A computational lens to interaction-shaping robotics

Sarah Gillet
KTH Royal Institute Of Technology
Stockholm, Sweden
sgillet@kth.se

I. Introduction

Positive human-human interactions contribute to feelings of happiness, security, and self-esteem [30, 16]. On the contrary, negative or nonexistent relationships can lead to feelings of social rejection, loneliness, and poorer physical health [20, 38]. Recent research in the Human-Robot Interaction (HRI) community has begun to explore robots interacting with groups of people [35, 34] and found that robots can influence people beyond the reciprocal interaction between a person and a robot. For example, robots can influence how people work [39], talk [41], or connect with each other [33, 42, 6]. Therefore, we coined the term interaction-shaping robotics to refer to robots that shape interactions between other agents, e.g., people [12].

My research vision aims to advance autonomous social robots that can interact with multiple people and adapt to human-human dynamics to positively shape interactions between people. For example, a robot might use subtle gaze behaviors to balance how much people participate in a group discussion [9]. Therefore, my research has addressed both computational approaches, e.g., reinforcement learning (RL), and explored their effects on human-human dynamics, i.e., the balance in participation, which we evaluated in user studies. Thereby, I am focusing on computational approaches that help the robot act in socially appropriate ways, i.e., avoiding inappropriate or biased actions in inappropriate moments.

Relevant prior work has often taken a social psychology approach to study robot behaviors that shape human-human interactions. These works have demonstrated that robots can improve conflict situations among adults [23] and children [37]. They can enhance emotional support [6], foster the expression of vulnerability [39], and support first contact between strangers [33] or the process of inclusion among adults [40]. Further, research has studied how robots shape participation behavior [5, 41, 27]. Additional works explored how to perceive groups and their dynamics [7], e.g., cohesion [36], or dominance [31], and how to apply machine learning techniques to other HRI applications. For example, imitation learning was used for kinesthetic teaching of manipulation skills [1, 26] or active listening behavior [22]. Other work explored RL to personalize robot behavior [25, 24, 43, 32, 14], increase engagement [21], or interact with bypassers [29].

The connection between the perception of group dynamics and robot shaping behaviors, for example, through heuristics or machine learning, remains largely unexplored. One prior work addressing this connection perceives the amount of speech in a discussion and uses a microphone-shaped robot to encourage the least active participant by turning toward them [41].

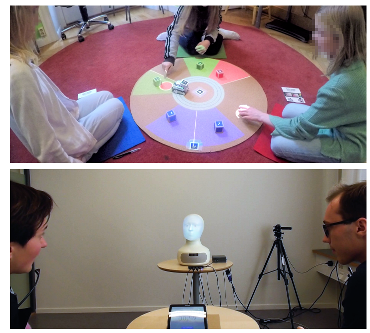

Fig. 1. Two examples of robots interacting in groups while perceiving and acting on the current group dynamics. Cozmo is fostering inclusion and collaboration through perceiving participation in the game (top). Using audio data, the robot Furhat utilizes gaze to encourage equal participation (bottom).

My contributions address this gap between the perception of group dynamics and interaction-shaping robot behaviors by exploring (1) how handcrafted behavior heuristics that perceive human-human dynamics can shape human-human interactions and (2) how machine learning techniques that reduce the need for handcrafting can be applied while ensuring that robot behaviors remain appropriate.

II. From handcrafting to learning behaviors

Fostering inclusion and collaboration: We studied whether a robot could help the inclusion of children that newly arrived in a country [8]. We developed an interactive music-mixing game played by a group of three children, which was mediated by the robot (see Figure 1-top) In this work, we approximated the group dynamics among children through their participation behaviors in the game. Using camera-based tracking of tangible game elements, the robot could collect these participation behaviors to actively prompt the least active child to take action in the game. In a control condition, the robot did not perceive the group dynamics and randomly chose which child to prompt. The analysis of children's interactions indicated that the robot could perceive the group's dynamic. Moreover, the robot could encourage newly arrived children to play more outgoing, and increase collaboration between the other children even beyond the interaction with the robot. In addition, already present children tended to be more prosocial when giving away stickers in a mini-dictator game [44, 15]. However, the robot's behavior and perception of the group

