Supporting Observability through Social Cues

Natalie Friedman
Cornell Tech
New York City, NY, USA
e-mail: nvf4@cornell.edu

Patricia L. McDermott, Jeff Stanley
The MITRE Corporation
McLean, VA, USA
{pmcdermott, jstanley}@mitre.org

***Abstract*—For improved acceptance of robots in social spaces, it is important to have a strong mental model of what the robot can do, what the robot is currently doing, or what the robot is about to do. How do social cues help people understand what is going on 'under the hood'? Imagine this: a machine perks up if someone enters the room. This lets you know it is socially aware, awake, and ready to interact. Drawing from a pre-existing taxonomy of social cues for conversational agents, we reviewed 40 papers with instances of robot or software agent personality traits influencing observability. This survey led us to elaborate on six particular cues, clarify their relationship to observability and provide examples, with the intent to advance discussion and encourage research on the relationship between social cues and observability.**

Keywords-observability; social cues; personality; human-robot interaction; human-computer interaction; trust; predictability

I. Introduction

How do social cues help people understand what is going on “under the hood” of robots? Many fields including Human-Computer Interaction (HCI), dance, and animation confirm that expressivity can reveal functionality [1][2]. However, there is limited research presenting design guidelines that link individual social cues with a robot’s internal state, i.e., what is happening “under the hood.” Observability is defined as appropriate transparency into what an automated partner can do and is doing relative to task progress [3]. Amy LaViers, a Human-Robot Interaction (HRI) researcher who studies movement, describes that in the Laban movement community there is an "indivorcibility of function and expression" and that "more expressive robots are more functional robots" [1]. HRI researchers, studying movement and expressivity, divide functional task motions (e.g., “grabbing the doorknob”) from expressive motions (e.g., “looking around the door handle and scratching its head”). While they attempt to separate the two, they write “we do not subscribe to the idea that these are completely separate concepts” [2]. Similarly, we argue that there is an indivorcibility of observability and expressivity. In this paper, we contribute an assessment of how an agent's social cues [4] can help improve robot observability. Discussion of social cues for robots is nothing new, and we aim to advance the discussion by clarifying ways to view social cues through the lens of observability.

A. *Observability*

Observability and transparency have been used synonymously [3]. A transparent system communicates feedback about the system reliability and situational factors; it can establish appropriate trust and improve team performance [5]. Through transparency, teammates generate a shared understanding of the task and calibrate trust based on the team members' capabilities [6]. It is especially important for humans to develop calibrated trust because inappropriate trust often leads to misuse or disuse of automation [5]. Transparency also gives the human team members situational awareness of the task, robot, and environment [7][8].

The benefits of observability come with challenges. Communication among human-machine teams (or teams of humans and automated agents) is restricted due to the limited ways that automated agents give and receive communication. For example, humans use non-verbal communication, and can adapt communication to situations outside the nominal task [8]. Some of the non-verbal ways that humans communicate are via social cues [4], which have been implemented on robots, as we will discuss.

B. *Personality and Social Cues*

Human personality is defined as “characteristic sets of behaviors, cognitions, and emotional patterns that evolve from biological and environmental factors” [9]. While robots are not influenced by biological factors, they have software and environmental sensors. In designing robots, personality can be operationalized by social behaviors and cues that are designed to be appropriate for the robot’s role and environmental factors.

A social cue is defined as “a cue that triggers a social reaction towards the emitter of the cue” [4][10][11][12]. In human-robot interaction, a social cue has been defined as “features that `act as channels of social information’” [4][13].

Which social cues should robots use? In human-computer interaction, performative behaviors [14] have been used to display a machine's reaction to its actions [2], system state, or to demonstrate its social standing [15]. For example, researchers in collaboration with animators designed a slouching behavior for a robot when the robot could not open the door [2]. This performative reaction demonstrates to the bystanders that it is aware and embarrassed of its failure. Slouching after failure is an example of how social cues can improve observability, and ultimately the human understanding of a robot.

 ISBN: 978-1-61208-896-9

Performative behaviors are also used to change the social dynamic within a group. For example, Jung et al. used repair interventions ("hey, not cool") to achieve an awareness of conflict within a human group dynamic to prevent personal attacks [15]. Similarly, Sebo et al. found that designing a robot to be vulnerable ("Sorry guys, I made the mistake this round. I know it may be hard to believe, but robots make mistakes too") can make others more likely to share their failure to the group and laugh together [16]. In this case, this performative behavior gives the user information about its knowledge of failure and displays personality. Here, we see performative social cues can help people understand that the robot is aware of its mistake by showing shame and vulnerability and aware of the social dynamic of the situation by reacting to it.

