Lifelong Learning and Personalization in Long-Term Human-Robot Interaction (LEAP-HRI): Adaptivity for All

Bahar Irfan
birfan@kth.se
KTH Royal Institute of Technology, Sweden

Aditi Ramachandran
aditi@myvanrobot.com
Van Robotics, USA

Mariacarla Staffa
mariacarla.staffa@uniparthenope.it
University of Naples Parthenope, Italy

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

ABSTRACT

Adaptation and *personalization* are critical elements when modeling robot behaviors toward users in real-world settings. Multiple aspects of the user need to be taken into consideration in order to personalize the interaction, such as their personality, emotional state, intentions, and actions. While this information can be obtained a priori through self-assessment questionnaires or in real-time during the interaction through user profiling, behaviors and preferences can evolve in long-term interactions. Thus, gradually learning new concepts or skills (i.e., "lifelong learning") both for the users and the environment is crucial to adapt to new situations and personalize interactions with the aim of maintaining their interest and engagement. In addition, adapting to individual differences autonomously through lifelong learning allows for inclusive interactions with all users with varying capabilities and backgrounds. The third edition[1] of the "Lifelong Learning and Personalization in Long-Term Human-Robot Interaction (LEAP-HRI)" workshop aims to gather and present interdisciplinary insights from a variety of fields, such as education, rehabilitation, elderly care, service and companion robots, for lifelong robot learning and adaptation to users, context, environment, and activities in long-term interactions. The workshop aims to promote a common ground among the relevant scientific communities through invited talks and in-depth discussions via paper presentations, break-out groups, and a scientific debate. In line with the HRI 2023 conference theme, "HRI for all", our workshop theme is "adaptivity for all" to encourage HRI theories, methods, designs, and studies for lifelong learning, personalization, and adaptation that aims to promote inclusion and diversity in HRI.

CCS CONCEPTS

• **Computer systems organization** → **Robotics**; • **Information systems** → **Personalization**; • **Computing methodologies** → **Lifelong machine learning**.

KEYWORDS

Lifelong Learning; Continual Learning; Personalization; Adaptation; Long-Term Interaction; Human-Robot Interaction; Long-Term Memory; User Modeling; Diversity; Inclusivity; Workshop

ACM Reference Format:
Bahar Irfan, Aditi Ramachandran, Mariacarla Staffa, and Hatice Gunes. 2023. Lifelong Learning and Personalization in Long-Term Human-Robot Interaction (LEAP-HRI): Adaptivity for All. In *Companion of the 2023 ACM/IEEE International Conference on Human-Robot Interaction (HRI '23 Companion), March 13–16, 2023, Stockholm, Sweden.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3568294.3579956

1 INTRODUCTION

While deep learning is rapidly advancing in fields that are valuable for robotics, such as natural language processing[2,3,4,5] and computer vision[6,7,8], these algorithms typically lack the ability to learn new information and adapt after deployment, known as "lifelong (or continual) learning" [15]. Lifelong learning techniques can help robots learn new tasks through transfer learning, as well as function well in diverse environments, such as a new kitchen. In addition, individuals may differ in their needs, preferences and behaviors, which can also change over time. Hence, adaptivity of robot behaviors for all users and personalization are necessary for usability, inclusivity, effectiveness in performance, and maintenance of user engagement over long-term interactions.

The third edition of the "Lifelong Learning and Personalization in Long-Term Human-Robot Interaction (LEAP-HRI)" Workshop aims to bring together interdisciplinary researchers for discussing the advances and challenges of the field, focusing on "adaptivity for all". The workshop consists of invited keynote speakers, presentations of early-stage researchers, a breakout session and a debate for thought-provoking discussions.