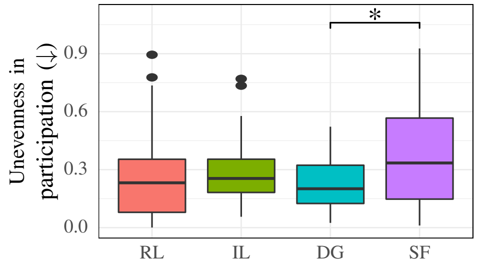

Fig. 2. Unevenness in participation for two handcrafted robot behaviors on the right (N=27 groups, DG: dynamic gaze, SF: speaker-follower) and gaze behaviors trained through RL or IL on the left (N=24 groups, RL: Reinforcement Learning, IL: Imitation Learning). ↓: Lower values are better.

dynamics were purposefully engineered and can only be used in comparable scenarios.

Balancing participation through gaze: We further explored suitable non-verbal behaviors that could balance participation in skill-imbalanced groups [9] to overcome the limitation of game-specific robot behavior and perception. In the second scenario we developed (see Figure 1-bottom), we used a human-focused approximation of group dynamics by collecting voice activation shares of the group members (collected automatically through individual close-talk microphones) to dynamically adjust handcrafted gaze patterns. To ensure a skill imbalance in the group, we paired language learners and native speakers to form a group with the fully autonomous robot. We used a skill-focused game in which participants describe words and the robot takes the unique role of guessing the words. We show that the adjusting and dynamic gaze behavior (DG), as opposed to a speaker-following (SF) gaze behavior, can lead to balanced participation. Figure 2 shows how the measure of unevenness in participation resulted in lower values, meaning more even participation in the DG condition.

Learning behavior policies for balancing participation: We formulated the problem of balancing participation through gaze as a Markov Decision Process (MDP) to be solved through imitation learning (IL) and RL [11]. The goal was to train adjusting and flexible gaze policies, overcoming the need for handcrafting robot behavior as done in previous work. We trained a gaze policy $\pi : \mathbf{s}_t \mapsto a_t$ that uses raw audio features (e.g. MFCC, pitch, ...) of participants as the state and the gaze options 'Look at speaker', ´'Look at other', 'Gaze aversion' and 'Do nothing' as the actions. For the policy π_{RL}, we computed the reward to be proportional to the balance of speech. Due to its potential to outperform behaviors from the dataset, we used offline RL in the form of Double Deep Q-learning [17] based on the data collected through the gaze heuristics in prior work for training.

In a user study, we compared the policy π_{RL} to the policy π_{IL} trained through IL, i.e., behavior cloning. For behavioral cloning, we used the same dataset with the actions that the heuristics chose as the ground truth. None of the behavior policies improved over the heuristics regarding the unevenness in participation as visualized in Figure 2. However, the user study showed promise for learning gaze behaviors for interacting in groups as interactions were not compromised. We were the first to formulate the task of shaping group interactions as an RL problem and showed that offline methods, which allow for evaluating the policy before deployment, can serve to overcome the limitations of handcrafting while ensuring socially appropriate behaviors.

III. Ongoing & Future work

GNNs to Model Human-Human Dynamics: The MDP in our prior work above still has handcrafted elements, such as specific 'Look at' actions encoding the addressees (speaker, other). Due to their ability to explicitly reason on the interactions between people, we aim to use Graph Neural Networks (GNNs) [13, 4] where humans are represented as nodes and their interactions as edges to choose *who* the robot should address through node-level prediction. In preliminary work, we explored choosing the addressee of the robot's action through behavioral cloning and compared the use of linear models and GNNs [10]. Given the complexity of the human-human dynamics, the F1 score was low for both approaches. However, the experiments showed that the GNNs outperformed the linear models while being smaller in network size. My future work will deepen the study of GNNs for modeling human-human dynamics of groups of varying sizes. With the goal of balancing participation, we will explore how to combine addressee selection and multimodal actions through GNNs and template-based action spaces [18].

