Social Signal Modeling in Human-Robot Interaction

Maia Stiber
Johns Hopkins University
Baltimore, Maryland, USA
mstiber@jhu.edu

Micol Spitale*
Politecnico di Milano
Milan, Italy
micol.spitale@polimi.it

Hatice Gunes
University of Cambridge
Cambridge, UK
hatice.gunes@cl.cam.ac.uk

Chien-Ming Huang
Johns Hopkins University
Baltimore, Maryland, USA
chienming.huang@jhu.edu

ABSTRACT

This workshop focuses on the understanding and modeling of social signals to create human-aware HRI. The three fundamental themes are: understanding social signals (gain insights into human internal states), modeling social signals for the generation of a human's mental state (translating social signals into actionable computational models), and operationalizing human models for human-aware applications (integrating these cognitive models into robotic systems to develop new human-aware capabilities). The invited speakers, paper presentations, and discussions will aim to focus on the social science background of social signals, acquisition and availability of benchmarking datasets, social signal modeling techniques, integration of models into real-time systems, usage of these models—such as error management, personalization, and mental model alignment—and applications of these models (i.e., healthcare, education, manufacturing). We expect these topics to demonstrate how modeling social signals, both explicit and implicit, is necessary for fluent, intuitive and trustworthy interactions.

CCS CONCEPTS

• **Human-centered computing**; • **Computer systems organization** → **Robotics**;

KEYWORDS

Social Signals, Human Behavior Modeling, HRI, Datasets

ACM Reference Format:
Maia Stiber, Micol Spitale, Hatice Gunes, and Chien-Ming Huang. 2024. Social Signal Modeling in Human-Robot Interaction. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24 Companion), March 11–14, 2024, Boulder, CO, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3610978.3638163

*This work was partially undertaken while Micol Spitale was a postdoctoral researcher at the University of Cambridge.

HRI '24 Companion, March 11–14, 2024, Boulder, CO, USA

ACM ISBN 979-8-4007-0323-2/24/03.
https://doi.org/10.1145/3610978.3638163

1 WORKSHOP MOTIVATION

Fluent human-robot interactions require robotic systems that can understand, predict, and adapt to human behavior [9, 15]. Robots without human-awareness can not anticipate human behaviors or adapt to their users' needs, causing interaction breakdowns or even hazardous interactions [10]. One solution to this problem is giving the robot an internal model of the human [13, 20]. Through understanding and subsequently modeling social signals, we can develop models of human intention, human mental models, and human internal cognitive states [7, 20].

There are three elements to developing these human-aware capabilities and enhancing fluency in HRI: identifying and understanding social signals, modeling social signals to develop human mental state estimation, and developing human-aware capabilities using the human estimation model.

Social signals, both implicit and explicit, reflect a user's intention, internal state, and mental model of the interaction that they are involved in (i.e., environment, task, robot) [4]. These signals, although noisy, serve as observable manifestations of these hidden human internal cognitive states. In HRI, social signals, such as body movements (e.g. [21]), facial expressions (e.g., [3], [14]), gaze (e.g., [1]), EEG (e.g., [8]), and verbal utterances (e.g., [12]), have been used to develop these human state approximations for human-aware capabilities. To use social signals effectively, one must determine what specific social signals should be employed for development of human mental state estimation and understand how social signals reflect these internal states. In addition, researchers need to develop collaborative repositories to collect a broader range of social signal data in HRI and sharing them with the community. This is crucial to ensuring that our understanding of social signals is comprehensive across various interaction scenarios.

By **modeling social signals**, we are generating an approximation of a human's mental state and mental models that can be operationalized in HRI systems. This estimation captures both the inherent variability of human responses during HRI and underlying consistent behavioral patterns across tasks and humans. To create these human estimation models, researchers must collaboratively determine what modeling techniques should be used, how we can share these models for the community, and what social signal dataset benchmarks to use. Moreover, we should consider whether the human models would be improved by including other inputs beyond social signals (such as context or human characteristics) as some social signals could by influenced by task type [12] or

personal characteristics (i.e., culture can affect what social signals are expressed [11]).

