INTERVIEW WITH DR. MALCOLM DOERING

Dr. Malcolm Doering with two of the robots he works with – Robovie (left) and TIAGo (right).

Biography

Dr. Malcolm DOERING is currently a Specially-appointed Assistant Professor at the Human-Robot Interaction Laboratory, Kyoto University, where he is mainly contributing to the JST AIP Trilateral project on Artificial Intelligence for Human-Robot Interaction (AI4HRI). His research focuses on data-driven imitation learning methods for human-robot interaction, i.e., how robots can learn interactive social behaviors automatically from data using machine learning, with minimal input from human designers or expensive manual annotation.

Malcolm is originally from Michigan, USA, where he graduated from Michigan State University in 2015 with an MS in computer science. In 2015 he began working as a research intern at Hiroshi Ishiguro Laboratories, ATR, Kyoto, Japan, where his work with the ERICA android (Fig. 1) won the Best Video Award at HRI 2017 [18]. In 2019 he graduated from Osaka University with a PhD focused on human-robot interaction and joined Prof. Takayuki Kanda's Human-Robot Interaction Lab at Kyoto University (https://www.robot.soc.i.kyoto-u.ac.jp/en/), where he supervises a small team of students. In 2022 he co-organized a workshop on Artificial Intelligence for Social Robots Interacting with Humans in the Real World (intellect4hri) at IROS [19]. He has served on the organizing committee of the HRI 2019 conference and has been a reviewer for many top tier robotics conferences and journals, including ACM/IEEE International Conference on Human-Robot Interaction (HRI), IEEE Robotics and Automation Letters (RA-L), and IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).

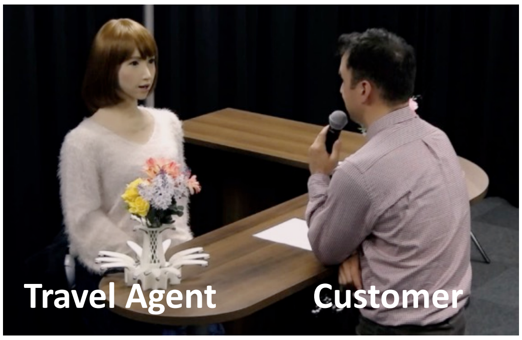

Figure 1. The ERICA android playing the role of a travel agent [12].

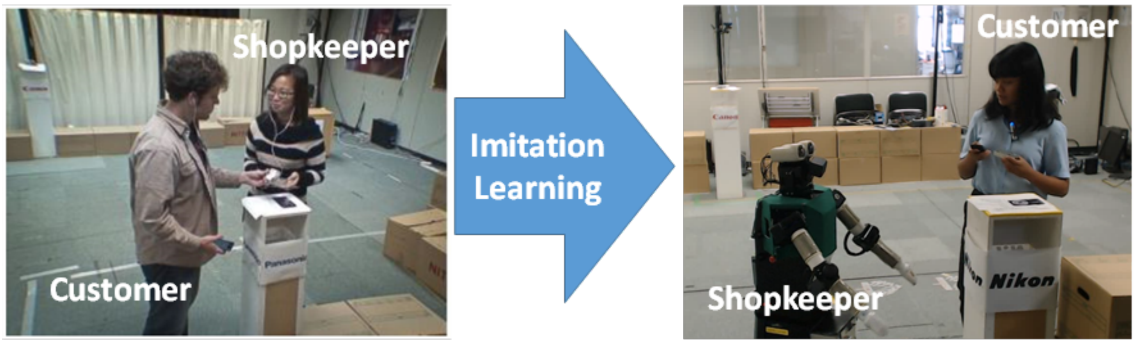

Figure 2. Natural human-human interaction data is collected with a passive sensor network and used to train a robot to replace the target human (e.g. a camera shopkeeper) [8,11].

Please briefly introduce you and your team.

Since a young age I have been inspired by science fiction works, like Isaac Asimov's I, Robot and the television program Start Trek, which depict futuristic societies with intelligent robots that are able to communicate fluently with humans. During my early studies I was inspired by these works to explore the fields of computer science and linguistics, where I was further motivated by provocative works like Alan Turing's seminal paper that asked, "Can a machine think?" and suggested that machines might learn to become intelligent, in the way that a child learns [20].

Through my research, I aim to realize socially beneficial robots, like those depicted in sci-fi, towards the goal of improving the human experience and quality of life. On another note, I have also long been interested in Japanese language and culture, such as Zen Buddhism and martial arts. That, combined with the advanced robotics research going on in Japan, led me to where I am today, working at the Human-Robot interaction Lab at Kyoto University, where I research imitation learning of social interaction behaviors and supervise a group of graduate students.

