

The 1st International Workshop on Human-in-the-Loop Applied Machine Learning (HITLAML)

BEREKET A. YILMA, University of Luxembourg, Luxembourg

SURAFEL M. LAKEW, AWS AI Labs, USA

YOGESH VIRKAR, AWS AI Labs, USA

ALINA KARAKANTA, Leiden University, Netherlands

LAURA PLEIN, Luxembourg Tech School asbl, Luxembourg

CCS Concepts: • **Computing methodologies** → **Machine learning algorithms; Machine learning approaches; Learning paradigms; Artificial intelligence**; • **Human-centered computing** → **HCI design and evaluation methods; Interactive systems and tools**.

Additional Key Words and Phrases: Interactive Machine Learning; Human-centered Artificial Intelligence; Adaptation; Personalization

ACM Reference Format:

Bereket A. Yilma, Surafel M. Lakew, Yogesh Virkar, Alina Karakanta, and Laura Plein. 2023. The 1st International Workshop on Human-in-the-Loop Applied Machine Learning (HITLAML). In . ACM, New York, NY, USA, 7 pages.

1 BACKGROUND AND MOTIVATION

Recent advances in applied Machine Learning (ML) are increasingly involving humans. Humans have turned out to be a crucial component in all stages of the development of robust technologies, from data processing, model training, inference, and system design to practical applications like disaster management robotics, especially in situations where physical limitations exist, or tasks are challenging to execute remotely. The involvement of humans ranges from data annotation [3, 26, 34], interpretability [13, 15, 19], bias mitigation [5, 7], adaptability [16, 17, 27] and learning from feedback [4]. For example, human feedback is today considered indispensable for leading large language models to follow the user’s intent [22]. This enables the model to achieve a more conceptual and symbolic understanding of objects to acquire the underlying reasoning behind the learned tasks. ML models incorporating human feedback and knowledge are deployed for designing conversational agents (CA) [14], adaptive user interfaces (AUIs)[6], human-centered artificial intelligence (HCAI) [2, 21, 24, 28], diverse human-computer interaction (HCI) and human-robot interaction applications. In the context of human-robot interaction, the integration of human feedback has become an essential practice for enhancing robotics capabilities. Acknowledging this approach as an effective means for training and leveraging human expertise to elevate skills related to grasping, locomotion, conversation, and human-like perception is crucial. A range of techniques, including a mix of self-supervised learning and learning through demonstrations, has been utilized to train robots in these aspects [4].

ML research, particularly computer vision (CV) and natural language processing (NLP) have enjoyed enormous success over the past decade. Advances in NLP have shown great relevance for various downstream tasks such as

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).

© 2023 Copyright held by the owner/author(s).

Manuscript submitted to ACM

language generation [1], personalization [32], and recommender systems [29–31]. Similarly, autonomous vehicles [9, 11, 12, 18, 23], robotics [8, 10], medical imaging [3], facial recognition, gesture tracking [20, 33] and interactive entertainment are among the key areas where the cross-domain adoption of human-in-the-loop techniques has gained momentum.

Improving model predictions, and creating a seamless interaction between humans and ML systems are the two main reasons for Human-in-the-Loop Applied ML (HITLAML). The aim of this interdisciplinary workshop is to put a spotlight on recent advances and practical applications of ML involving human feedback and identify new paths for refining ML algorithms, promoting transparency, improving model interpretability, and ensuring ethically sound AI.

1.1 Theme of Submission

The workshop has an interdisciplinary theme and welcomes submissions of works involving practical applications of Machine learning in HCI context. The relevant themes include but not limited to:

- Human-in-the-loop NLP
 - Language Modeling for downstream tasks
 - Machine Translation
 - Conversational Agents
- Human-in-the-loop CV
 - Gesture recognition
 - Motion Detection
 - Pose Estimation
- HCI
 - Emotion recognition using EEG, FNIRS, Eye tracking
 - Brain computer interfaces (BCI)
 - Bio-signal processing
 - Mouse tracking
 - Adaptive User Interfaces (AUI)
- Applied ML
 - Recommender Systems, Personalization
 - Applied ML in Health care
 - Applied ML in digital arts and cultural Heritage
 - Applied ML in Economics and Finance
 - Applied ML in Transportation and Logistics
- Ethical considerations in applied ML.