SafeRL for Socially Appropriate Exploration: In previous work, we showed that offline methods allow for ensuring socially appropriate behavior. However, these methods always require a dataset to be used for training. Therefore, we are exploring shielding [2] to allow only "safe", i.e., socially appropriate behaviors. We decided to focus on one-on-one interactions to reduce the complexity of the problem and attentive listening in the form of backchanneling. We use a conversational dataset [3] in combination with the concept of backchanneling relevant spaces [19] to explore if a data-driven approach to shielding could generate a shield that limits a randomly exploring RL agent to backchannel only in appropriate moments. This work will lay the foundation for investigating how other types of appropriateness relating to action types or addressees could allow for exploring and learning more complex robot behavior online.

Bias in robot behavior: While striving for less handcrafting of robot behavior, we risk overlooking biases in human behavior that we captured in datasets, e.g., resulting in different robot behavior toward men and women [28]. Therefore, one goal of my future work is to further explore the risk of copying human biases into robot behavior and find strategies to mitigate those risks. For example, we could leverage methods from previous work, such as offline learning techniques, to evaluate the robot's behavior before deployment or explore if we can generate shields that lead to appropriate behavior and avoid developing biases. A combination of these methods might be necessary to ensure that robot behaviors remain unbiased.

REFERENCES

[1] Baris Akgun, Maya Cakmak, Jae Wook Yoo, and Andrea Lockerd Thomaz. Trajectories and keyframes for kinesthetic teaching: A human-robot interaction perspective. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*, pages 391–398, 2012.

[2] Mohammed Alshiekh, Roderick Bloem, Rüdiger Ehlers, Bettina Könighofer, Scott Niekum, and Ufuk Topcu. Safe reinforcement learning via shielding. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), Apr. 2018.

[3] Andrew J. Aubrey, David Marshall, Paul L. Rosin, Jason Vandeventer, Douglas W. Cunningham, and Christian Wallraven. Cardiff conversation database (ccdb): A database of natural dyadic conversations. In *2013 IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 277–282, 2013.

[4] Peter Battaglia, Jessica Blake Chandler Hamrick, Victor Bapst, Alvaro Sanchez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. 2018.

[5] Vicky Charisi, Luis Merino, Marina Escobar, Fernando Caballero, R Gomez, and E Gómez. The effects of robot cognitive reliability and social positioning on child-robot team dynamics. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021.

[6] Hadas Erel, Denis Trayman, Chen Levy, Adi Manor, Mario Mikulincer, and Oren Zuckerman. Enhancing emotional support: The effect of a robotic object on human–human support quality. *International Journal of Social Robotics*, pages 1–20, 2021.

[7] Daniel Gatica-Perez. Automatic nonverbal analysis of social interaction in small groups: A review. *Image and Vision Computing*, 27(12):1775–1787, 2009. ISSN 0262-8856. Visual and multimodal analysis of human spontaneous behaviour:.

[8] Sarah Gillet, Wouter van den Bos, and Iolanda Leite. A social robot mediator to foster collaboration and inclusion among children. In *Proceedings of Robotics: Science and Systems*, Corvalis, Oregon, USA, July 2020. doi: 10.15607/RSS.2020.XVI.103.

[9] Sarah Gillet, Ronald Cumbal, André Pereira, José Lopes, Olov Engwall, and Iolanda Leite. Robot Gaze Can Mediate Participation Imbalance in Groups with Different Skill Levels. In *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, New York, NY, USA, 3 2021. ACM. ISBN 9781450382892. doi: 10.1145/3434073.3444670. URL https://doi.org/10.1145/3434073.3444670.

[10] Sarah Gillet, Iolanda Leite, and Marynel Vázquez. A graph neural network approach for choosing robot addressees in group human-robot interactions. 2022.

[11] Sarah Gillet, Maria Teresa Parreira, Marynel Vázquez, and Iolanda Leite. Learning gaze behaviors for balancing participation in group human-robot interactions. In *Proceedings of the 2022 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '22, page 265–274. IEEE Press, 2022.

