Lifelong Learning and Personalization in Long-Term Human-Robot Interaction (LEAP-HRI): Open-World Learning

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

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

Andreea Bobu
abobu@theaiinstitute.com
Boston Dynamics AI Institute, USA

Nikhil Churamani
nikhil.churamani@cl.cam.ac.uk
University of Cambridge, UK

ABSTRACT

The complex and largely unstructured nature of real-world situations makes it challenging for conventional *closed-world* robot learning solutions to adapt to such interaction dynamics. These challenges become particularly pronounced in long-term interactions where robots need to go beyond their past learning to continuously evolve with changing environment settings and *personalize* towards individual user behaviors. In contrast, *open-world learning* embraces the complexity and unpredictability of the real world, enabling robots to be "lifelong learners" that continuously acquire new knowledge and navigate novel challenges, making them more context-aware while intuitively engaging the users. Adopting the theme of "open-world learning", the fourth edition of the "Lifelong Learning and Personalization in Long-Term Human-Robot Interaction (LEAP-HRI)"[1] workshop seeks to bring together interdisciplinary perspectives on real-world applications in human-robot interaction (HRI), including education, rehabilitation, elderly care, service, and companionship. The goal of the workshop is to foster collaboration and understanding across diverse scientific communities through invited keynote presentations and in-depth discussions facilitated by contributed talks, a break-out session, and a debate.

CCS CONCEPTS

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

KEYWORDS

Lifelong Learning; Continual Learning; Personalization; Adaptation; Human-Robot Interaction; Open-World Learning; Workshop

ACM Reference Format:
Bahar Irfan, Mariacarla Staffa, Andreea Bobu, and Nikhil Churamani. 2024. Lifelong Learning and Personalization in Long-Term Human-Robot Interaction (LEAP-HRI): Open-World Learning. In *Companion of the 2024 ACM/IEEE International Conference on Human-Robot Interaction (HRI '24 Companion), March 11–14, 2024, Boulder, CO, USA.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3610978.3638159

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

HRI '24 Companion, March 11–14, 2024, Boulder, CO, USA

ACM ISBN 979-8-4007-0323-2/24/03.
https://doi.org/10.1145/3610978.3638159

1 INTRODUCTION

Traditional approaches to robot learning operate under the assumption that the training dataset encompasses all the required information for the model to make accurate predictions or classifications. This assumption, referred to as the *closed-world assumption* [2, 7], implies a static and fixed state of the world. However, the real world is dynamic, introducing new classes, tasks, and environments that may emerge at any given moment. The need for the robot to autonomously identify and learn these new concepts is crucial, especially in applications requiring long-term human-robot interaction (HRI), such as homes, rehabilitation, and education. To meet the expectations of users in these environments, robots must learn to become active partners in their interactions with humans, not only perceiving and understanding their environment, including socio-emotional human behaviors, but also responding with personalized context-appropriate behaviors [6]. *Open-world learning* approaches aim to address that need without requiring retraining which could be costly in long-term applications [4]. For example, a companion robot should be capable of detecting and learning new users (e.g., new friend) and objects (e.g., new souvenir) to recognize them in future interactions without undergoing a complete relearning process of prior knowledge, such as that of 10 years, each time [2, 10]. The continuous acquisition of new knowledge and skills throughout the robot's lifespan, referred to as *lifelong (continual) learning* [12], enables the robot to offer personalized experiences to users by adjusting to their needs and preferences.

The fourth edition of the "Lifelong Learning and Personalization in Long-Term Human-Robot Interaction (LEAP-HRI)" workshop focuses on "open-world learning" to highlight the challenges of learning and personalization for long-term HRI in the real world and learn about the solutions through invited keynotes, contributed talks, a break-out session, and a debate that bring together perspectives from industry and academia.

