Imitation-Based Robot Learning for Human-Robot Social Greeting: A Process Study

Muhan Hou
Department of Computer Science
Vrije Universiteit Amsterdam
Amsterdam, the Netherlands
m.hou@vu.nl

Koen Hindriks
Department of Computer Science
Vrije Universiteit Amsterdam
Amsterdam, the Netherlands
k.v.hindriks@vu.nl

A.E. Eiben
Department of Computer Science
Vrije Universiteit Amsterdam
Amsterdam, the Netherlands
a.e.eiben@vu.nl

Kim Baraka
Department of Computer Science
Vrije Universiteit Amsterdam
Amsterdam, the Netherlands
k.baraka@vu.nl

I. INTRODUCTION

Human-human interaction often starts with a form of social greeting: handshakes, fist-bumps, bows, hugs, etc. Social greetings communicate deeply rooted social structures and roles. Studies have shown that human-human greeting benefits emotion sharing [1]–[4] and strengthens social connection between individuals [5]. Despite fundamental differences in the perception of humans versus robots [6], [7], some work has shown that robot greeting produced more positive user perceptions of robots and helped maintain social ties [8]. Therefore, we believe it is beneficial to equip robots with advanced greeting skills for more effective initiation of situated interactions.

Although humans engage in social greetings intuitively with almost simultaneous observation and action [9], it remains challenging for robots to master this embodied interaction skill. Such human-centered interactive task requires spatial and temporal coordination between agents in real-time, tackling physical limitation from constrained robot bodies, and decoding ambiguity in subjective evaluation for concrete learning objective. Some rule-based methods [10], [11] hardcoded interactive responses via empirical observation, but they can hardly be robust enough for unseen situations and universal for different social agents. Alternatively, data-driven methods [12]–[15] trained the robot with demonstration data, aiming to imitate the policy of human experts. However, these imitation-based approaches mainly focused on improving the algorithm itself and neglected the potential benefits if more careful attention is paid to data collection and subjective evaluation metrics. Specifically, our work aims to answer these two questions: *1) how should demonstration data be collected under limited resources so as to maximize learning performance?*, and *2) how to elicit objective factors that can help interpret subjective evaluations of a learned policy and inform algorithm improvement?*

These are indeed non-trivial questions. Regarding the *data collection*, the performance of imitation-based methods relies heavily on the amount and quality of training data [16], [17]. Unfortunately, conditions often only afford limited amounts of data (e.g., due to scarcity of resources in lab-based HRI research). How to collect diverse yet representative demonstration data within dataset size constraints remains an open question. Regarding *evaluation metrics*, it is a common practice to evaluate robot performance using subjective metrics [18], [19] in human-centered tasks like social greetings. However, just knowing how satisfied the human partner feels hardly speaks for how the robot could concretely improve its policy. It would be more improvement-informative if more robot-centered factors that are objective and tuneable could be identified from general subjective metrics. These factors would better support refining learning algorithms to adapt to user experience and predicting the subjective experience for early intervention. Despite all these benefits, how to identify objective and robot-centered factors from subjective metrics has not yet been fully investigated.

Therefore, our approach (summarized in figure 1) contributes a process study to, first, inform the collection of demonstration data and, second, identify robot-centered factors for imitation-based robot learning in social greeting tasks. Our contributions can be summarized as follows:

- ***A guideline to collect diverse and representative demonstration data within amount limitation***. We identified representative behavioral dimensions of human greeting behaviors (referred to as Human Behavioral Features or HBFs) using a Wizard-of-Oz setup and a filter-based unsupervised feature selection method. The outcome is a semi-structured procedure to collect expert demonstrations in instances that are considered diverse along the previously validated dimensions.
- ***An experimental method to elicit objective robot-centric***