In this paper, we assess specific social cues and how they can make a robot's internal state more observable. The paper has five sections. Section II briefly describes the methodology and sources that informed our analysis. Section III introduces six social cues and illustrates how they can contribute to observability, as summarized in Table I at the end of that section. In Section IV we highlight some overall themes and open questions. Section V concludes the paper.

II. METHODS

We intend to develop a dataset of social cues that are useful to designers for improving observability of robots and automated agents. Although we primarily refer to robots in this paper, the examples and findings are also applicable to embodied agents or digital assistants. We draw from the paper "A Taxonomy of Social Cues for Conversational Agents" by Feine et al. [4], which systematically identifies 48 social cues from the literature. In parallel, we have searched for these social cues in relationship to observability. We used a combination of ACM Digital Library and Google Scholar searches to learn about the effects of personality on observability. Searches for articles included keywords like, "observability", "mental model", "status", "common ground", "predictability", and "machine personality." These search terms originated from literature about observability [3].

Once we identified research that investigated how these social cues could support observability, we narrowed down the list of social cues using the following criteria:

- Cue should display system state dynamically (as opposed to static design choices like gender and name).
- Cue can be applied to both physical and virtual agents.

These criteria enabled us to focus on cues that were applicable to designers of a wide variety of systems, from assistive robots to digital agents in planning applications. After reviewing 40 sources, we selected cues with sources illustrating their relationship to observability.

We included cues from three different modality types: posture, voice, and dialogue. See Table 1 for the social cues, the definition of the cues from Feine et al. [4] and their relationship to observability.

III. ANALYSIS OF SOCIAL CUES

A. *Head Movement*

Head movement refers to a gesture or position of the head and can include nodding, shaking, tilting, looking towards or looking away from a human. Common straightforward cues include nodding to indicate agreement and shaking head to indicate disagreement [17]. Pairing a nod with an affirmative statement can make an agent's behavior seem consistent and reassure a human observer. Conversely, if the head movement is inconsistent with verbal statements, such as nodding while denying a request, the human may perceive the robot unreliable.

Poggi [18] presents more nuanced cues that can indicate a speaker's beliefs and goals, such as: nodding head to show certainty, looking up to show careful thinking, looking down obliquely to indicate trying to remember. A human may have an expectation that an agent is able to respond instantaneously. If an agent provides a cue that it is thinking, either by looking up or looking obliquely down, the human can understand that the agent needs time to construct a response. The human can adjust his or her expectations and be less frustrated by a pause in the dialogue. Humans are more likely to be patient if they know that their problem or question is being carefully considered.

Head movement may not occur in isolation, as the only non-verbal cue. Often head movement cues are combined with facial expression and eye movement to communicate beliefs and intentions. An example is a robot that combines a fixed stare with raising the inner parts of the eyebrows and a bent head to indicate "I implore you." [18].

Takeaway: Head movement can be used for more than just agreement (nod), disagreement (shake). Head movements can convey that an agent is thinking (looking upwards), is certain (nodding) or trying to remember (looking obliquely downwards).

B. *Facial Expression*

A facial expression, a form of nonverbal communication, is a movement of the face to communicate internal state like surprise, sadness, happiness, anger, etc. Internal state is demonstrated through specific facial expressions including a nod, smile, shake, frown, tension of the lips, tilt, or raise of the eyebrows [19].

For example, Cassell designs an avatar to be expressive during turn taking; when the agent is letting the person speak, it raises its eyebrows and relaxes its hands [20].

Bevacqua et al. studied the efficacy of backchanneling (a non-interrupting acoustic or visual signal demonstrated by the listener during a speaker's turn) in creating meaning through different facial expressions and acoustic cues [17]. They tried to learn which cues together could communicate interest and understanding. They found that 'interest' was conveyed by the following combination of cues: a smile and a verbal "okay," a nod and a verbal "okay," and a nod and a verbal "ooh." Likewise, 'understanding' was conveyed by nodding as well as the following combinations: raising

 ISBN: 978-1-61208-896-9

eyebrows and a verbal “ooh,” nodding and a verbal “ooh,” nodding and a verbal “really,” nodding and a verbal “yeah.”

