

HUMAN-ROBOT INTERACTION

Robots for the people, by the people: Personalizing human-machine interaction

Caitlyn Clabaugh* and Maja Mataric

Multimodal, interactive, and multitask machine learning can be applied to personalize human-robot and human-machine interactions for the broad diversity of individuals and their unique needs.

Copyright © 2018
The Authors, some
rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim
to original U.S.
Government Works

New frontiers in human-centered robotics and artificial intelligence (AI) hold the potential to realize a human-machine symbiosis similar to that envisioned by the world's first computer programmer, Ada Lovelace, in the early 19th century. Today's robots aim to assist humans in a variety of contexts with tasks ranging from the purely functional, such as manufacturing, to the truly personal, such as in-home assistance. Advances in AI, and more specifically machine learning (ML), equip robots with increasingly powerful mechanisms for understanding and interacting with the world. More recently, these mechanisms are being developed to understand and to interact with humans in long-term, real-world settings, exposing many challenges and opportunities to learn from and for the diversity of humanity.

The field of human-robot interaction (HRI) is a multidisciplinary field that seeks to understand and to develop technologies to support robots' coexistence with humanity. Grounded in humanist values, a subset of HRI focuses on robots that improve individuals' quality of life; such research includes medical haptics, rehabilitation robotics, and socially assistive robotics (SAR). SAR specifically aspires to supplement and augment the efforts of educators, parents, caregivers, and clinicians through combined applications of engineering, computer science, and social science (1). The field has shown that human-robot partnerships can mitigate critical health challenges that require socially mediated, personalized, long-term support, such as special needs education, elderly care, and rehabilitation. Human learning, development, and care all follow nonlinear trajectories unique to each individual. Computational personalization of HRI to the hidden and continuously changing needs of individual users presents real-world complexities that are largely avoided

by the greater robotics and ML research communities.

Fortunately, in the era of big data, human-machine interactions are abundant; people regularly interact with computational models, whether hidden behind their favorite digital media websites or overtly powering interactions with AI assistants in their homes. Twenty-first century users, willingly or unknowingly, trade their personal data for "free" services. As a result, personalized robotic and virtual agent services to improve quality of life are becoming a reality in industry—commercially available instances include Woebot, the therapeutic chatbot based on cognitive behavior therapy, and Catalia Health's robotic medication-adherence coach, Mabu. However, many open challenges remain in leveraging these data to better understand and interact with human users over time, including the inherent noisiness of real-world data, the broad spectrum of individuals' needs and preferences, unintended model bias, privacy and security barriers, and the trade-offs between learning more personalized or generalizable models of HRI. To address these challenges, it is essential for the paradigms of multimodal, interactive, and multitask ML to be applied to and evaluated in real-world assistive settings.

Human interaction is inherently multimodal: While interacting, we are continually producing and interpreting a rich mixture of data. Depending on the data sources and modalities available, one may recognize a friend's joy from an audio signal such as laughter, a visual image of the friend's smile, a video showing the friend's expressive body language, or symbolic information such as a happy emoji sent by the friend. The paradigm of multimodal ML integrates features from multiple modalities to build more comprehensive models of human communication

such that they are more robust to personal differences and real-world noise (2). In research, multimodal ML has been successfully applied to affect recognition, prediction of various psychological disorders, and more classical computer vision tasks like image classification. Yet, most of today's conversational and interactive agents (e.g., AI assistants and chatbots) are unimodal, relying on voice only. Although the composition and the joint representation of multiple modalities yield many technical challenges, models that include a range of communicational features are well-situated for understanding and personalizing to human interaction in real-world environments.

