

Familiar Acoustic Cues for Legible Service Robots

Georgios Angelopoulos¹, Francesco Vigni¹, Alessandra Rossi²,
Giuseppina Russo², Mario Turco², and Silvia Rossi²

Abstract—When navigating in a shared environment, the extent to which robots are able to effectively use signals for coordinating with human behaviors can ameliorate dissatisfaction and increase acceptance. In this paper, we present an online video study to investigate whether familiar acoustic signals can improve the legibility of a robot’s navigation behavior. We collected the responses of 120 participants to evaluate their perceptions of a robot that communicates with one of the three used non-verbal navigational cues (an acoustic signal, an acoustic in pair with a visual signal, and a dissimilar frequency acoustic signal). Our results showed a significant legibility improvement when the robot used both light and acoustic signals to communicate its intentions compared to using only the same acoustic sound. Additionally, our findings highlighted that people also perceived differently the robot’s intentions when they were expressed through two frequencies of the mere sound. The results of this work suggest a paradigm that can help the development of mobile service robots in public spaces.

I. INTRODUCTION

Service robots aim to assist people in various scenarios and environments by performing tasks to support day-to-day work, domestic, and care activities [1]. The autonomous mobile robot market was valued at 1.61 billion USD in 2021, and experts predicted that it will reach 22.15 billion USD by 2030 with a compound annual growth rate of 34.3% from 2022-2030 [2]. Therefore, people will need to gain familiarity with robots and their behaviors in private and public spaces. In this direction, researchers are increasingly working in deploying service robots both in controlled environments and in public spaces, such as shopping malls [3], airports [4], and hospitals [5].

In these domains, robots need to move in environments shared with people, furniture, and pets, where collisions may happen, especially when people and robots are not able to correctly infer the other’s intentions ending with conflicting trajectories. Nevertheless, collisions between human pedestrians can also occur every day. While people can resolve these conflicts by passively communicating through non-verbal cues [6], [7], a robot’s navigational intention can be more difficult for people to interpret. If robots are not

interactive and have legible behavior, they will be simply perceived by people as obstacles rather than helpful assistants [8]. Robot manufacturers are creating robots that can use non-verbal communication systems to tackle the problem of legibility but in non-intuitive fashions. Little effort has been made to investigate whether the familiarity of a cue, and its association with motion, can influence the users’ understanding.

In most situations, the everyday environment consists of individuals walking through hallways and corridors immersed in their busy lives. People may not have enough time to focus and concentrate in this situation; thus, a rich design robot with plenty of information will be ineffective for interacting with people. Any communication method utilized to express directional purpose should be straightforward and intuitive to use. For this reason, robots should be designed to express navigational intentions using familiar cues to elicit natural and easy communication with humans.

The European Union, with a new Commission Delegated Regulation 2017/1576¹, for example, understood the problem of legibility and mandated the use of a familiar acoustic signal emitted by the e-car to make people aware of the presence and intention of a car that otherwise would be almost silent. Our common definition of familiar road knowledge needs to be expressed with formal tools to define interfaces in mobile robotics. Bearing that in mind, the present work designs, tests, and compares the explainability of human-familiar signaling mechanisms that do not require a demonstration to be understood. Particular attention has been given to the use of acoustic signals, as stand alone or in combination with visual cues. Therefore, we tested whether and how variations of acoustic signals can easily express the robot’s intentions without ambiguity during the robot’s navigation in a shared environment.

II. RELATED WORK

Due to the increasing deployment of mobile robots in our everyday lives, several researchers are investigating how to create legible robot behaviors to communicate their navigational intentions. As described by [9], non-verbal communication is fundamental to sharing the robot’s internal state and making it legible to people. Moreover, authors in [10] highlighted that non-verbal signals, when used by service robots in social contexts, should be designed to match the expectations of the surrounding users. To develop simple and

*This work has been supported by the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 955778.