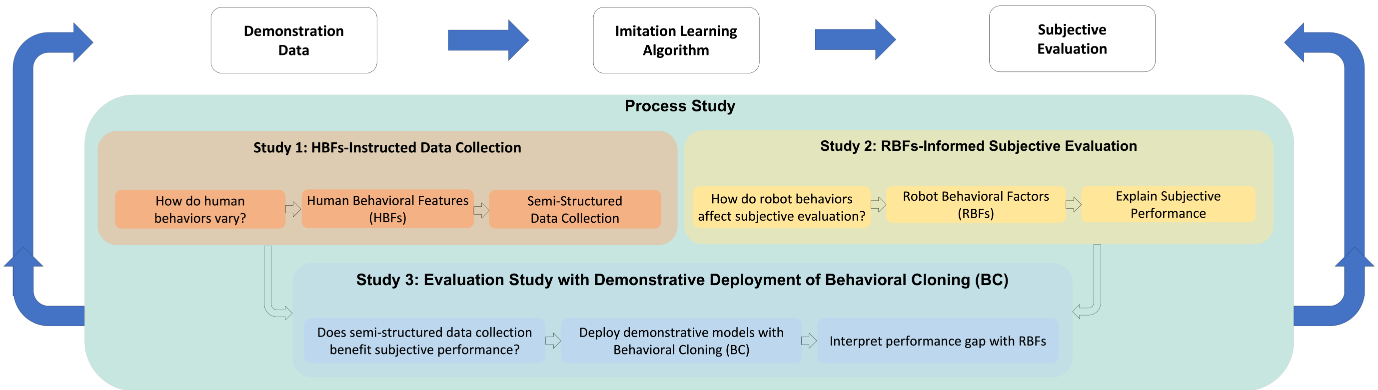

Fig. 1. Overview of the approach taken in this paper.

factors from subjective evaluation to better inform improvement. We proposed and validated robot behavioral factors (RBFs), i.e., tunable robot control variables that influence high-level subjective evaluation metrics. These factors are obtained via a Wizard-of-Oz method and are helpful to explain subjective performance gap and inform algorithm improvement.

- ***A demonstrative deployment of the aforementioned process study in a real human-robot social greeting scenario.*** With demonstration data collected following our proposed guideline, we demonstrated its application in a fist-bump greeting task trained with a typical imitation learning algorithm (i.e., Behavioral Cloning). Performance analysis followed to indicate potential benefits of RBFs in explaining subjective performance results and generating insights to potentially refine algorithm design for better performance.

II. Methodology

A. Study 1: HBF-Informed Data Collection

1) Unstructured Data Collection: To identify the representative HBFs, we collected diverse sample data on human behavior with unstructured data-collection experiments. Each human subject conducted a fist-bump with the robot in N_{div} ($N_{div} > 1$) different ways. In each trial, human subjects played the role of "initiator" to initialize the fist-bump and the robot responded to the greeting teleoperated by the same expert human operator. Human subjects first implemented the fist-bump with the robot in their most intuitive way. Then they altered their ways of doing it in any aspect as long as they still felt natural and comfortable. We refer to this data collection process as "unstructured", since we did not instruct humans on how to change their greeting behaviors and they took full charge of it.

2) Semi-Structured Data Collection: We applied Feature Selection techniques (i.e., Laplacian Score) to the data collected above and selected the top N_{rep} ($N_{rep} \geq 1$) from candidate HBFs as the representative HBFs to instruct the semi-structured data collection. Like in II-A1, human subjects first conducted fist-bumps with the robot in their most intuitive ways. Then we instructed them to vary their behaviors with respect to each of the representative HBFs. We refer to this way of data collection as "semi-structured", since we did not instruct humans how exactly they should alter their greeting behavior to reproduce certain expected distribution (i.e., "fully structured"). Alternatively, we only indicated the dimensions along which they changed their greeting behaviors (i.e., "semi-structured").

B. Study 2: RBF-Informed Subjective Evaluation

We conducted N_{RBFs} sections of experiments to identify the RBFs that impacted human subjective evaluation of the robot performance in the social greeting task. The amount of sections corresponds to the total number of potential RBFs. Each section consisted of two trials. The same expert human operator controlled the "puppet" robot to do bump-fist with human subjects twice in the same natural way. In one trial, the "performer" robot followed the exact command from the "puppet" robot (i.e., *"normal mode"*) while in the other trial we hardcoded it to deliberately divert from the original command (i.e., *"diversion mode"*) with regards to one specific RBF. After two trials finished, we immediately asked human subjects to evaluate the robot greeting in each trial with a 5-point Likert Scale, where 1 represents "very poor" and 5 represents "excellent". A one-minute break followed for fatigue concerns.

C. Study 3: Evaluation Study with Demonstrative Deployment of Behavioral Cloning (BC)

We trained Behavior Cloning (BC) models respectively with data sets from the II-A1 and II-A2. To train with an equal amount of data, for each human subject we sampled $N_{train}(N_{train} = \min(N_{div}, N_{rep})$ trials of demonstration data from both data sets, with 80% of them for training and 20% for validation.

With the trained BC models, we invited another 16 human participants to evaluate robot performance in doing fist-bump. The experiment included 3 sections and in each section the robot was controlled by a different model (i.e. one of the trained BC models, or Wizard-of-Oz). Human subjects conducted fist-bump to the robot in their most intuitive ways and scored robot performance via a 5-point Likert Scale where 1 represents "very poor" and 5 represents "excellent". A 1-minute break followed after each section. Our results indicated that more structured data collection improved subjective performance. Furthermore, such performance was correlated to influential RBFs, which could be used to refine learning algorithms.