2 BACKGROUND

Open-world learning entails that robots dynamically expand their capabilities, both in terms of improving their understanding of the environment that they perceive as well as expanding their skills by adapting to changing environmental and interaction dynamics. In terms of perception, they should be able to discern between instances of known concepts, as well as yet unknown concepts [13, 16, 17]. Learning these boundaries enables robots to detect new information and update their knowledge incrementally [4]. Lifelong (continual) learning [8, 15] approaches address such application settings where agents, throughout their lifetime, can continue to learn new concepts while ensuring past knowledge is not forgotten [5, 12]. However, most lifelong learning approaches focus on balancing past vs. novel knowledge, mitigating catastrophic interference [14]. Open-world learning, on the other hand, requires robots to identify and distinguish novel knowledge from past experiences, moving beyond knowledge preservation.

Traditional *reinforcement learning* (RL) methods struggle to accomplish long-term learning, especially in open-world settings [19] as they assume *stationarity* in data distributions where robot interactions with the environment are modeled in *episodes* of learning, as samples of the entire data distribution. This assumption is violated in open-world settings [18], especially in HRI context, as the environment, user, and robot adapt towards each other. Recent advances in *continual reinforcement learning* (CRL) [1, 11] offer a positive direction for continual expansion of robot skills, in terms of learning amidst dynamic shifts in environments, learning new tasks or adapting with changing reward dynamics. CRL approaches that focus on explicit knowledge retention [11] address the challenges of *stationarity*, by modulating model plasticity and retaining learned representations [3], or through experience or memory replay buffers [9]. By perceiving the social and contextual settings, and adapting to the evolving dynamics of the environment and user interactions, these approaches can contribute to the personalized learning of the robot to align with user preferences, needs, and expectations over time, and prolong user engagement.

3 WORKSHOP OVERVIEW

LEAP-HRI is a half-day workshop on the topics of lifelong learning and personalization in long-term HRI that focuses on learning and adapting to new concepts from interactions in the real world. The workshop is offered in a hybrid format to accommodate participants who cannot attend the conference in person, following the same format as the HRI conference. The workshop will consist of:
Keynotes: Sonia Chernova and Silvia Rossi will share their insights and thoughts on lifelong learning and personalization in HRI, focusing on their experience on deploying robots to the real-world.
Debate: The debate will revolve around a controversial statement on real-world applications in HRI, bringing together experts from industry and academia: Siddhartha Srinivasa, Maja Matarić, Georgia Chalvatzaki, and Hae Won Park, moderated by Tony Belpaeme.
Contributed Talks: Authors of accepted research papers (3-4 pages) will provide short presentations about their ongoing work to encourage feedback from the audience.
Break-out Session: The workshop participants will be divided into groups of 4-5 individuals for a break-out session. During this session, they will briefly introduce their work and engage in discussions pertaining to the debate statement.
Target Audience and Approach for Recruiting Participants: We invite scientific papers ranging from 3 to 4 pages, with additional space for references and appendices. Submissions can encompass various types of work, including ongoing projects with preliminary findings, technical reports, case studies, and surveys that address lifelong learning and personalization. These topics span diverse fields in real-world applications, such as education, rehabilitation, elderly care, collaborative tasks, and companion robots, as well as long-term studies. We encourage submissions that align with the overarching theme of the workshop, "open-world learning". All submitted papers will undergo a thorough review process to assess their relevance, originality, and scientific and technical robustness. The workshop has consistently attracted an audience of 70 to 100 attendees in the previous years, and we expect to maintain this level of engagement with the hybrid format, with 5-6 contributed talks. Details about the workshop are disseminated through the LEAP-HRI website, mailing lists, and social networking channels.
List of Topics: Topics of interest include but are not limited to:

- Lifelong personalization and/or adaptation
- Lifelong learning or personalization for open-world learning
- Incremental and/or online learning 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
- Continual/lifelong machine learning
- Long-term memory (episodic, semantic, associative)
- Privacy and ethical considerations in lifelong learning/ personalization in HRI

Plan for documenting the workshop: The accepted papers will be published on the workshop website, as well as in arXiv.