They note that only cues together could compose meaning. In general, it is not possible to select one cue from their findings and expect it to communicate interest or understanding on its own [17].

A risk of using a facial expression could be using the wrong expression for the environment or role. For example, if a robot is showing a happy facial expression in a serious or sad environment, like a funeral, the robot might be deemed as socially inappropriate. Conversely, we speculate that a more serious facial expression in a librarian role, for example, might help make the robot belong in the role and context.

Takeaway: Facial expressions, like raising eyebrows (in conjunction with relaxing hands) can show that an agent is letting a person speak. Smiling and saying “okay,” can show interest.

C. *Voice Tempo*

Voice tempo refers to speech rate and pacing. It can be measured as seconds per syllable, for example, or as the length of pauses between words when spoken by a synthesized voice.

Speech rate, as well as pauses in speech, can communicate confidence level, which is important for observability of machines. Research on human speech perception has found that a moderately fast speech rate (ideally similar to or slightly faster than the listener’s speech rate) conveys competence, while slower pacing and pauses longer than 5 seconds indicate that the speaker is not sure of the content [21]. Therefore, voice tempo could be a valuable social cue for machines that need to communicate confidence levels and manage the human’s trust in content because trust can be calibrated appropriately by altering the tempo.

Abnormally fast speech can be associated with nervousness or urgency. Jang has shown that the speed of computerized speech does in fact convey urgency of a situation in a predictable way [22]. On the other hand, a slower speech rate might help humans to remain calm during an emergency. For instance, when a semi-automated car notices an approaching obstacle, a fast speech rate may be needed to communicate urgency, followed by a slower speech rate to orient and support the human.

Speech rates that are too fast or too slow can contribute to comprehension problems [23]. When manipulating voice tempo, in non-extreme situations, synthetic voices should tend away from extremes and toward everyday average human speech rates to be appropriate interaction partners.

Takeaway: Voice tempo can speed up to communicate urgency, but in non-emergency situations should be similar to a human’s voice tempo for easy understandability. Slowness or pauses can be used to communicate lower confidence in information. This is valuable because correctly communicating uncertainty can improve the human’s overall calibrated trust in the machine’s judgments.

D. *Pitch Range*

Pitch range refers to how high and low a synthesized voice varies from its average pitch frequency.

One study found evidence that for both Italians and Americans, having a wider pitch range made communications for people seem more exciting, interesting, and credible [24]. Similarly, a study of consumers’ impressions indicated that a wider pitch range contributes to more exciting and memorable commercials [25]. These findings suggest that an exaggerated pitch variation should be used to present important high-confidence information.

While pitch range increases the perceived competence of a message, research on vocal styles has found an inverse relationship between competence and benevolence; and there is some evidence that exaggerated pitch variation decreases perceived benevolence of a speaker, along with respect and fairness [26]. Therefore, a wider-than-normal pitch range should only be used when needed to draw the user’s attention; and should perhaps be avoided especially for systems that need to be seen as fair, such as recidivism predictors [27].

Cowell and Stanney [28] tested the effects of non-verbal cues expressed by a conversational digital assistant for helping to sort photos into albums. Characters with an appropriate pitch variation, a moderately fast speech rate (50-70 words per minute), facial expressions, and eye gaze were rated as significantly more trustworthy and credible than characters without any intelligent management of these features.

Takeaway: Use exaggerated pitch range only to draw attention to important or high-confidence information. Otherwise, use humanlike pitch range as appropriate to convey the trustworthiness and competence of the machine.

E. *Greetings and Farewells*

Greetings and farewells are expressions, or “ritual behaviors” [29], marking an agent entering or leaving an interaction. These social cues can improve trust [30] and have been found to help users to perceive an agent as more reliable, competent, and knowledgeable. Examples include saying, “Welcome”, “Nice to meet you”, “Hello”, “See you later.”

A greeting demonstrates that the target has entered the agent’s realm of activation and is being sensed or tracked. We speculate that the timing and manner of an agent greeting a person could show awareness of when a person is bored or busy. A well-timed farewell is an effective way for a person to know that the robot understands that the interaction is over.