¹ are co-first authors and with the Interdepartmental Center for Advances in Robotic Surgery - ICAROS, University of Naples Federico II, Naples, Italy [georgios.angelopoulos, francesco.vigni]@unina.it

² are with the Department of Electrical Engineering and Information Technologies - DIETI, University of Naples Federico II, Naples, Italy [alessandra.rossi, silvia.rossi]@unina.it

¹<https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32017R1576&from=EN>

intuitive signals, many implementations revolve around using lights and sound cues in different ways to reveal the robot's internal state. Indeed, light is a mechanism by which the robot can communicate helpful state information to people by varying their intensity, color, or frequency. On the contrary, sound as a communication tool works more effectively across cultures and language groups [11]. Jee et al. [12] in their study showed that sound as a cue for robots was effective at conveying the intention but also at expressing their emotions.

Fernandez et al. [13] presented a study in which a mobile robot navigates in a hallway and signals its intention of passing a human participant using a strip of LEDs as a turning signal. The authors observed that people did not interpret the LED turning signals when interacting with a non-humanoid robot, but a brief and passive signal demonstration is sufficient to disambiguate its meaning.

Shrestha et al. [14] compared arrow-like shapes on a display screen to flashing lights as motion legibility cues in a head-on interaction between a walking person and a mobile robot. The turn indicator was rated significantly better than the screen indication.

Hart et al. [15] used a virtual agent head on a mobile robot to measure the relevance of the robot gaze in coordinating the navigation with people. Their results highlighted that people found this modality easy to interpret, and, therefore, the robot's intended navigational direction was legible.

Fiore et al. [16] evaluated the effects of the social cues, such as gaze and proxemic, on people's perceived social presence and emotional state attributions to a non-humanoid robot. They observed that the gaze was not an effective communicative modality for a non-humanoid robot.

Watanabe et al. [17] presented a navigational intention approach that allows a wheelchair to communicate its intentions to its passenger and the people in the shared environment using lights projected on the floor. Such modality provided comfortable navigation both for the passenger and the people encountered along the way.

Most approaches presented in this section need demonstration to be understood by humans or depend on the level of anthropomorphism of the robot. Moreover, the majority of the state-of-art studies considered sound like a single means to attract people's attention "in the wild" settings. We believe that mobile robots need to use clearer, more familiar, and efficient navigational cues for avoiding collisions while maneuvering in shared spaces with people. Therefore, in our work, we investigated whether the familiarity of a sound generated by a mobile robot can also clearly infer the robot's intention to the users.

III. METHODS

This work aims to exploit ordinary users' experience with vehicles to improve the legibility of a robot's motion. We believe that a critical aspect is that users need to be familiar with the cues used by a robot to communicate its intention without explicitly training people to read them. For example, commercial passenger vehicles are equipped with turning sounds and lights to signal the intention to change the

vehicle's travel direction. Many vehicles can also rely on an intermittent sound that changes its frequency to signal the approaching of an obstacle. Such simple communication modalities, which have been around for many years, are unambiguously recognized worldwide.

Therefore, we decided to investigate whether it is possible to seamlessly transfer the semantic knowledge from vehicles to mobile robots. To this extent, we selected a subset of non-verbal communication modalities used by common vehicles.

Two non-verbal signals (i.e., audio and light) have been identified and combined to express navigational cues:

Navigational Cue 1 (NC1): The robot's produced through one of the speakers an intermittent tone with a constant frequency similar to the turning indicator sound of a vehicle. We believe that the use of a turning indicator sound could elicit a directional intention in people.

Navigational Cue 2 (NC2): The robot used a red LED and a speaker to produce a paired intermittent switch with constant frequency. This cue resulted in a synchronous use of light and sound (i.e., when the light is on, the speaker produces a tone). NC2 has been designed to provide turning direction (via the right led) while producing the same sound profile presented in NC1. In this cue, the sound is not intended for conveying directional intent, but to attract attention. The red color usually attracts the focus of people's attention [18].