III. Conclusion

In this work, we presented a process study of imitation-based robot learning for social greeting task. It provided guidance for collecting higher-quality demonstration data within limited budget and bridged the gap between uninformative subjective evaluation and concrete algorithm refinement. However, we only built our studies around the task of fist-bump as one of the most typical greeting behaviors. In the future, we also plan to apply the results from this work to a broader range of cooperative task settings and further investigate the potential of imitation learning algorithms in social interactive scenarios.

References

[1] J. B. Kellett, R. H. Humphrey, and R. G. Sleeth, "Empathy and the emergence of task and relations leaders," *The Leadership Quarterly*, vol. 17, no. 2, pp. 146–162, 2006.

[2] D. Keltner and J. Haidt, "Social functions of emotions at four levels of analysis," *Cognition & Emotion*, vol. 13, no. 5, pp. 505–521, 1999.

[3] D. T. Tepper and R. F. Haase, "Verbal and nonverbal communication of facilitative conditions.," *Journal of counseling psychology*, vol. 25, no. 1, p. 35, 1978.

[4] P. Ekman, W. V. Friesen, M. O'Sullivan, and K. Scherer, "Relative importance of face, body, and speech in judgments of personality and affect.," *Journal of personality and social psychology*, vol. 38, no. 2, p. 270, 1980.

[5] M. W. Kraus, C. Huang, and D. Keltner, "Tactile communication, cooperation, and performance: an ethological study of the nba.," *Emotion*, vol. 10, no. 5, p. 745, 2010.

[6] L. D. Riek, T.-C. Rabinowitch, P. Bremner, A. G. Pipe, M. Fraser, and P. Robinson, "Cooperative gestures: Effective signaling for humanoid robots," in *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pp. 61–68, IEEE, 2010.

[7] F. Eyssel and D. Kuchenbrandt, "Social categorization of social robots: Anthropomorphism as a function of robot group membership," *British Journal of Social Psychology*, vol. 51, no. 4, pp. 724–731, 2012.

[8] A. M. Sabelli, T. Kanda, and N. Hagita, "A conversational robot in an elderly care center: an ethnographic study," in *2011 6th ACM/IEEE international conference on human-robot interaction (HRI)*, pp. 37–44, IEEE, 2011.

[9] S. de La Rosa, Y. Ferstl, and H. H. Bülthoff, "Visual adaptation dominates bimodal visual-motor action adaptation," *Scientific Reports*, vol. 6, no. 1, pp. 1–8, 2016.

[10] A. Zaraki, L. Wood, O. Novanda, B. Robins, and K. Dautenhahn, "Toward autonomous child-robot interaction: development of an interactive architecture for the humanoid kaspar robot," 2017.

[11] E. Okamura and F. Tanaka, "Design of a robot that is capable of high fiving with humans," in *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pp. 704–711, IEEE, 2017.

[12] H. B. Amor, D. Vogt, M. Ewerton, E. Berger, B. Jung, and J. Peters, "Learning responsive robot behavior by imitation," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3257–3264, IEEE, 2013.

[13] H. B. Amor, G. Neumann, S. Kamthe, O. Kroemer, and J. Peters, "Interaction primitives for human-robot cooperation tasks," in *2014 IEEE international conference on robotics and automation (ICRA)*, pp. 2831–2837, IEEE, 2014.

[14] W.-R. Ko, J. Lee, M. Jang, and J. Kim, "End-to-end learning of social behaviors for humanoid robots," in *2020 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pp. 1200–1205, IEEE, 2020.

[15] J. Bütepage, A. Ghadirzadeh, Ö. Öztimur Karadağ, M. Björkman, and D. Kragic, "Imitating by generating: Deep generative models for imitation of interactive tasks," *Frontiers in Robotics and AI*, vol. 7, p. 47, 2020.

[16] Z. Han and T. Williams, "Towards formalizing hri data collection processes," *arXiv preprint arXiv:2203.08396*, 2022.

[17] Z. Gong, P. Zhong, and W. Hu, "Diversity in machine learning," *IEEE Access*, vol. 7, pp. 64323–64350, 2019.

[18] A. Steinfeld, T. Fong, D. Kaber, M. Lewis, J. Scholtz, A. Schultz, and M. Goodrich, "Common metrics for human-robot interaction," in *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, pp. 33–40, 2006.

[19] Y. Mizuchi and T. Inamura, "Optimization of criterion for objective evaluation of hri performance that approximates subjective evaluation: a case study in robot competition," *Advanced Robotics*, vol. 34, no. 3-4, pp. 142–156, 2020.