Risks of using this social cue could be saying a farewell too early in an interaction, which could be perceived as rude or socially unaware. This is a challenge, because of the lack of social intuition that robots have. It is often difficult for a robot to know when to interrupt, which has been explored by Semmens et al., in which researchers in a car periodically asked, “Is now a good time?” They found that a system can access automotive data for knowing when to ask if it is a good time [31]. This goes to show that while robots have trouble sensing social situations, there is an opportunity to

TABLE I. SOCIAL CUES AND RELATIONSHIP TO OBSERVABILITY

Social Cue	Feine's Definition [4]	Relationship to Observability	Source
Head movement [Posture]	The agent moves its head. (I.e., nodding and turning)	In addition to more obvious indicators of agreement and disagreement, head movement can be used to indicate nuanced beliefs and goals such as confidence, thinking, and remembering	[18]
Facial expression [Posture]	The agent expresses a gesture by executing one or more motions with his facial muscles (i.e., smile or eyebrow raise)	Facial expressions, like raising eyebrows (in conjunction with relaxing hands) can show that an agent is letting a person speak.	[30]
Voice Tempo [Voice]	The pace of the agent's voice.	The speed of computerized speech conveys urgency of a situation in a predictable and systematic way, and speech pacing conveys confidence.	[21][22]
Pitch Range [Voice]	The degree of variation from the agent's average pitch. (I.e., monotone, animate voice)	Exaggerated pitch range can draw attention to important or high-confidence information. Humanlike pitch range should be appropriate to the trustworthiness and competence of the machine.	[25]
Greetings and farewells [Dialogue]	The agent expresses a word of welcome or marks someone's departure.	Small talk, which include greetings and farewells, can improve perception of an agent's good will and credibility.	[30]
Ask to start/ pursue dialogue [Dialogue]	The agent requests the user's permission to start, continue, or end the conversation	Asking to start or pursue a dialogue communicates that the human is in charge and is in support of the Human-Machine Teaming theme of Directability	[3][32]

leverage other sensors, like proximity sensors or lidar, for detecting things like if someone has left the space.

Takeaway: Greetings and farewells signal awareness of the user and that a new interaction is beginning/ending.

F. *Ask to Start / Pursue Dialogue*

Ask to start/pursue dialogue refers to behavior in which the agent seeks permission to interact with a human partner. This could take the form of initiating a conversation, as in, "My name is Indira and I can help plan tourist activities. Would you like me to look for available excursions?" It could also involve the seeking of approval to continue an interaction, such as "Would you like me to keep searching?" This social cue is tightly related to the Human-Machine Teaming theme of Directability [3], by which humans are easily able to direct and redirect an automated partner's resources, activities, and priorities. Asking to start or pursue dialogue signals that the human is in control, which is important as the human should not be removed from the command role [32]. When humans perceive a lack of control, they can become frustrated. Letting the human know that they can discontinue the agent's help shows that the agent is directable, potentially decreasing frustration and increasing user adoption.

Takeaway: When an agent asks to start or continue dialogue, it signals that the human is in control of the interaction. Provide multiple choice points in which the human can decide whether to continue dialoguing with the agent to minimize user frustration.

IV. DISCUSSION

In assessing the social cues and their relationship to observability, we found that one common risk of using social cues is setting wrong expectations. For example, if there are moving eyes on a robot, it might be perceived that a robot can see. We speculate that if it cannot see, but has eyes, then that could lead to a mistrust of an agent, which could be worse than not including the social cue at all. The use of social cues can increase the "human-ness" of a robot or software agent, but if expectations are violated it would be better to not include the social cue. Another danger of making a robot human-like is that it can enter the "uncanny valley" [33], in which robots that look a lot like humans, but not quite human, are perceived as creepy and cause revulsion. Examples include the characters in the Polar Express movie and the version of Sophia from Hanson Robotics that debuted at the 2016 South by Southwest (SXSW) conference. We posit that the social cues described in this article would not by themselves enter the uncanny valley. Rather, their use on a highly humanoid (but not convincingly human) robot platform could be disconcerting.

We also noticed that in most research about social cues, one social cue alone often does not express one piece of information; but instead, it is multiple cues in a row, or different modalities demonstrated in parallel, which communicate the desired observable behavior. For instance, one facial expression will not convey interest, but instead, a facial expression in parallel with other gestures will show interest.

Social cues may be sensitive to the cultural context in which they are used. Some cues may be universal in humans, such as the combination of lowered eyebrows, lips firmly pressed, and bulging eyes to convey anger. Other cues could be interpreted differently in separate cultures. What is perceived as friendly in the United States could be perceived as intrusive in other countries. Interpretability of social cues should be tested across cultures to ensure the gesture, movement, or vocal characteristic conveys the intended information. This could lead to the identification robot social cues that are universal.