4 ORGANIZERS

Bahar Irfan, KTH Royal Institute of Technology, Sweden, birfan@kth.se. Bahar Irfan is the founder and coordinator of the LEAP-HRI workshop series. She 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. Currently, she is working on lifelong learning in large language models for developing a personalized companion robot for older adults. Previously, she was a Research and Development Associate at Evinoks Service Equipment Industry and Commerce Inc., developing customizable software for industrial robots, virtual reality applications, and smart buffets. Prior to that, she worked as an R&D Lab Associate at Disney Research Los Angeles on emotional language adaptation in multiparty interactions. She has a diverse background in robotics, from personalization in long-term HRI during her PhD at the University of Plymouth and SoftBank Robotics Europe as a Marie Skłodowska-Curie Actions fellow 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.

Mariacarla Staffa, University of Naples Parthenope, Italy, mariacarla.staffa@uniparthenope.it. Mariacarla Staffa is an Assistant Professor in Human-Computer/Robot Interaction, Artificial Intelligence and Cognitive Robotics at the Department of Science

and Technologies of the University of Naples Parthenope, Italy. She received the M.Sc. degree in Computer Science from the University Federico II with honors, in 2008. She got a Ph.D. in Computer Science and Automation Engineering from the University Federico II in 2011. She was a visiting researcher at the "Institute de Système Intelligentes et de Robotique" at the University of Paris "Pierre et Marie Curie". She is a senior member of the Institute of Electrical and Electronics Engineers (IEEE) and of the EUCognition - European Society for Cognitive Systems (ID: 2037). She is also part of the IEEE RAS Technical Committee for Cognitive Robotics. She serves as Expert Reviewer for the European Commission Framework Programme for Research and Innovation (Area: AI and Robotics). She is the Coordinator of the virtual BRAIN Laboratory (the "Bioinspired Robotics and Artificial intelligence Networking Lab" of the University of Naples Parthenope) working in the fields of Cognitive Robotics, Artificial Intelligence and Social and Assistive Robotics. She is mainly interested in exploring computational neuroscience and cognitive robotics to generate innovative strategies and solutions for scientific problems and technological limitations. She authored several works on Social Assistive Robots, Adaptive Human Robot Interaction, Human Behavior and Emotion interpretation, etc. She is the Principal Scientific Coordinator of the Project RESTART - Robot Enhanced Social abilities based on Theory of mind for Acceptance of Robot in assistive Treatments and Unit Scientifc Coordinator of the SPECTRA Project (Supporting schizophrenia PatiEnts' Care wiTh Robotics and Artificial Intelligence) both funded by the Italian Ministry of University and Research.

Andreea Bobu, Boston Dynamics AI Institute, USA,
abobu@theaiinstitute.com. Andreea Bobu is a Research Scientist at the Boston Dynamics AI Institute, and an incoming Assistant Professor at MIT. She works at the intersection of robotics, mathematical human modeling, and deep learning, and studies algorithmic human-robot interaction with a focus on how robots and humans can efficiently arrive at shared representations of their tasks for more seamless and reliable interaction. She obtained her Ph.D. in Electrical Engineering and Computer Science at University of California Berkeley with Anca Dragan. Previously, she received a B.S. in Computer Science and Engineering at MIT. She was the recipient of the Apple AI/ML Ph.D. fellowship, is a Rising Star in EECS and an R:SS and HRI Pioneer, and has won best paper award at HRI 2020.

Nikhil Churamani, University of Cambridge, UK,
nikhil.churamani@cl.cam.ac.uk. Nikhil Churamani is a Postdoctoral Researcher at the Affective Intelligence and Robotics (AFAR) Lab of the Department of Computer Science and Technology, University of Cambridge. His PhD research at the University of Cambridge focused on Continual Learning for Affective Robotics, funded by EPSRC, UKRI. His current research investigates Continual Lifelong Learning of Affect for social robots, focused on affect-driven learning for Human-Robot Interaction as well as Federated Continual Learning of socially appropriate robot behaviors in human-centered environments. He has published in several top journals and conferences such as PMLR, IEEE Transactions of Affective Computing, Frontiers in Robotics & AI, ACM/IEEE HRI, IEEE FG, IEEE RO-MAN, IEEE IROS amongst others.