The more modalities or features considered by a ML model, the more examples, computational power, and memory are needed to learn to discriminate between those features. Recent advances in deep learning accept this "curse of dimensionality" and train on massive amounts of data. For instance, one deep reinforcement learning (RL) algorithm out of Google's DeepMind, AlphaGo Zero, was trained on 29 million example games, learning through millions of trials and errors to master the strategy game Go (3). Although this is a significant instance of RL solving problems with high complexity, the complexity of Go is far from that of real-world human interaction, which is multimodal, noisy, incomplete, stochastic, contextual, and difficult to obtain data for at scale. The paradigm of interactive ML (IML) offers one solution: leverage input from human users and domain experts to further inform, guide, and expedite the ML process (4). Research in interactive RL (IRL) or policy shaping has shown that direct human feedback, even if infrequent and inconsistent, improves policy optimization (5). In addition, studies in HRI have shown that transparency, such as a robot's confession of its own incompetence or request for advice, could increase the robot's likeability and the human-robot relationship by the Pratfall Effect (6). Thus, by applying

Department of Computer Science, University of Southern California, 3710 McClintock Avenue, Los Angeles, CA 90089, USA.

*Corresponding author. Email: clabaugh@usc.edu

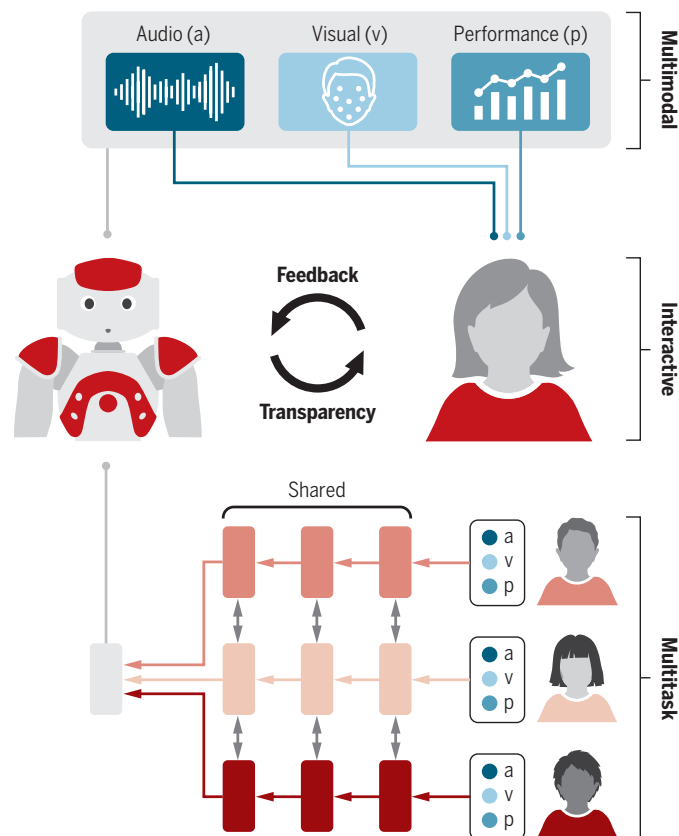


Fig. 1. Personalizing human-machine interactions. The paradigms of multimodal, interactive, and multitask ML could enable robots to personalize their interactions with human users by integrating a broad spectrum of modalities of human communication, including interactive feedback from human users and sharing knowledge across multiple, similar users.

the “human-in-the-loop” approach of IML, robots could simultaneously leverage and foster their human relationships.

Another ML paradigm that has gained traction in solving complex problems is multitask learning (MTL). Whereas many applications of ML optimize for a singular problem or task, MTL optimizes for more than one, preserving aspects of what is learned for one task and transferring it to a related task (7). MTL has been applied to a wide range of problem domains from personalized affective computing to object manipulation in robotics. MTL may be similarly applied to personalized HRI across many different users; a robot’s knowledge of how to interact with one user could

be transferred to or shared between its interactions with similar users. MTL for HRI may enable robots to more readily and intelligently interact with new users, avoiding the adverse effects of failed social interaction, even if using a data-hungry method such as RL. MTL and multimodal ML have been applied to automatically recognize engagement in children with autism spectrum disorders (ASD) (8), and we are currently studying its application in conjunction with RL to personalize long-term, in-home SAR for children with ASD (9). ASD is a developmental disorder defined by a wide variety of symptom combinations and severities related to social and communication skills; thus, inductive transfer of HRI strategies across children must

be particularly sensitive. Going beyond classical MTL applications, it is important to carefully define the relatedness or transferability between users to avoid detrimental inductive or model bias toward personalizing HRI.