Navigational Cue 3 (NC3): The robot produced an intermittent tone with a frequency inversely proportional to the read of the proximity sensor. This cue aims to mimic the tone produced by a vehicle (e.g., connected to parking sensors) that is approaching an obstacle. The tone employed frequency variations to convey the remaining distance to the obstacle.


All navigational cues have been designed based on simplicity and are known by people in similar contexts.

We believed that sound plays an essential role in the communication of intention, however, it conveys a clearer directional motion when combined with the use of the lights. Therefore, our hypotheses were: **Hypothesis H1** the navigational cue NC2 is more legible than NC1; and **Hypothesis H2** varying the frequency of the sound in NC3 will improve the clarity of the robot's intention compared to NC1.

A. Robotic Platform

Robots that are already highly integrated into society are more likely to be easily understood and consequently accepted [19]. For this reason, we used a simple robot, resembling a vacuum cleaner, for domestic usage as our platform. To test our navigational cues, we selected an *arduino* microcontroller [20] and designed a prototype printed circuit board (PCB) to control small magnetic speakers, a buzzer, a proximity sensor, and two independent LEDs positioned on top of the robot, as shown in Figure 1a.

The speaker and the buzzer differ in tone and modulation. The buzzer was used to elicit an interaction that could recall



(a) Schematics of the electronics.



(b) Front view of the assembled prototype.



(c) Rear view of the assembled prototype.

Fig. 1: The platform.

a vehicle during a parking maneuver. The buzzer emitted an intermittent tone whose frequency increased as the robot approached the user. The sound produced by the speaker, instead, was designed to mimic the noise of an active turning light as perceived inside a vehicle. Each LED was installed in an opaque glass-shaped container so that its light could use a higher surface. Each container was located on the top-side of the robot, mimicking the typical lateral position of turning lights in a vehicle. A version of the final design can be found in Figure 1. The figure is obtained by augmenting an open-source 3D robot model with shapes representing the LEDs containers and the electronic box.

IV. USER STUDY

We conducted a user study to test the legibility of the proposed navigational cues. We designed a between-subject study where different participants tested each condition so that each person was only exposed to a single navigational cue.

A. Procedure

The COVID-19 pandemic interrupted our plans for in-person trials; therefore, we conducted an online user study using video clips of the robot’s navigational cues and a digital survey. The video clips had the same shooting angle, duration, environment, and lighting conditions. In particular, the videos were recorded in a long corridor (145cm wide, 320cm long) with the camera positioned at the opposite end, facing the robot. The camera was at a fixed position in all conditions to avoid variations in the field of view. The video clips lasted 10 seconds, and during the first 5 seconds, the robot moved in a straight line starting at the beginning of the corridor towards the camera. Then, the robot moved forward while signaling using navigational cues. The robot used only the cue but did not complete the next movement (i.e., turning or moving forward) to hide the navigational goal of the robot since we intended to evaluate participants’ understanding of the robot’s intent. The online study was distributed to participants via social media and to the University’s community members.

B. Measurements

At the beginning of the experiment, participants were asked to provide demographic information about their age, gender, level of education, and to rate their previous experience with robots. To evaluate people’s perceptions and to understand the legibility of the designated cues, a brief post-interaction survey comprising 5-point Likert and nine scale questions was provided to the participants. The questions presented in the user study are similar to [21] that were used for evaluating navigational cues used by a mobile robot. We removed the question “I trust the robot” from the set of questions, as we believe it was out of scope for this work.

The questions can be clustered as follows: (in brackets are the labels as used in Figure 2):

1) Comprehensibility

- The robot’s behavior was misleading (*Misleading*)
- I quickly understood the robot’s behavior (*Understandable*)
- It is difficult to understand what the robot intended to do (*Unclear*)

2) Reliability

- The robot was deceptive (*Deceptive*)
- I am weary of the robot (*Draining*)

3) Social compatibility, comfort, and friendliness

- The robot’s behavior would be socially compatible with a pedestrian’s environment (*Socially compatible*)
- The robot’s behavior made me feel comfortable (*Pleasant*)
- I liked the robot (*Likable*)

Furthermore, a direct question was added to the survey “What do you think the robot will do?” to further analyze whether participants correctly perceived the intention of the robot behind the navigational cue. The participant could choose an answer between the following: “The robot is turning Left”, “The robot is turning right”, “The robot continues straight”, “I do not know”. Finally, participants were asked to rate the confidence level of their previous




Fig. 2: Average of responses to the questionnaire, significant differences between navigational cues have been indicated with * for $p < 0.05$ and ** for $p \leq 0.001$.

answer using a 5-point Likert scale.