We recommend that future work focus on the combination of social cues to convey information. In human-to-human encounters, social cues typically occur simultaneously and across modalities and are a natural part of communication. Mapping combinations of social cues to

 ISBN: 978-1-61208-896-9

observability would likewise enhance the richness of robot-to-human communication.

V. Conclusion

In this paper, we began with discussing the importance of people understanding what is going on "under the hood" of machines and the opportunities of social cues to help uncover system status for a user. Next, we shared our process of finding social cues that have the potential to improve observability in both physical and virtual agent design. We focused on cues that display system state dynamically and had supporting literature. Lastly, we did a deep dive into six social cues that met our criteria by including definitions of the cue, examples of the cue, the relationship to observability, and the associated risks of using the cue inappropriately. See table 1 for the cues organized. We found that for maximum success at communicating observability, the cues should be used in parallel with other cues.

Acknowledgment

Thank you to MITRE for funding this work; thank you to Kit Mandeville, Kelly Horinek, Alexandra Valiton, Dr. Ozgur Eris, and Dr. Beth Elson for their discussions about these topics.

Approved for Public Release; Distribution Unlimited. Public Release Case Number 21-1848.

References

[1] I. Pakrasi, N. Chakraborty, and A. LaViers, "A design methodology for abstracting character archetypes onto robotic systems," in *Proceedings of the 5th International Conference on Movement and Computing*, 2018, pp. 1–8.

[2] L. Takayama, D. Dooley, and W. Ju, "Expressing thought: improving robot readability with animation principles," in *Proceedings of the 6th international conference on Human-robot interaction*, 2011, pp. 69–76.

[3] P. McDermott et al., "Human-machine teaming systems engineering guide," The MITRE Corporation, Bedford, MA, USA, 2018. [Online]. Available: https://www.mitre.org/publications/technical-papers/human-machine-teaming-systems-engineering-guide [Accessed Sept. 13, 2021].

[4] J. Feine, U. Gnewuch, S. Morana, and A. Maedche, "A Taxonomy of Social Cues for Conversational Agents," *International Journal of Human-Computer Studies*, vol. 132, pp. 138–161, Dec. 2019, doi: 10.1016/j.ijhcs.2019.07.009.

[5] K. A. Hoff and M. Bashir, "Trust in automation: integrating empirical evidence on factors that influence trust," *Hum Factors*, vol. 57, no. 3, pp. 407–434, May 2015, doi: 10.1177/0018720814547570.

[6] R. Parasuraman and V. Riley, "Humans and automation: Use, misuse, disuse, abuse," *Human factors*, vol. 39, no. 2, pp. 230–253, 1997.

[7] J. B. Lyons, "Being transparent about transparency: A model for human-robot interaction," presented at the AAAI Spring Symposium: Trust and Autonomous Systems, 2013.

[8] J. Joe, J. O'Hara, H. Medema, and J. Oxstrand, "Identifying requirements for effective human-automation teamwork," presented at PSAM, 2014.

[9] P. J. Corr and G. Matthews, Eds., *The Cambridge Handbook of Personality Psychology*, 2nd ed. Cambridge: Cambridge University Press, 2020. doi: 10.1017/9781108264822.

[10] A. Vinciarelli and G. Mohammadi, "A Survey of Personality Computing," *IEEE Transactions on Affective Computing*, vol. 5, no. 3, pp. 273–291, Jul. 2014, doi: 10.1109/TAFFC.2014.2330816.

[11] B. J. Fogg, "Persuasive technology: using computers to change what we think and do," *Ubiquity*, vol. 2002, no. December, p. 2, 2002.

[12] C. Nass and Y. Moon, "Machines and mindlessness: Social responses to computers," *Journal of Social Issues*, vol. 56, pp. 81–103, 2000.

[13] E. J. Lobato, S. F. Warta, T. J. Wiltshire, and S. M. Fiore, "Varying social cue constellations results in different attributed social signals in a simulated surveillance task," presented at FLAIRS, 2015.

[14] L. Holzman, "Lev Vygotsky and the new performative psychology: Implications for business and organizations," *The social construction of organization*, pp. 254–268, 2006.

[15] M. F. Jung, N. Martelaro, and P. J. Hinds, "Using robots to moderate team conflict: the case of repairing violations," in *Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction*, 2015, pp. 229–236.