As advances in robotics and AI affect increasingly more personal environments and sensitive human needs, it is paramount to consider how human interactions can and should shape technology. The paradigms of multimodal, interactive, and multitask ML hold promise for addressing the open challenges in personalizing human-robot and human-machine interactions through the inclusion of a broad spectrum of communicative modalities, interactive user feedback, and

knowledge sharing across users. The advancement and the proper application of these paradigms require greater collaboration between the fields of ML and HRI. Only through such multidisciplinary efforts can human-propelled and human-centered personalization democratize the design and development of emerging technologies for all people.

REFERENCES AND NOTES

1. M. J. Mataric, Socially assistive robotics: Human augmentation vs. automation. *Sci. Robot.* **2**, eaam5410 (2017).
2. T. Baltrusaitis, C. Ahuja, L.-P. Morency, Multimodal machine learning: A survey and taxonomy. *IEEE Trans. Pattern Analysis Machine Intelligence*, 10.1109/TPAMI.2018.2798607, published online 25 January 2018.
3. D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, L. Sifre, G. van den Driessche, T. Graepel, D. Hassabis, Mastering the game of Go without human knowledge. *Nature* **550**, 354–359 (2017).
4. S. Amershi, M. Cakmak, W. Bradley Knox, T. Kulesza, Power to the people: The role of humans in interactive machine learning. *AI Mag.* **35**, 105–120 (2014).
5. S. Griffith, K. Subramanian, J. Scholz, C. L. Isbell, A. L. Thomaz, Policy shaping: Integrating human feedback with reinforcement learning, in *Advances in Neural Information Processing Systems* (Curran Associates Inc., 2013), pp. 2625–2633.
6. N. Mirnig, G. Stollnberger, M. Miksch, S. Stadler, M. Giuliani, M. Tscheligi, To err is robot: How humans assess and act toward an erroneous social robot. *Front. Robot. AI* **4**, 21 (2017).
7. R. Caruana, Multitask learning, in *Learning to Learn* (Springer, 1998), pp. 95–133.
8. O. Rudovic, J. Lee, L. Mascarell-Maricic, B. W. Schuller, R. W. Picard, Measuring engagement in robot-assisted autism therapy: A cross-cultural study. *Front. Robot. AI* **4**, 36 (2017).
9. C. Clabaugh, D. Becerra, E. Deng, G. Ragusa, M. Mataric, Month-long, in-home case study of a socially assistive robot for children with autism spectrum disorder, in *Companion of the 2018 ACM/IEEE International Conference on Human-Robot Interaction (ACM/IEEE, 2018)*, pp. 87–88.

Competing interests: M.M. is co-founder and chief science officer of Embodied Inc., and C.C. is an employee of Embodied Inc.

10.1126/scirobotics.aat7451

Citation: C. Clabaugh, M. Mataric, Robots for the people, by the people: Personalizing human-machine interaction. *Sci. Robot.* **3**, eaat7451 (2018).

Robots for the people, by the people: Personalizing human-machine interaction

Caitlyn Clabaugh and Maja Mataric

Sci. Robotics **3**, eaat7451.
DOI: 10.1126/scirobotics.aat7451

ARTICLE TOOLS

<http://robotics.sciencemag.org/content/3/21/eaat7451>

REFERENCES

This article cites 5 articles, 0 of which you can access for free
<http://robotics.sciencemag.org/content/3/21/eaat7451#BIBL>

PERMISSIONS

<http://www.sciencemag.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of Service](#)

Science Robotics (ISSN 2470-9476) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. 2017 © The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. The title *Science Robotics* is a registered trademark of AAAS.