1.2 Call for Participation

Submissions for contributing papers must be original and may consist of 4–6 pages of content excluding references and appendices by following the CEUR single-column style [template](#) Previously published articles, preliminary work, and demos can be submitted as extended abstracts (2 pages + references). Submissions are to be made via [Microsoft CMT](#).

Papers will be submitted to a double-blind review process. Accepted papers and abstracts will be presented as oral or poster presentations and will be published with [CEUR-Workshop Proceedings](#). Presentations will be made freely

The 1st International Workshop on Human-in-the-Loop Applied Machine Learning (HITLAML)

available on the workshop [website](#). All papers accepted to HITLAML must be presented at the workshop to appear in the proceedings. HITLAML requires that at least one author of each accepted submission must attend the workshop.

Important Dates:

- Paper and Demo Submission deadline: July 14, 2023
- Travel Grant Application deadline: July 20, 2023
- Notification of acceptance: August 3, 2023
- Camera-ready Submission: August 18, 2023
- Registration opens: July 26, 2023
- Workshop dates: September 04 - 06, 2023

All deadlines are in Anywhere on Earth (AoE) time zone.

2 ORGANIZERS

- **Bereket A. Yilma:** is a Computer Scientist specializing in optimization and applied Machine Learning. He holds a PhD in Automatic Signal and Image Processing, and Computer Engineering. Currently, he works as a researcher within the Computational Interaction (COIN) research group at the University of Luxembourg and is a Lecturer in the Department of Computer Science. He also serves as the instructor of Recommender Systems in the Doctoral School of Computer Science and Computer Engineering. His research focuses on various aspects of Human-Centered Artificial Intelligence, including Recommender Systems, Adaptive User Interfaces, and Personalization in the context of Smart Interactive Environments aka Cyber-Physical-Social Systems (CPSS). As an Associate Chair of SIGCHI, he actively contributes to the HCI community. He is an Instructor at the SIGCHI CIX Summer Schools and serves as a PC member and reviewer for flagship HCI/ML venues.
- **Surafel M. Lakew:** is an applied scientist at Amazon AWS AI Labs, since 2020. He completed his Ph.D. at the Fondazione Bruno Kessler HLT-MT research lab and at the University of Trento ICT-Doctoral School, advised by Marcello Federico. His thesis, Multilingual Neural Machine Translation for Low-Resource Languages, investigates approaches to improve the performance of neural machine translation models for language pairs without large-scale training data. His current research focuses on natural language and speech processing.
- **Yogesh Virkar:** is an applied scientist at AWS AI Labs working on speech and language technologies. He graduated with a Ph.D. from the University of Colorado, Boulder in 2017. In his thesis titled "Effects of multilayer network interactions on neural network dynamics", he explored stability and learning in computational models of complex biological neural networks. Before joining Amazon, he worked for Agile SDE, LLC., a small technology startup in the healthcare industry where he worked on problems related to natural language processing of medical text.
- **Alina Karakanta:** is an Assistant Professor in Machine Translation at Leiden University. Her research has been navigating disciplinary boundaries between translation studies, computational linguistics and machine learning, with the aim of leveraging the symbiotic relationship between humans and artificial intelligence to enhance machine translation systems. Her research interests include speech translation, NLP for low-resource languages

and scenarios, corpus-based translation & interpreting studies, and media accessibility.

- **Laura Plein:** is a researcher and tech educator at the Luxembourg Tech School asbl. Currently, her research focuses on machine learning and its applications in cybersecurity. More specifically, she specializes in the development of AI verification systems for use in privacy-related use cases, as well as the ethical implications of such technologies. In addition to this current research, she has previous experience working with natural language processing and has developed Luxembourgish language models. Beyond her work on Luxembourgish models, she is also working on generating formal test cases from informal bug reports using NLP towards improving the process of automated program repair in collaboration with the SnT.