2 STATE OF THE ART & NEW DIRECTIONS

If we aim to achieve *adaptivity for all* via machine learning (ML), lifelong or continual learning (CL) is the way forward because it has the potential to enable ML models to continually learn and adapt by balancing past knowledge preservation with incremental learning

[1] All editions of the workshop are available at: https://leap-hri.github.io

[2] ChatGPT: https://openai.com/blog/chatgpt/
[3] PaLM-SayCan: https://sites.research.google/palm-saycan
[4] BLOOM: https://bigscience.huggingface.co/blog/bloom
[5] Whisper: https://openai.com/blog/whisper/
[6] DALL-E 2: https://openai.com/dall-e-2/
[7] Stable Diffusion: https://stability.ai/blog/stable-diffusion-public-release
[8] Imagen: https://imagen.research.google/

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
HRI '23 Companion, March 13–16, 2023, Stockholm, Sweden
© 2023 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-9970-8/23/03.
https://doi.org/10.1145/3568294.3579956

of new information [7, 16]. Introductory background to the topics of continual and lifelong learning have already been provided in our previous workshop editions and summary papers [11, 12]. In this year's edition we would like to take a forward-facing approach and focus specifically on the emerging trends in CL, namely causality and resource-efficiency.

The majority of ML/CL models focus on finding correlations in the data and predicting outcomes rather than understanding causality [17, 23]. Causality can be used for gaining insights into the causal structure of the data and the ML problem at hand (e.g., [2]) and can be used to fortify CL models by re-adjusting model learning to cope with changing data distributions [18]. One promising direction is the causal replay for knowledge rehearsal in CL [3].

Using ML models in resource-constrained environments (e.g., robots) requires additional considerations [8, 19] such as how to minimize the high resource and energy consumption of deep learning frameworks [26] while reducing the carbon footprint associated with training and deploying these models [24]. CL frameworks also face these criticisms. Regularization and replay are important components of incremental and sequential learning [16], however, their advantages come at the cost of additional computation and memory. Various researchers have started proposing solutions to address these challenges (e.g., [25]).

In addition to the adaptivity of robot behaviors, personalization is also necessary to increase inclusivity, effectiveness in performance, and maintenance of user engagement over long-term interactions in a vision of universal inclusion paving the way to an "HRI for all" [13]. To allow adaptivity for all, it is of paramount importance to tailor the parameters of the interaction to accommodate individual differences among users, such as needs, age, gender, abilities, literacy, culture, and emotions, from physical to cognitive and sociocultural levels. This ambitious objective poses many challenges due to the complexity unpredictability and uniqueness of each user's progress during the interaction, which are also affected by possible different contexts. Furthermore, individual differences are difficult to define and to generalize from one context to another, especially when dealing with vulnerable individuals such as elderly adults or children with neurodivergent emotional, cognitive, and social responses.

While the benefits of personalization are well documented [1, 5, 6, 9, 10, 14, 22], it is not yet affordable at all. Personalization is often achieved through human-in-the-loop or Wizard of Oz (WoZ) frameworks, wherein intervention strategies are mapped to individuals a priori or in situ via human input. However, each individual has her/his unique own path of development and growth that may not be linear and thus exceptionally complicated [27]. Therefore, robot behavior adaptation needs to be tailored towards the specific needs and preferences of each participant over time by intervention strategies which must be adapted in situ via untrained human feedback [4, 20]. Reinforcement learning (RL) methods have been successfully applied to adapt to a user's learning habits over time, particularly in early child development studies [21].