[12] Sarah Gillet, Marynel Vázquez, Sean Andrist, Iolanda Leite, and Sarah Sebo. Interaction-shaping robotics: Robots that influence interactions between other agents. *J. Hum.-Robot Interact.*, 13(1), mar 2024. doi: 10.1145/3643803. URL https://doi.org/10.1145/3643803.

[13] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International conference on machine learning*, pages 1263–1272. PMLR, 2017.

[14] Goren Gordon, Samuel Spaulding, Jacqueline Korywestlund, Jin Joo Lee, Luke Plummer, Marayna Martinez, Madhurima Das, and Cynthia Breazeal. Affective personalization of a social robot tutor for children's second language skills. *30th AAAI Conference on Artificial Intelligence, AAAI 2016*, (2011):3951–3957, 2016.

[15] Berna Güroglu, Wouter van den Bos, and Eveline A. Crone. Sharing and giving across adolescence: An experimental study examining the development of prosocial behavior. *Frontiers in Psychology*, 5(APR):1–13, 2014.

[16] Michelle A Harris and Ulrich Orth. The link between self-esteem and social relationships: A meta-analysis of longitudinal studies. *Journal of personality and social psychology*, 119(6):1459, 2020.

[17] Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI'16, page 2094–2100. AAAI Press, 2016.

[18] Matthew Hausknecht, Prithviraj Ammanabrolu, Marc-Alexandre Côté, and Xingdi Yuan. Interactive Fiction Games: A Colossal Adventure. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7903–7910, Apr. 2020.

[19] Mattias Heldner, Anna Hjalmarsson, and Jens Edlund. Backchannel relevance spaces. In *Nordic Prosody XI, Tartu, Estonia, 15-17 August, 2012*, pages 137–146. Peter Lang Publishing Group, 2013.

[20] Julianne Holt-Lunstad. Why social relationships are important for physical health: A systems approach to understanding and modifying risk and protection. *Annual review of psychology*, 69:437–458, 2018.

[21] Nusrah Hussain, Engin Erzin, T. Metin Sezgin, and Yücel Yemez. Training socially engaging robots: Modeling backchannel behaviors with batch reinforcement learning. *IEEE Transactions on Affective Computing*, 13(4): 1840–1853, 2022.

[22] Vidit Jain, Maitree Leekha, Rajiv Ratn Shah, and Jainendra Shukla. Exploring semi-supervised learning for predicting listener backchannels. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–12, 2021.

[23] Malte F. Jung, Nikolas Martelaro, and Pamela 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*, page 229–236, Portland, Oregon, USA, 2015. Association for Computing Machinery. ISBN 9781450328838.

[24] Iolanda Leite, André Pereira, Ginevra Castellano, Samuel Mascarenhas, Carlos Martinho, and Ana Paiva. Modelling Empathy in Social Robotic Companions. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 7138 LNCS, pages 135–147. 2012.

[25] Noriaki Mitsunaga, Christian Smith, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. Robot Behavior Adaptation for Human-Robot Interaction based on Policy Gradient Reinforcement Learning. *Journal of the Robotics Society of Japan*, 24(7):820–829, 2006. ISSN 0289-1824.

[26] Carl L Mueller and Bradley Hayes. Safe and robust robot learning from demonstration through conceptual constraints. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pages 588–590, 2020.

[27] Bilge Mutlu, Toshiyuki Shiwa, Takayuki Kanda, Hiroshi Ishiguro, and Norihiro Hagita. Footing in human-robot conversations: how robots might shape participant roles using gaze cues. In *Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*, pages 61–68. ACM, 2009.

[28] Maria Teresa Parreira, Sarah Gillet, Katie Winkle, and Iolanda Leite. How did we miss this? a case study on unintended biases in robot social behavior. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction*, HRI '23, page 11–20, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9781450399708. doi: 10.1145/3568294.3580032. URL https://doi.org/10.1145/3568294.3580032.