The integration of these **human estimation models** into robotic systems enables development of fluent human-aware interactions. By applying these models, we can achieve functionalities such as detecting robot mistakes (e.g., [18], [17], [2]), determining the right time to engage a human, discerning user intentions, and establishing models of user trust. These models can also be used to generate robot behaviors (e.g., [16], [17], [2]) for fluent interactions. Notably, these models can be used to both broaden capabilities to be generalized across users, tasks, and environments (e.g., [19]) and to be personalized to each user, task, and environment (e.g., [22], [14]). Consistent modeling of human behaviors also lays the foundation for continually updating internal state models [6], further strengthening prolonged engagements in systems where humans are active participants [5]. That being said, it is still an open question what HRI capabilities can best benefit from these human models.

The objective of this workshop is to demonstrate that modeling social signals can promote fluent HRI through development of human-aware capabilities. Thoughtful discussions through invited speakers, and a poster session and presentations from researchers (but not limited to) will address acquisition and availability of benchmarking datasets, models and systems, human behavior model generation, applications for these models, and real-time systems incorporating these models.

2 ORGANIZERS

Maia Stiber (Johns Hopkins University, USA) is a PhD student in the Intuitive Computing Lab at Johns Hopkins University advised by Professors Chien-Ming Huang and Russell Taylor. Her research focuses on understanding and modeling implicit behavioral responses to robot actions to allow the robot to detect robot errors. She has been awarded the JHU Computer Science Department Fellowship and the Jay D. Samstag Engineering Fellowship.

Micol Spitale (Politecnico di Milano, Italy) is currently an Assistant Professor at the Department of Electronics, Information and Bioengineering at the Politecnico di Milano (Polimi), as well as a Visiting Affiliated Researcher at the University of Cambridge. Her research activities are grounded in the Social Robotics area. Her current research focuses on developing socio-emotionally adaptive robots that can foster wellbeing through coaching and psychologically proven interventions. She has been awarded "cum laude" a Ph.D. in Information Technology, Computer Science and Engineering Area at the Politecnico di Milano, co-funded by IBM Italy and EIT Digital, in October 2021. During her Ph.D., she spent several months at the University of Southern California (USC) in the Interaction Lab as a visiting Ph.D. student.

Hatice Gunes (University of Cambridge, UK) is a Full Professor of Affective Intelligence and Robotics (AFAR) and the Director of the AFAR Lab at the University of Cambridge. Her expertise is in the areas of affective computing and social signal processing cross-fertilizing research in multimodal interaction, human behaviour understanding, machine learning and social robotics. She has published over 175 papers in these areas (H-index=39, citations > 8,300), with most recent works on nonverbal generative AI, graph representation learning and lifelong learning for affective computing and affective robotics, bias mitigation and fairness, human-robot interaction (HRI) for service robotics and longitudinal HRI for wellbeing.

Chien-Ming Huang (Johns Hopkins University, USA) is the John C. Malone Assistant Professor in the Department of Computer Science at the Johns Hopkins University. His research focuses on designing interactive AI aimed to assist and collaborate with people. Huang completed his postdoctoral training at Yale University and received his Ph.D. in Computer Science at the University of Wisconsin–Madison. He is a recipient of the NSF CAREER award.

3 WORKSHOP OVERVIEW

The workshop will consist of paper presentations, invited talks, and discussion sessions. Paper presentations, from researchers whose papers were chosen for this medium, will consist of seven minute presentations and three minutes Q&A sessions. They will be in groups of three. Invited speakers time slots will consist of 30 minute presentations and 15 minute Q&A and discussion session after a set of two have gone.