3 WORKSHOP OVERVIEW

LEAP-HRI is a half-day workshop on the topics of lifelong learning and personalization in long-term HRI. The workshop will be hybrid to accommodate the needs of the participants that might not be able to attend the in-person conference. The workshop consists of: **Keynotes:** The invited researchers, Dorsa Sadigh (Stanford University, USA) and Pablo Barros (Sony, Belgium), will present their experiences and perspectives on the topics related to the workshop. The presentations will last 30 minutes, followed by 10 minutes of question session. **Debate:** Inspired by the Robotics Debates[9] at the IEEE International Conference on Robotics and Automation (ICRA), the 1-hour debate will have a controversial statement related to the topics of the workshop, where two debaters will argue their views for the statement and two debaters against it. The debaters will be Brian Scassellati (Yale University, USA), Séverin Lemaignan (PAL Robotics, Spain), Eric Eaton (University of Pennsylvania, USA), and Karol Hausman (Google Brain, USA and Stanford University, USA). **Full talks:** The authors of the accepted research papers of full length (3-4 pages) will give 7 minutes presentation followed by 3 minutes question session. **Break-out session:** To encourage constructive discussions, the workshop attendees will be divided into groups of 4-5 for a break-out session (15 minutes). **Discussion:** To improve the LEAP-HRI workshop series, there will be a discussion with the audience coverage, topics, speakers, and other issues in the previous workshops (10 minutes).

Target Audience & Attracting Participants. We invite papers of 3-4 pages (plus additional pages for references and appendices), including work in progress containing preliminary results, technical reports, case studies, surveys and state-of-the-art research in lifelong learning and personalization in a variety of fields (e.g., education, rehabilitation, elderly care, collaborative tasks, customer-oriented service and companion robots) and long-term studies. We will encourage submissions to address the theme of the workshop "Adaptivity for All". Papers will be reviewed for their relevance, novelty, and scientific and technical soundness. Researchers from HRI, robotics, cognitive science, rehabilitation and educational backgrounds will be invited to contribute. Based on the accepted submissions for LEAP-HRI in the past two years, we expect around 5-6 accepted full papers. The workshop achieved to attract 70 to 100 attendees in the previous years, we expect to reach to the same audience with the hybrid format. The workshop will be announced through a dedicated website, a call for papers on robotics mailing lists and on social network channels.

Documenting the Workshop. The accepted papers will be published on the workshop website, as well as in arXiv. The proceedings may be made as a single submission, or as a set of individual papers with an index submission.

List of Topics. Topics of interest include but are not limited to:

- Lifelong personalization and/or adaptation
- Lifelong learning or personalization for inclusivity in HRI
- Modeling user(s) and/or user behavior(s) in multi-session/ long-term HRI
- Modeling robot behavior in multi-session/ long-term HRI
- Modeling context in multi-session/ long-term HRI
- Agent/robot architectures for personalization/adaptation
- Lifelong (long-term) human-agent or multi-user/ multi-agent interactions
- Lifelong (long-term) multimodal interactions

[9] Website for Robotics Debates: https://www.roboticsdebates.org/

- Continual/lifelong machine learning
- Long-term memory (episodic, semantic, associative)
- Privacy and ethical considerations in lifelong learning/ personalization in HRI

4 ORGANIZERS

Bahar Irfan (KTH Royal Institute of Technology, Sweden) is a Postdoctoral Researcher and Digital Futures fellow at KTH Royal Institute of Technology. Her research focuses on creating personal robots that can continually learn and adapt to assist in everyday life. She has a diverse background in robotics, from personalization in long-term human-robot interaction during her PhD at the University of Plymouth and SoftBank Robotics Europe to user-centered task planning for household robotics during her MSc in computer engineering, and building robots for BSc in mechanical engineering at Boğaziçi University.

Aditi Ramachandran (Van Robotics, USA) is the Chief Technology Officer at Van Robotics where she works on building educational robots and oversees all software development at the company. She received a PhD from the Social Robotics Lab at Yale University where her research focused on personalized social robot tutors for children.

Mariacarla Staffa (University of Naples Parthenope, Italy) is an Assistant Professor in Human-Robot Interaction, Artificial Intelligence and Cognitive Robotics at the Department of Science and Technologies of the University of Naples Parthenope. She received her M.Sc. degree in Computer Science from the University Federico II with honors, in 2008. She received her Ph.D. in Computer Science and Automation Engineering from the University Federico II in 2011. She is mainly interested in exploring computational neuroscience and cognitive robotics to generate innovative strategies and solutions for scientific problems and technological limitations.