ACKNOWLEDGMENTS

B. Irfan's work for this workshop is supported by the KTH Digital Futures Research Center, Sweden. M. Staffa's work for this workshop is supported by the PRIN2022 research project "RESTART" (Robot Enhanced Social abilities based on Theory of mind for Acceptance of Robot in assistive Treatments), funded by the Next Generation EU (NGEU) Programme and by the Italian Ministry of University and Research (MIUR). N. Churamani's work is supported by Google under the GIG Funding Scheme.

REFERENCES

[1] David Abel, Andre Barreto, Benjamin Van Roy, Doina Precup, Hado van Hasselt, and Satinder Singh. 2023. A Definition of Continual Reinforcement Learning. In *Thirty-seventh Conference on Neural Information Processing Systems*.
[2] Abhijit Bendale and Terrance Boult. 2015. Towards Open World Recognition. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE.
[3] Diana Borsa, Thore Graepel, and John Shawe-Taylor. 2016. Learning shared representations in multi-task reinforcement learning. *arXiv preprint:1603.02041* (2016).
[4] Zhiyuan Chen and Bing Liu. 2018. *Lifelong Machine Learning*. Springer International Publishing, Cham.
[5] Nikhil Churamani, Sinan Kalkan, and Hatice Gunes. 2020. Continual Learning for Affective Robotics: Why, What and How?. In *29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*. IEEE, 425–431.
[6] Kerstin Dautenhahn. 2004. Robots we like to live with?! - A developmental perspective on a personalized, life-long robot companion. In *13th IEEE International Workshop on Robot and Human Interactive Communication*. IEEE, 17–22.
[7] Geli Fei, Shuai Wang, and Bing Liu. 2016. Learning Cumulatively to Become More Knowledgeable. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (San Francisco, California, USA) *(KDD '16)*. Association for Computing Machinery, New York, NY, USA, 1565–1574.
[8] 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.
[9] Tyler L Hayes, Nathan D Cahill, and Christopher Kanan. 2019. Memory efficient experience replay for streaming learning. In *International Conference on Robotics and Automation (ICRA)*. IEEE, 9769–9776.
[10] Bahar Irfan, Michael Garcia Ortiz, Natalia Lyubova, and Tony Belpaeme. 2021. Multi-Modal Open World User Identification. *Transactions on Human-Robot Interaction* 11, 1 (2021).
[11] Khimya Khetarpal, Matthew Riemer, Irina Rish, and Doina Precup. 2022. Towards Continual Reinforcement Learning: A Review and Perspectives. *Journal of Artificial Intelligence Research* 75 (2022), 1401–1476.
[12] 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.
[13] Martin Mundt, Yongwon Hong, Iuliia Pliushch, and Visvanathan Ramesh. 2023. A wholistic view of continual learning with deep neural networks: Forgotten lessons and the bridge to active and open world learning. *Neural Networks* 160 (2023), 306–336.
[14] Martin Mundt, Iuliia Pliushch, Sagnik Majumder, Yongwon Hong, and Visvanathan Ramesh. 2022. Unified Probabilistic Deep Continual Learning through Generative Replay and Open Set Recognition. *Journal of Imaging* 8, 4 (2022).
[15] 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.
[16] Jitendra Parmar, Satyendra Chouhan, Vaskar Raychoudhury, and Santosh Rathore. 2023. Open-World Machine Learning: Applications, Challenges, and Opportunities. *ACM Comput. Surv.* 55, 10, Article 205 (feb 2023), 37 pages.
[17] Walter J. Scheirer, Anderson de Rezende Rocha, Archana Sapkota, and Terrance E. Boult. 2013. Toward Open Set Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 7 (2013), 1757–1772.
[18] Annie Xie, James Harrison, and Chelsea Finn. 2021. Deep Reinforcement Learning amidst Continual Structured Non-Stationarity. In *Proceedings of the 38th International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 11393–11403.
[19] Haoqi Yuan, Chi Zhang, Hongcheng Wang, Feiyang Xie, Penglin Cai, Hao Dong, and Zongqing Lu. 2023. Plan4MC: Skill Reinforcement Learning and Planning for Open-World Minecraft Tasks. *arXiv preprint:2303.16563* (2023).