V. RESULTS

One hundred twenty participants took part in the study. The pool of participants was composed of 73 males, 46 females, and one preferred not to say (age: $M=25.03$, $SD=7.46$), which would allow us to detect an effect size of $d=0.25$ with .90 power at an alpha level of .05. Participants were distributed for the three navigational cues conditions as follows: forty-three participants in NC3, 39 and 38 participants in NC1 and NC2, respectively.

Over two-thirds of the participants (69.2%) stated to be enrolled at the University either as master's or bachelor's students, 23.1% were employees, and 7.7% were unemployed. Furthermore, 59.2% of the participants stated that they had never interacted with robots before.

A. Questionnaire Ratings

An Independent Samples T-test with 95% confidence intervals was used to determine if a difference exists between the navigational cues. Figure 2 shows the average responses grouped by navigational cues.

In the question on *Misleading*, NC1 scored significantly higher than NC2 with a mean of 2.770 (0.442 to 1.518), $t(73.958) = 3.636$, $p = .001$. There was also a statistically significant difference between NC2 and NC3, with NC3 scoring higher than NC2, 2.470 (-1.192 to -0.159), $t(79.000) = -2.604$, $p = .011$. Although the averages scored below the mean of 3.0 (on a 5-point Likert scale), users find NC2 (i.e., the robot using light and sound to convey directional intent) to be the less misleading navigational cue.

In the question on *Understandable*, NC1 scored significantly higher than NC2 with a mean of 3.740 (-2.247 to -1.176), $t(74.985) = -6.368$, $p < .0005$. A statistically significant difference between NC1 and NC3 was observed, with NC3 scoring higher than NC1, 3.230 (-1.734 to -0.680), $t(79.079) = -4.557$, $p < .0005$. Hence, varying the sound profile and intermittent frequency (NC3) can significantly contribute to conveying directional intent.


A significant difference is also found on the *Unclear* scale between NC1 and NC2, with NC1 scoring higher than NC2, 3.080 (0.379 to 1.617), $t(74.723) = 3.214$, $p = .002$. These are in line with the results in the *Understandable* scale, and people found clearer a robot that employs light and sound (NC2) to communicate directional intent rather than one using only sound (NC1).

Considering the responses to the question *Deceptive*, NC1 scored significantly higher than NC2 with a mean of 2.260 (1.142 to 1.002), $t(72.476) = 2.651$, $p = .010$, and a statistically significant difference between NC2 and NC3, with NC3 scoring higher than NC2, 2.350 (-1.137 to -0.192), $t(73.504) = -2.804$, $p = .006$. These results show that participants found NC2 to be the less deceptive navigational cue. Moreover, no significant difference was observed between NC1 and NC3 on this scale.

Finally, on both questions *Socially compatible* and *Pleasant*, NC2 scored significantly higher than NC1 respectively with a mean of 3.500 (-1.457 to -0.414; $t(74.867) = -3.575$, $p = .001$), and 3.320 (-1.415 to -0.345; $t(74.473) = -3.279$, $p = .002$). These outcomes show that the users considered a robot that signals its motions using an acoustic tone synced with a blinking LED (NC1) more congruous for a pedestrian environment rather than one using only the acoustic tone (NC2).

Considering participants' responses to the *Draining* and *Likable* questions, we did not observe any significant difference between the three navigational cues.