[29] Ahmed Hussain Qureshi, Yutaka Nakamura, Yuichiro Yoshikawa, and Hiroshi Ishiguro. Intrinsically motivated reinforcement learning for human–robot interaction in the real-world. *Neural Networks*, 107:23–33, 2018. ISSN 18792782.

[30] Harry T Reis, W Andrew Collins, and Ellen Berscheid. The relationship context of human behavior and development. *Psychological bulletin*, 126(6):844, 2000.

[31] Rutger Rienks and Dirk Heylen. Dominance detection in meetings using easily obtainable features. In *International Workshop on Machine Learning for Multimodal Interaction*, pages 76–86. Springer, 2005.

[32] Hannes Ritschel, Tobias Baur, and Elisabeth Andre. Adapting a Robot's linguistic style based on socially-Aware reinforcement learning. *RO-MAN 2017 - 26th IEEE International Symposium on Robot and Human Interactive Communication*, 2017-Janua:378–384, 2017.

[33] Ofir Sadka, Alon Jacobi, Andrey Grishko, Udi Lumnitz, Benny Megidish, and Hadas Erel. " by the way, what's your name?": The effect of robotic bar-stools on human-human opening-encounters. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*, pages 1–6, 2022.

[34] Eike Schneiders, EunJeong Cheon, Jesper Kjeldskov, Matthias Rehm, and Mikael B. Skov. Non-dyadic interaction: A literature review of 15 years of human-robot interaction conference publications. *J. Hum.-Robot Interact.*, 11(2), 2022.

[35] Sarah Sebo, Brett Stoll, Brian Scassellati, and Malte F Jung. Robots in Groups and Teams: A Literature Review. *Proc. ACM Hum.-Comput*, 4(October):37, 2020.

[36] Garima Sharma, Shreya Ghosh, and Abhinav Dhall. Automatic Group Level Affect and Cohesion Prediction in Videos. *2019 8th International Conference on Affective Computing and Intelligent Interaction Workshops and Demos, ACIIW 2019*, (April 2020):161–167, 2019.

[37] Solace Shen, Petr Slovak, and Malte F. Jung. "Stop. I See a Conflict Happening.". In *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, pages 69–77, New York, NY, USA, 2 2018. ACM. ISBN 9781450349536.

[38] George M Slavich. Social safety theory: a biologically based evolutionary perspective on life stress, health, and behavior. *Annual review of clinical psychology*, 16:265, 2020.

[39] Sarah Strohkorb Sebo, Margaret Traeger, Malte F. Jung, and Brian Scassellati. The Ripple Effects of Vulnerability: The Effects of a Robot's Vulnerable Behavior on Trust in Human-Robot Teams. *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction - HRI '18*, (February):178–186, 2018. ISSN 21672148.

[40] Sarah Strohkorb Sebo, Ling Liang Dong, Nicholas Chang, and Brian Scassellati. Strategies for the Inclusion of Human Members within Human-Robot Teams. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, pages 309–317, New York, NY, USA, 3 2020. ACM. ISBN 9781450367462.

[41] Hamish Tennent, Solace Shen, and Malte Jung. Micbot: a peripheral robotic object to shape conversational dynamics and team performance. In *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 133–142. IEEE, 2019.

[42] Sylvaine Tuncer, Sarah Gillet, and Iolanda Leite. Robot-mediated inclusive processes in groups of children: From gaze aversion to mutual smiling gaze. *Frontiers in Robotics and AI*, 9, 2022. ISSN 2296-9144. doi: 10.3389/frobt.2022.729146.

[43] Klaus Weber, Hannes Ritschel, Ilhan Aslan, Florian Lingenfelser, and Elisabeth André. How to shape the humor of a robot - Social behavior adaptation based on reinforcement learning. *ICMI 2018 - Proceedings*

of the 2018 International Conference on Multimodal Interaction, pages 154–162, 2018.

[44] Jing Yu, Liqi Zhu, and Alan M. Leslie. Children's sharing behavior in mini-dictator games: The role of in-group favoritism and theory of mind. *Child Development*, 87 (6):1747–1757, 2016.