should be triggered only when a new (next) elevator arrives to avoid confusing people in the process of entering the elevator.

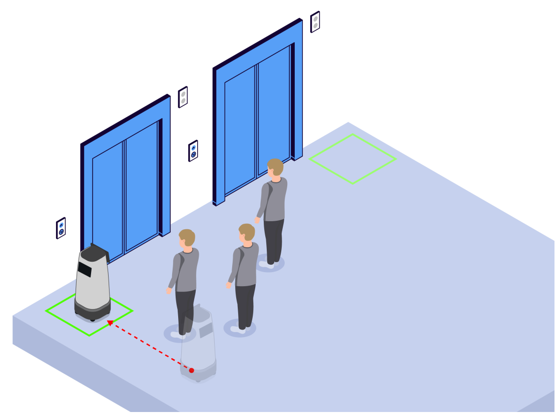

Figure 3. The robot takes a designated waiting position regardless of the position of the humans waiting for the elevator.

D. Considerations regarding experiment design

The social navigation behavior that we have designed based on the above considerations has resulted in the following hypothesis that we will assess with several user experiments, each tailored to the specific level of design under consideration.

- In certain cases (especially related to positioning and movements), people will understand and prefer *machine-like yet human-friendly* robot **actions** rather than those mimicking humans, such as the subtle movements exhibited by people taking elevators.

- Our proposed robot's behaviors for **communicating intent** (calling the elevator, waiting, entering) will be as understandable and less intrusive than fully human-like ones, for the first use as well as for an extended period.

In our elevator-taking scenario, the actions of *waiting* and *entering* have to be tested under the conditions of *machine-like yet human-friendly* and *human-like* to provide proof for our first hypothesis.

While to test the second hypothesis, we need to test the understandability of different alternate communication modalities used by the robot against the baseline of verbal interaction modality commonly used in existing service robots.

E. The layered testing approach

Our design proposal combines a range of navigation policies and different communication modalities to convey the robot's intent. To effectively validate the impact of these different elements, we need first to evaluate them independently to conduct then tests that integrate them into a comprehensive service. For this reason, we decided to adopt a layered testing approach consisting of three steps: online experiments, in-situ experiments, and naturalistic observations of the robot working in context. We also add naturalistic observations to our testing toolkit for a better ecological validity of our robotic service.

Online experiment. While online experiments may have limitations related to participants' profile, level of engagement, and quality of results, they can provide preliminary feedback without the constraints of participants' time and exhaustion. We will compare several alternatives for each defined action through a perceptual, low fidelity experiment in which videos of the situations identified around elevators with each condition (alternatives for specific features, e.g., position, interfaces, etc.) will be shown to participants. We will collect and analyze objective and subjective measurements of participants' understanding and preference for each condition.

In-situ experiment. In-situ evaluations are seen to be more valid, especially for performance-related metrics like response time, completion time, or tasks completed. We will evaluate the pre-selected alternate behaviors/features (from the results of the online experiment) in a realistic set-up. Participants will be requested to perform the actions related to using elevators while sharing the waiting space with a robot. We will collect objective measurements of response time and task completion and subjective measurements of understanding and preference.

Naturalistic observation. Evaluations conducted in lab-based controlled settings often lack ecological validity. To counter this, experimenters rely on experimental realism or simulating the context of use. While this is effective, it can still produce bias as the participants are recruited and briefed about the experiment [19], and the introduced novelty can also influence them. Hence, naturalistic observations can counter these effects. We will deploy the robot in the wild with the selected features from the in-situ test. We will capture and analyze the reaction of passers-by and people taking a shared elevator for first encounters and an extended period. Through video analysis, the person or group's understanding and preference will be analyzed.

III. Conclusion

The real-world deployment of autonomous robots presents several complexities. In an indoor environment, robots sharing an elevator with people can be considered one of the scenarios that will soon be a reality. To design for near-future deployable robots, the illusion of social competence of the robots must be carefully managed, as well the robot harmoniously blending with the social norms of a setting. Indeed, the technology does not provide fully reliable solutions to understand and model the complex and potentially ambiguous social situations we observed. Moreover, it is still an open question if a fully human-like behavior is desirable in autonomous agents. Hence, we propose a *machine-like* yet *human-friendly* approach to the design of robot navigation behaviors and a layered testing approach. Through these experiments, we aim to validate our assumption that *machine-like* yet *human-friendly* interactions are preferable for the robot's social behaviour.

References

[1] R. Simmons, R., J. Forlizzi and R. Kirby. 2010. "Social robot navigation."