3.1 Workshop Topics

We will discuss state-of-the-art approaches in modeling social signals as well as applications for these models. We aim to promote discussion and assemble researchers in this domain to foster collaboration, share datasets, and develop new opportunities in the field. The workshop welcomes contributions across a wide range of topics including, but not limited to:

- **Identifying and Understanding Social Signals**:
  - Social science perspective to social signals
  - Behavioral signal modalities
  - Social signal datasets and repositories
- **Modeling Social Signals** for developing human state estimation:
  - Modeling techniques
  - Datasets and benchmarking
  - Public code and repositories
- **Operationalizing Human State Models** for human-aware HRI:
  - Usage of social signal models for:
    - Generalization
    - Personalization
    - Mental model alignment
    - Robot error management and repair
    - Long-term Interactions
  - Applications of social signal models for:
    - Home Assistance
    - Education
    - Healthcare
    - Manufacturing
    - Retail and Warehouse
  - Real-time system integration
- Future directions and opportunities

3.2 Invited Speakers

We have invited the following speakers:

- **Henny Admoni**, Associate Professor, Carnegie Mellon University
- **Chung Hyuk Park**, Associate Professor, George Washington University
- **Yukie Nagai**, Project Professor, The University of Tokyo, Japan

3.3 Tentative Schedule

The tentative schedule is as follows:

- 1:00 - 1:05pm: Welcoming Remarks
- 1:05 - 1:35pm: Invited Speaker #1
- 1:35 - 2:05pm: Invited Speaker #2
- 2:05 - 2:20pm: Session Discussion
- 2:25 - 2:40pm: Coffee Break (aligned with the conference)
- 2:40 - 3:10pm: Invited Speaker #3
- 3:10 - 3:40pm: Invited Speaker #4
- 3:40 - 3:55pm: Session Discussion
- 4:00 - 4:55pm: Papers selected for oral presentation
- 4:55 - 5:00pm: Closing Remarks

The above schedule is flexible and we will accommodate for changes depending on the number of accepted papers and conference schedule with respect to timing of coffee breaks.

4 TARGET AUDIENCE AND EXPECTED PARTICIPANTS

This workshop aims to bring together researchers—who are interested in advancing the understanding of social signal modeling and its application in HRI systems, designing and development of robotic systems, and discussing data-driven approaches to modeling and interpreting social signals. We want to encourage an interdisciplinary approach and include participants involved in HRI, machine learning, and affective computing.

4.1 Approach for Recruiting Participants

We will publicize this workshop both within our networks, through several mailing lists (e.g., robotics-worldwide, hri-announcements), and through social media (e.g., linkedin). In addition, we will reach out to investigators in human behavior modeling, machine learning, cognitive science, and robotics.

5 WORKSHOP DOCUMENTATION

We will record all presentations and discussions and make them available online after being granted permission by the speakers. In addition, we plan to make the accepted papers available on ArXiv given permission from the authors. Building upon this workshop, we will seek to organize a special issue for a journal or a special session at a conference.

ACKNOWLEDGEMENTS

M. Stiber and C.-M. Huang are supported by the National Science Foundation award #2143704. H. Gunes is supported by the EPSRC/UKRI under grant ref. EP/R030782/1 (ARoEQ).