Hatice Gunes (University of Cambridge, UK) is a 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, computer vision, machine learning and social robotics. She has published over 150 papers in these areas (H-index=36, citations > 7,100), with most recent works on lifelong learning for facial expression recognition, fairness and affective robotics; and longitudinal HRI for wellbeing.

ACKNOWLEDGEMENTS: H. Gunes' work for this workshop is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) & UK Research and Innovation through the Adaptive Robotic EQ for Well-being (ARoEQ) Project (grant ref. EP/R030782/1). B. Irfan's work for this workshop is supported by the Digital Futures Research Center, Sweden.

REFERENCES

[1] Olivier A. Blanson Henkemans, Bert P.B. Bierman, Joris Janssen, Mark A. Neerincx, Rosemarijn Looije, Hanneke van der Bosch, and Jeanine A.M. van der Giessen. 2013. Using a robot to personalise health education for children with diabetes type 1: A pilot study. *Patient Education and Counseling* 92, 2 (2013).

[2] Jiaee Cheong, Sinan Kalkan, and Hatice Gunes. 2023. Causal Structure Learning of Bias for Fair Affect Recognition. In *Proc. IEEE/CVF Winter Conf. on Applications of Computer Vision (WACV) Workshops*. 340–349.

[3] Nikhil Churamani, Jiaee Cheong, Sinan Kalkan, and Hatice Gunes. 2023. Towards Causal Replay for Knowledge Rehearsal in Continual Learning. In *Proc. AAAI Bridge Program on Continual Causality*.

[4] Caitlyn Clabaugh, Kartik Mahajan, Shomik Jain, Roxanna Pakkar, David Becerra, Zhonghao Shi, Eric Deng, Rhianna Lee, Gisele Ragusa, and Maja Matarić. 2019. Long-Term Personalization of an In-Home Socially Assistive Robot for Children With Autism Spectrum Disorders. *Frontiers in Robotics and AI* 6 (2019).

[5] Caitlyn Clabaugh, Gisele Ragusa, Fei Sha, and Maja Mataric. 2015. Designing a socially assistive robot for personalized number concepts learning in preschool children. In *2015 Joint IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)*. 314–319.

[6] Alessandro Di Nuovo, Frank Broz, Ning Wang, Tony Belpaeme, Angelo Cangelosi, Ray Jones, Raffaele Esposito, Filippo Cavallo, and Paolo Dario. 2018. The multi-modal interface of Robot-Era multi-robot services tailored for the elderly. *Intelligent Service Robotics* 11 (01 2018), 1–18.

[7] Raia Hadsell, Dushyant Rao, Andrei A. Rusu, and Razvan Pascanu. 2020. Embracing Change: Continual Learning in Deep Neural Networks. *Trends in Cognitive Sciences* 24, 12 (Dec. 2020), 1028–1040.

[8] Charlie Hewitt and Hatice Gunes. 2018. CNN-based Facial Affect Analysis on Mobile Devices. *arXiv preprint arXiv:1807.08775* (2018).

[9] Bahar Irfan, Natalia Céspedes, Jonathan Casas, Emmanuel Senft, Carlos A. Cifuentes, Luisa F. Gutiérrez, Mónica Rincon-Roncancio, Tony Belpaeme, and Marcela Múnera. 2022. Personalised Socially Assistive Robot for Cardiac Rehabilitation: Critical Reflections on Long-Term Interactions in the Real World. *User Modeling and User-Adapted Interaction* (2022).

[10] Bahar Irfan, Mehdi Hellou, Alexandre Mazel, and Tony Belpaeme. 2020. Challenges of a Real-World HRI Study with Non-Native English Speakers: Can Personalisation Save the Day?. In *Companion of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 272–274.

[11] Bahar Irfan, Aditi Ramachandran, Samuel Spaulding, Sinan Kalkan, German Ignacio Parisi, and Hatice Gunes. 2021. Lifelong Learning and Personalization in Long-Term Human-Robot Interaction (LEAP-HRI). In *Companion of the 2021 ACM/IEEE International Conference on Human-Robot Interaction, HRI*. 724–727.