[2] C. Lichtenthäler, T. Lorenzy and A. Kirsch. 2012. "Influence of legibility on perceived safety in a virtual human-robot path crossing task." IEEE RO-MAN: The 21st IEEE International Symposium on Robot and Human Interactive Communication, 2012, pp. 676-681, doi: 10.1109/ROMAN.2012.6343829.

[3] A. Sandygulova, M. Dragone, and G. M. P. O'Hare. 2016. "PRIveT- A portable ubiquitous robotics testbed for adaptive human-robot interaction". Journal of Ambient Intelligence and Smart Environments, 8(1), 5-19. https://doi.org/10.3233/AIS-150356

[4] M. Veloso, B. Joydeep, B. Coltin, S. Rosenthal, T. Kollar, C. Mericli, M. Samadi, S. Brandao, and R. Ventura. 2012. "CoBots: Collaborative Robots Servicing Multi-Floor Buildings". In Proceedings of IROS'12, the IEEE/RSJ International Conference on Intelligent Robots and Systems, Vilamoura, Portugal.

[5] J.-G. Kang, S.-Y. An, and S.Y. Oh. 2007. "Navigation strategy for the service robot in the elevator environment". ICCAS 2007 - International Conference on Control, Automation and Systems. 1092 - 1097. 10.1109/ICCAS.2007.4407062.

[6] B. Brown and E. Laurier. 2017. "The trouble with autopilots: Assisted and autonomous driving on the social road". In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, Denver, United States.

[7] E. Vinkhuyzen and M. Cefkin. 2016. "Developing Socially Acceptable Autonomous Vehicles. 2016 Ethnographic Praxis in Industry Conference Proceedings." https://www.epicpeople.org

[8] M. Joosse, M. Lohse, and V. Evers. 2014. "Lost in Proxemics: Spatial Behavior for Cross-Cultural HRI". 2014. HRI'14 Workshop on Culture Aware Robotics, Bielefeld, Germany.

[9] J. Liebner, A. Scheidig, and H.-M. Gross. 2019. "Now I Need Help! Passing Doors and Using Elevators as an Assistance Requiring Robot," in Social Robotics, vol. 11876, M. A. Salichs, S. S. Ge, E. I. Barakova, J.-J. Cabibihan, A. R. Wagner, Á. Castro-González, and H. He, Eds. Cham: Springer International Publishing, 2019, pp. 527–537. doi: 10.1007/978-3-030-35888-4_49.

[10] R. Ichinose, I. Takeuchi, T. Teramoto. 2015. "Elevator System that Autonomous Mobile Robot Takes Together with Person." US Patent US 8,958,910 B2. Feb. 17, 2015.

[11] H. Garfinkel. 1967. Studies in Ethnomethodology. Polity Press.

[12] R. Ayaß. 2020. "Doing Waiting: An Ethnomethodological Analysis." Journal of Contemporary Ethnography, Vol 49(4) 419-455.

[13] A. Vega-Magro, L. Manso, P. Bustos, P. Nunez and D. G. Macharet. 2017. "Socially Acceptable Robot Navigation over Groups of People." 26th

IEEE Internation Symposium on Robot and Human Interactive Communication. Lisbon, Portugal.

[14] L. Suchman. 2007. “Human-Machine Reconfigurations: Plans and Situated Actions.” Cambridge: Cambridge University Press.

[15] M. Boden. 2006. Robots and anthropomorphism. AAAI Workshop - Technical Report. 69-74.

[16] L. Oliveira, K. Proctor, C. G. Burns, and S. Birrell. 2019. "Driving Style: How Should an Automated Vehicle Behave?" Information 10, no. 6: 219. https://doi.org/10.3390/info10060219

[17] R. Rousi. 2013. “The experience of no experience Elevator UX and the role of unconscious experience”. In Proceedings of International Conference on Making Sense of Converging Media - AcademicMindTrek ’13, ACM Press, Tampere, Finland, 289–292. DOI:https://doi.org/10.1145/2523429.2523455

[18] E. T. Hall. 1966. “The hidden dimension: man’s use of space in public and private.” The Bodley Head Ltd, London.

[19] T. Kruse, A. K. Pandey, R. Alami, A. Kirsch. 2013. “Human-aware robot navigation: A survey“. Robotics and Autonomous Systems, Volume 61, Issue 12, 2013, Pages 1726-1743, ISSN 0921-8890, https://doi.org/10.1016/j.robot.2013.05.007.