REFERENCES

[1] Reuben M. Aronson and Henny Admoni. 2018. Gaze for error detection during human-robot shared manipulation.
[2] Minja Axelsson, Micol Spitale, and Hatice Gunes. 2024. "Oh, Sorry, I Think I Interrupted You": Designing Repair StrategiFes for Robotic Longitudinal Well-being Coaching. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction.*
[3] Chandrayee Basu and Mukesh Singhal. 2016. Trust dynamics in human autonomous vehicle interaction: a review of trust models. In *2016 AAAI spring symposium series.*
[4] Alexandra Bremers, Alexandria Pabst, Maria Teresa Parreira, and Wendy Ju. 2023. Using Social Cues to Recognize Task Failures for HRI: A Review of Current Research and Future Directions. *arXiv preprint arXiv:2301.11972* (2023).
[5] Nikhil Churamani, Minja Axelsson, Atahan Caldır, and Hatice Gunes. 2022. Continual learning for affective robotics: A proof of concept for wellbeing. In *2022 10th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW).* IEEE, 1–8.
[6] Nikhil Churamani, Sinan Kalkan, and Hatice Gunes. 2020. Continual learning for affective robotics: Why, what and how?. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN).* IEEE, 425–431.
[7] Zoran Duric, Wayne D Gray, Ric Heishman, Fayin Li, Azriel Rosenfeld, Michael J Schoelles, Christian Schunn, and Harry Wechsler. 2002. Integrating perceptual and cognitive modeling for adaptive and intelligent human-computer interaction. *Proc. IEEE* 90, 7 (2002), 1272–1289.
[8] Stefan Ehrlich, Agnieszka Wykowska, Karinne Ramirez-Amaro, and Gordon Cheng. 2014. When to engage in interaction—And how? EEG-based enhancement of robot's ability to sense social signals in HRI. In *2014 IEEE-RAS International Conference on Humanoid Robots.* IEEE, 1104–1109.
[9] Guy Hoffman and Cynthia Breazeal. 2007. Cost-based anticipatory action selection for human–robot fluency. *IEEE transactions on robotics* 23, 5 (2007), 952–961.
[10] Guy Hoffman and Cynthia Breazeal. 2007. Effects of anticipatory action on human-robot teamwork efficiency, fluency, and perception of team. In *Proceedings of the ACM/IEEE international conference on Human-robot interaction.* 1–8.
[11] Rachael E Jack, Oliver GB Garrod, Hui Yu, Roberto Caldara, and Philippe G Schyns. 2012. Facial expressions of emotion are not culturally universal. *Proceedings of the National Academy of Sciences* 109, 19 (2012), 7241–7244.
[12] Dimosthenis Kontogiorgos, Minh Tran, Joakim Gustafson, and Mohammad Soleymani. 2021. A systematic cross-corpus analysis of human reactions to robot conversational failures. In *Proceedings of the 2021 International Conference on Multimodal Interaction.* 112–120.
[13] Joseph B Lyons and Paul R Havig. 2014. Transparency in a human-machine context: Approaches for fostering shared awareness/intent. In *Virtual, Augmented and Mixed Reality. Designing and Developing Virtual and Augmented Environments: 6th International Conference, VAMR 2014, Held as Part of HCI International 2014, Heraklion, Crete, Greece, June 22-27, 2014, Proceedings, Part I 6.* Springer, 181–190.
[14] Emily McQuillin, Nikhil Churamani, and Hatice Gunes. 2022. Learning Socially Appropriate Robo-waiter Behaviours through Real-time User Feedback. In *ACM/IEEE International Conference on Human-Robot Interaction, HRI 2022, Sapporo, Hokkaido, Japan, March 7 - 10, 2022.* 541–550.
[15] Amit Kumar Pandey, Lavindra de Silva, and Rachid Alami. 2016. A human-robot competition: Towards evaluating robots' reasoning abilities for hri. In *Social Robotics: 8th International Conference, ICSR 2016, Kansas City, MO, USA, November 1-3, 2016 Proceedings 8.* Springer, 138–147.
[16] Hannes Ritschel and Elisabeth André. 2017. Real-time robot personality adaptation based on reinforcement learning and social signals. In *Proceedings of the companion of the 2017 acm/ieee international conference on human-robot interaction.* 265–266.
[17] Micol Spitale, Minja Axelsson, and Hatice Gunes. 2023. VITA: A Multi-modal LLM-based System for Longitudinal, Autonomous, and Adaptive Robotic Mental Well-being Coaching. arXiv:2312.09740 [cs.RO]
[18] Micol Spitale, Minja Axelsson, Neval Kara, and Hatice Gunes. 2023. Longitudinal Evolution of Coachees' Behavioural Responses to Interaction Ruptures in Robotic Positive Psychology Coaching. In *2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN).* IEEE.
[19] Maia Stiber, Russell H Taylor, and Chien-Ming Huang. 2023. On using social signals to enable flexible error-aware HRI. In *Proceedings of the 2023 ACM/IEEE International Conference on Human-Robot Interaction.* 222–230.
[20] Aaquib Tabrez, Matthew B Luebbers, and Bradley Hayes. 2020. A survey of mental modeling techniques in human–robot teaming. *Current Robotics Reports* 1 (2020), 259–267.
[21] Samuele Vinanzi, Angelo Cangelosi, and Christian Goerick. 2020. The role of social cues for goal disambiguation in human-robot cooperation. In *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN).* IEEE, 971–977.
[22] Baijun Xie and Chung Hyuk Park. 2021. A MultiModal Social Robot Toward Personalized Emotion Interaction. *arXiv preprint arXiv:2110.05186* (2021).