[16] S. S. Sebo, M. Traeger, M. F. Jung, and B. Scassellati, "The Ripple Effects of Vulnerability: The Effects of a Robot's Vulnerable Behavior on Trust in Human-Robot Teams," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, New York, NY, USA, Feb. 2018, pp. 178–186. doi: 10.1145/3171221.3171275.

[17] E. Bevacqua, S. Pammi, S. J. Hyniewska, M. Schröder, and C. Pelachaud, "Multimodal backchannels for embodied conversational agents," in *International Conference on Intelligent Virtual Agents*, 2010, pp. 194–200.

[18] I. Poggi, "Mind markers," *The Semantics and Pragmatics of Everyday Gestures. Berlin Verlag Arno Spitz*, 2001.

[19] D. Heylen, E. Bevacqua, M. Tellier, and C. Pelachaud, "Searching for prototypical facial feedback signals," in *International Workshop on Intelligent Virtual Agents*, 2007, pp. 147–153.

[20] J. Cassell, "Embodied conversational interface agents," *Communications of the ACM*, vol. 43, no. 4, pp. 70–78, 2000.

[21] R. L. Street and R. M. Brady, "Speech rate acceptance ranges as a function of evaluative domain, listener speech rate, and communication context," *Communication Monographs*, vol. 49, no. 4, pp. 290–308, Dec. 1982, doi: 10.1080/03637758209376091.

[22] P.-S. Jang, "Designing acoustic and non-acoustic parameters of synthesized speech warnings to control perceived urgency," *International Journal of Industrial Ergonomics*, vol. 37, no. 3, pp. 213–223, Mar. 2007, doi: 10.1016/j.ergon.2006.10.018.

[23] E. Rodero, "Do Your Ads Talk Too Fast To Your Audio Audience?: How Speech Rates of Audio Commercials Influence Cognitive and Physiological Outcomes," *Journal of Advertising Research*, vol. 60, no. 3, pp. 337–349, Sep. 2020, doi: 10.2501/JAR-2019-038.

[24] M. Urbani, "The Pitch Range of Italians and Americans. A Comparative Study," University of Padua, 2013. [Online].

Available: http://paduaresearch.cab.unipd.it/5976/ [Accessed Sept. 13, 2021].

[25] E. Rodero, R. F. Potter, and P. Prieto, "Pitch Range Variations Improve Cognitive Processing of Audio Messages," *Human Communication Research*, vol. 43, no. 3, pp. 397–413, Jul. 2017, doi: 10.1111/hcre.12109.

[26] P. Rockwell, "The effects of vocal variation on listener recall," *J Psycholinguist Res*, vol. 25, no. 3, pp. 431–441, May 1996, doi: 10.1007/BF01727001.

[27] J. Dressel and H. Farid, "The accuracy, fairness, and limits of predicting recidivism," *Science Advances*, vol. 4, no. 1, p. eaao5580, Jan. 2018, doi: 10.1126/sciadv.aao5580.

[28] A. J. Cowell and K. M. Stanney, "Manipulation of non-verbal interaction style and demographic embodiment to increase anthropomorphic computer character credibility," *International Journal of Human-Computer Studies*, vol. 62, no. 2, Art. no. 2, Feb. 2005, doi: 10.1016/j.ijhcs.2004.11.008.

[29] J. Cassell and H. Vilhjálmsson, "Fully embodied conversational avatars: Making communicative behaviors autonomous," *Autonomous agents and multi-agent systems*, vol. 2, no. 1, pp. 45–64, 1999.

[30] J. Cassell and T. Bickmore, "External manifestations of trustworthiness in the interface," *Communications of the ACM*, vol. 43, no. 12, pp. 50–56, 2000.

[31] R. Semmens, N. Martelaro, P. Kaveti, S. Stent, and W. Ju, "Is now a good time? An empirical study of vehicle-driver communication timing," in *Proceedings of the 2019 chi conference on human factors in computing systems*, 2019, pp. 1–12.

[32] C. E. Billings, *Aviation automation: The search for a human-centered approach*. CRC Press, 2018.

[33] M. Mori, K. F. MacDorman, and N. Kageki, "The uncanny valley [from the field]," *IEEE Robotics & Automation Magazine*, vol. 19, no. 2, pp. 98–100, 2012.

 ISBN: 978-1-61208-896-9