What is the current focus of you and your team's research?

The overall goal of the Human-Robot Interaction Lab is human-robot interaction broadly; but, my students and I are focusing mostly on imitation learning of social interaction behaviors for robots (Fig. 2).

Recently social robots have been applied in a variety of domains, such as elder care [1], personal companions [2], hotel concierges [3], workout partners [4], and in day-to-day interaction [5], which require various social interaction behaviors. One approach to developing interaction behaviors is to manually code them or use integrated development environments such as choreograph [6] or interaction composer [7], but this is tedious, time-consuming, and requires the developer to anticipate a myriad social scenarios. Another approach, which overcomes these limitations, is to use data-driven imitation learning to learn social interaction behaviors automatically from example interaction data without manual data annotation. With this approach, interaction data can be collected via passive sensor networks in places that people frequently interact and the repeatable, formulaic behaviors, which characterize many domains where social service robots might be useful (e.g. retail, restaurants, and museum tours), can be learned via machine learning with reduced effort by developers.

Our previous work on data-driven imitation learning has focused on one-to-one interactions [8], one-to-many interactions [9], learning proactive behaviors [10], curiosity-driven learning [11], resolving ambiguity [12], remembering customer preferences [13], and adapting to changing product inventory [14].

Currently, we are exploring the application of imitation learning to various problems, including interacting with fellow staff members as well as customers [15], interactive motions (e.g. gestures and body postures, and their relation to the environment) [16], and social interaction behaviors around object manipulation. We are also exploring how to automatically identify and prevent errors that imitation learning systems make [17].

Why do you choose to focus on data-driven imitation learning instead of reinforcement learning for human-robot interaction?

Reinforcement learning is very hot right now, but imitation learning has some advantages for human-robot interaction.

Reinforcement learning typically requires a lot of trial and error (exploration and exploitation), which is not always appropriate for a robot interacting with humans. Reinforcement learning works well when we can simulate the agent's interaction with the environment, because a large number of training sessions can be simulated very quickly. But, when it comes to social interaction, it is not so simple to design a simulation of a human for the robot to practice interacting with. If we could design such a simulation, the human-robot interaction problem would already largely be solved! Furthermore, it is usually not appropriate to perform RL for social learning in the real world either – people are probably not patient enough to deal with a robot that is frequently making mistakes, causing more trouble than it's worth.

Moreover, RL requires some reward to indicate when the system is performing well. In social interaction, it's not always clear where this reward signal should come from, and whether it can be accurately perceived by the system. For example, consider a tutor robot trying avoid frustrating a student – it would require accurate sensing of the student's facial expression and tone of voice, possibly in noisy environments with various lighting conditions.

In contrast, with imitation learning we can collect datasets of experts performing in some target scenario and use that data to directly train the robot. The robot can learn to perform the same behaviors as the expert without trial and error.

Of course, some combination of reinforcement learning and imitation learning might be the best for social robots interacting in the real world: Imitation learning can be used to learn a foundation of behaviors and interaction logic on which to build and improve with reinforcement learning during live interactions.

How do you think big models, especially large language models could benefit robot learning?

I believe large language models (LLMs) have a lot of potential for improving robot capabilities. For example, research is already going on about how knowledge embedded in LLMs can be used to teach robots to perform new tasks [21,22]. There is also a lot of knowledge about the unwritten rules of social interaction and social norms implicitly stored in these models. It will be exciting to see in the coming years how this knowledge will be used to enable robots to interact more naturally with humans.

What kinds of directions do you think are most promising in the next 5 years?

Possibly one of the most promising directions will be the application of large language models to robotics and human-robot interaction. To do this though, some problems need to be solved, such as how to ground these models to the robot's sensors and actuators, i.e. how to deal with the 'situatedness' of the robot – its embeddedness in a complex, dynamic environment. Towards solving these problems, it may be good to create large models trained on other data in addition to text. For example, training large predictive models on audio and video data may help to better ground the robot's perception of the environment to the knowledge embedded in the models.

Overall, it is very exciting to be involved in such rapidly evolving field, and I watch with hopeful anticipation of where things will go in the future.