The questionnaire can also be grouped into *positive* (i.e., Understandable, Socially compatible, Pleasant, and Likable) and *negative* (i.e., Misleading, Unclear, Deceptive, and Draining) scales. If the values are higher in these questions, the first group associates them with a better robot's behavior, while the second associates them with the worst behavior. The strength of the NC2 cue is further reflected both in the positive effects and the negative effects of the post-interaction survey (see Figure 3a and Figure 3b respectively). Figures 3a and 3b provide useful insights into the relative quality of the navigational cues. In particular, we can observe



(a) Participants' outcomes demonstrated a trend of higher mean ratings in Social compatible, Pleasant, and Understandable for the NC2 and Likeable for the NC3. (b) Participants' outcomes demonstrated a trend of higher mean ratings in Misleading, and Unclear for the NC1, Draining and Deceptive for the NC3.

Fig. 3: Positive and Negative effects of the Navigational Cues.

a trend of higher mean ratings for the positive questions in the NC2, showing that is the preferred one since it is evaluated more favorably. It is also interesting to notice that the inverted coaxial order of navigational cues between the negative and positive questions marked NC2 as the preferred navigational cue, followed by NC3 and NC1.

B. Robot's Legibility

To further test the cues' legibility, we analyzed the question: "What do you think the robot will do?". As can be observed from Table I, the clearer behavior was NC2 with 81.6% of correct answers, and followed by NC3 and NC1 respectively with 67.4% and 12.8%. A Chi-square test showed a significant difference between the correct answers of NC2 and NC1, with $\chi^2(1) = 16.537; p < .001$, and at the same time a significant difference between the correct answers of NC3 and NC1, with $\chi^2(1) = 8.159; p < .001$. We also observed a statistically significant difference between the correct answers of NC3 and NC2.

Figure 4 illustrates three heat maps of the responses (one for each navigational cue) where the darkest shades correspond to a higher frequency of the sureness. For NC2 and N3, users perceived the intention of the turning signal and also the direction of it, even if they did it more clearly in NC2. The heat map obtained for NC1 shows that participants did not identify a clear perceived motion, even though most users (48.7% of the participants) identified such intended motion as "The robot continues straight" instead of turning.

TABLE I: Frequency of the answers to the question "What do you think the robot will do?".

	<i>Answer</i>	<i>Frequency</i>	<i>Percentage</i>
NC1	The robot is turning (Correct Answer)	14	35.8
	Other false answers	25	64.2
NC2	The robot is turning (Correct Answer)	31	81.6
	Other false answers	7	18.4
NC3	The robot continues straight (Correct Answer)	29	67.4
	Other false answers	14	32.6

C. Discussion of the Results

Our results exceed the hypothetical expectations, especially considering the first Hypothesis (H1). Results indicate that a robot that uses light and an acoustic signal is preferred and more effective than a robot that uses just an acoustic signal. Moreover, we obtained higher percentage values on false answers on NC1, but not on NC3 cue (Hypothesis (H2)). As expected for the acoustic cues, better performances are found in the navigational cue NC3 with frequency variation that depends on proximity measures, due to the motion's familiarity of the users. These findings revealed that further research is needed on investigating different frequency variations of the sound for the acoustic signals, since humans are getting more and more familiar with these signals in their everyday lives.

In contrast with the results of Fernandez et al. [13] in which the visual signal (LEDs) needed a demonstration to be understood, our findings showed that the NC2 cue proposed in this work ameliorates this problem as participants find it more comprehensible. Furthermore, even though previous studies stated that sound had no positive effect on navigational cues, our results showed that humans accept sound cues that are familiar to them, such as the sound of an interface for a potential collision. In our study, however, people were not able to distinguish a directional intent from sounds unless it was linked to the use of LEDs.

Even though Woods et al. [22] showed that video results were equivalent to in-person results in a study about which direction was the most appropriate for a robot to approach a human, future works will include in-person interactions between people and the robot. These would help us to further investigate whether people are able to similarly discern sounds even if they do not use headphones or are closer to the speakers.

VI. CONCLUSIONS

As robots move into human-populated spaces, it is more vital than ever to navigate in shared spaces with humans