[12] Bahar Irfan, Aditi Ramachandran, Samuel Spaulding, German Ignacio Parisi, and Hatice Gunes. 2022. Lifelong Learning and Personalization in Long-Term Human-Robot Interaction (LEAP-HRI). In *ACM/IEEE International Conference on Human-Robot Interaction, HRI*. 1261–1264.

[13] Min Kyung Lee, Jodi Forlizzi, Sara Kiesler, Paul Rybski, John Antanitis, and Sarun Savetsila. 2012. Personalization in HRI: A Longitudinal Field Experiment. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction* (Boston, Massachusetts, USA) *(HRI '12)*. Association for Computing Machinery, New York, NY, USA, 319–326.

[14] Iolanda Leite, Ginevra Castellano, André Pereira, Carlos Martinho, and Ana Paiva. 2012. Modelling empathic behaviour in a robotic game companion for children: An ethnographic study in real-world settings. *HRI'12 - Proceedings of the 7th Annual ACM/IEEE International Conference on Human-Robot Interaction* (03 2012).

[15] Timothée Lesort, Vincenzo Lomonaco, Andrei Stoian, Davide Maltoni, David Filliat, and Natalia Díaz-Rodríguez. 2020. Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges. *Information Fusion* 58 (2020), 52–68.

[16] German I. Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. 2019. Continual lifelong learning with neural networks: A review. *Neural Networks* 113 (May 2019), 54–71.

[17] Judea Pearl. 2009. *Causality*. Cambridge university press.

[18] Judea Pearl. 2019. The seven tools of causal inference, with reflections on machine learning. *Commun. ACM* 62, 3 (2019), 54–60.

[19] Davy Preuveneers, Ilias Tsingenopoulos, and Wouter Joosen. 2020. Resource usage and performance trade-offs for machine learning models in smart environments. *Sensors* 20, 4 (2020), 1176.

[20] Laurel D. Riek. 2012. Wizard of Oz Studies in HRI: A Systematic Review and New Reporting Guidelines. *J. Hum.-Robot Interact.* 1, 1 (jul 2012), 119–136.

[21] Raquel Ros, Marco Nalin, Rachel Wood, Paul Baxter, Rosemarijn Looije, Yiannis Demiris, Tony Belpaeme, Alessio Giusti, and Clara Pozzi. 2011. Child-robot interaction in the wild: Advice to the aspiring experimenter. *ICMI'11 - Proceedings of the 2011 ACM International Conference on Multimodal Interaction*, 335–342.

[22] Silvia Rossi, Daniela Conti, Federica Garramone, Gabriella Santangelo, Mariacarla Staffa, Simone Varrasi, and Alessandro Di Nuovo. 2020. The role of personality factors and empathy in the acceptance and performance of a social robot for psychometric evaluations. *Robotics* 9, 2 (2020).

[23] Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. 2021. Toward causal representation learning. *Proc. IEEE* 109, 5 (2021), 612–634.

[24] Roy Schwartz, Jesse Dodge, Noah A Smith, and Oren Etzioni. 2020. Green AI. *Commun. ACM* 63, 12 (2020), 54–63.

[25] Samuil Stoychev, Nikhil Churamani, and Hatice Gunes. 2023. Latent Generative Replay for Resource-Efficient Continual Learning of Facial Expressions. In *Proc. IEEE Int'l. Conf. on Automatic Face and Gesture Recognition*.

[26] Samuil Stoychev and Hatice Gunes. 2022. The Effect of Model Compression on Fairness in Facial Expression Recognition. (2022). arXiv:2201.01709

[27] Lev Vygotsky. 1980 (1930). *Interaction between Learning and Development*. Vol. 2003. Harvard University Press, Cambridge, MA, Chapter 6.