References

[1] K. Kuwamura, S. Nishio, and H. Ishiguro, Designing Robots for Positive Communication with Senior Citizens. Springer, 2016, pp. 955–964.
[2] H. A. Samani, A. D. Cheok, F. W. Ngiap, A. Nagpal, and M. Qiu, "Towards a formulation of love in human-robot interaction," in RO-MAN. IEEE, 2010, Conference Proceedings, pp. 94–99.
[3] S. Guo, J. Lenchner, J. H. Connell, M. Dholakia, and H. Muta, "Conversational bootstrapping and other tricks of a concierge robot," in HRI, 2017, Conference Proceedings, pp. 73–81.
[4] D. J. Rea, S. Schneider, and T. Kanda, "" is this all you can do? harder!" the effects of (im) polite robot encouragement on exercise effort," in Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction, 2021, pp. 225–233.
[5] S. Rosenthal, J. Biswas, and M. Veloso, "An effective personal mobile robot agent through symbiotic human-robot interaction," in AAMAS. IFAAMAS, 2010, Conference Proceedings, pp. 915–922.
[6] E. Pot, J. Monceaux, R. Gelin, and B. Maisonnier, "Choregraphe: a graphical tool for humanoid robot programming," in RO-MAN. IEEE, 2009, Conference Proceedings, pp. 46–51.
[7] D. F. Glas, T. Kanda, and H. Ishiguro, "Human-robot interaction design using interaction composer: Eight years of lessons learned," in HRI. IEEE Press, 2016, Conference Proceedings, pp. 303–310.
[8] P. Liu, D. F. Glas, T. Kanda, and H. Ishiguro, "Data-driven hri: Learning social behaviors by example from human–human interaction," IEEE T-RO, vol. 32, no. 4, pp. 988–1008, 2016.
[9] A. Nanavati, M. Doering, D. Brščić, and T. Kanda, "Autonomously learning one-to-many social interaction logic from human-human interaction data," in Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction, ser. HRI '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 419–427. [Online]. Available: https://doi.org/10.1145/3319502.3374798
[10] P. Liu, D. F. Glas, T. Kanda, and H. Ishiguro, "Learning proactive behavior for interactive social robots," Autonomous Robots, vol. 42, no. 5, pp. 1067–1085, 2017.
[11] M. Doering, P. Liu, D. F. Glas, T. Kanda, D. Kulić, and H. Ishiguro, "Curiosity did not kill the robot: A curiosity-based learning system for a shopkeeper robot," ACM Transactions on Human-Robot Interaction (THRI), vol. 8, no. 3, p. 15, 2019.
[12] M. Doering, D. F. Glas, and H. Ishiguro, "Modeling interaction structure for robot imitation learning of human social behavior," IEEE Transactions on Human-Machine Systems, 2019.
[13] M. Doering, T. Kanda, and H. Ishiguro, "Neural-network-based memory for a social robot: Learning a memory model of human behavior from data," J. Hum.-Robot Interact., vol. 8, no. 4, Nov. 2019. [Online]. Available: https://doi.org/10.1145/3338810
[14] M. Doering, D. Brščić, and T. Kanda, "Data-driven imitation learning for a shopkeeper robot with periodically changing product information," J. Hum.-Robot Interact., vol. 10, no. 4, jul 2021. [Online]. Available: https://doi.org/10.1145/3451883
[15] M Doering, D Brščić, T Kanda. Learning Social Robot Behaviors for Interacting with Staff and Customers. Workshop on Machine Learning in Human-Robot Collaboration: Bridging the Gap, 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI 2022), March 2022.
[16] Y Jiang, M Doering, T Kanda. Towards Imitation Learning of Human Interactive Motion. Workshop on Artificial Intelligence for Social Robots Interacting with Humans in the Real World (intellect4hri), 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2022), October 2022.
[17] J Ravishankar, M Doering, T Kanda. Analysis of Robot Errors in Social Imitation Learning. Workshop on Artificial Intelligence for Social Robots Interacting with Humans in the Real World (intellect4hri), 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2022), October 2022.
[18] Dylan F. Glas, Malcolm Doering, Phoebe Liu, Takayuki Kanda, and Hiroshi Ishiguro. 2017. Robot's Delight: A Lyrical Exposition on Learning by Imitation from Human-human Interaction. In Proceedings of the Companion of the 2017 ACM/IEEE International Conference on Human-Robot Interaction (HRI '17). Association for Computing Machinery, New York, NY, USA, 408. https://doi.org/10.1145/3029798.3036646
[19] E Repiso, A Clodic, M Doering, R Alami, T Kanda, D Brščić, M Beetz, M Abdel-Keream, G Sarthou. Artificial Intelligence for Social Robots Interacting with Humans in the Real World (intellect4hri), 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2022), October 27, 2022.
[20] Turing, A. M. (1950). Can a machine think. Mind, 59(236), 433-460.
[21] J. R. Kirk, R. E. Wray, P. Lindes, and J. E. Laird, "Evaluating diverse knowledge sources for online one-shot learning of novel tasks," arXiv preprint arXiv:2208.09554, 2022.
[22] Huang, W., Abbeel, P., Pathak, D., & Mordatch, I. (2022, June). Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In International Conference on Machine Learning (pp. 9118-9147). PMLR.