3 WEBSITE

Details about the workshop, including key dates, will be available via the workshop website: <https://hitlaml.uni.lu>. This website together with the workshop's official [Twitter](#) and [LinkedIn](#) pages will act as the information catalyst in the weeks leading up to the workshop. It would also act as a repository for participants' submissions, and participants will be encouraged to familiarize themselves with the other attendees before the day of the workshop.

4 WORKSHOP PROGRAM

The 1st edition of HITLAML workshop¹ took place at the University of Luxembourg, Belval Campus from september 04 - 06, 2023. The event garnered a lot of interest from the community with over 115 participants, 11 submissions, 2 tutorials, 3 keynotes, 1 demo session and 1 panel discussion that fostered a lot of interdisciplinary communication. Additionally, we also published proceedings with CEUR². The event took place as a hybrid format to encourage collaboration, knowledge sharing, and meaningful interactions among participants, whether they choose to attend in-person or virtually. Remote attendees had access to a dedicated virtual platform, where they can join live-streamed sessions, interact with speakers, and engage in discussions. Recognizing the diverse time zones and scheduling constraints of our global audience, the workshop offered asynchronous engagement. All sessions, including keynote talks, paper presentations, and panel discussions, were recorded and made available for later viewing via the workshop [website](#) and our official [YouTube](#) channel. This allowed participants who could not attend in real-time to access the workshop materials at their convenience, ensuring that the knowledge shared during the workshop remains accessible to a broad and varied audience. We also encouraged and made sure all authors prepare their papers and presentations following the accessible presentation guidelines of SIGACCESS[25].

5 WORKSHOP STRUCTURE

During the 3-day event, keynotes, tutorials, paper presentations, interactive demos and a panel of discussions were delivered by world-renowned experts including Prof. Jean Vanderdonckt from UCLouvain, Belgium, Prof. Jonathan K. Kummerfeld from the University of Sydney, Australia, Prof. Janin Koch from ExSitu Lab at Inria Paris-Saclay, Dr. Ludovik Coba from Expedia Group, London, United Kingdom. The workshop activities and timeline are shown in [Table 1](#).

The workshop created significant opportunities for scientists to share their latest research findings, allowing the general public to learn about the practical applications of AI. The panel discussion, a highlight of the conference, focused

¹<https://hitlaml.uni.lu/>

²<https://ceur-ws.org/Vol-3524/>

The 1st International Workshop on Human-in-the-Loop Applied Machine Learning (HITLAML)

Day 1	
Workshop Activity	Timeline
Registration and welcome coffee	09:00 - 09:30
Tutorial 1: by <i>Dr. Ludovik Coba</i> : Interpretability of Machine Learning Models	10:00 - 11:00
Paper presentations	11:00 - 12:00
Lunch break and networking	12:00 - 2:00
Opening Keynote by <i>Prof. Djamila Aouada</i> : Enabling 3D reverse engineering	2:00 - 3:00
Spotlight Presentation	3:00 - 3:45
Break and networking	3:45 - 4:00
Demo session	4:00 - 6:00
Dinner	7:00 - 8:00
Day 2	
Keynote by <i>Prof. Jean Vanderdonckt</i> : To the end of our possibilities with adaptive user interfaces	09:15 - 10:15
Break and networking	10:15 - 10:30
<i>Prof. Jonathan K. Kummerfeld</i> : Collaborative Human-AI Systems for Databases, Diplomacy, and more	10:30 - 12:00
Lunch break and networking	12:00 - 2:00
Paper presentations	2:00 - 3:00
Panel discussion	3:00 - 3:45
Break and networking	3:45 - 4:00
Closing Keynote by <i>Prof. Janin Koch</i> : Human-AI collaboration Shifting the focus to interaction	4:00 - 5:00
Closing Ceremony and Best paper award	5:00 - 5:30
Day 3	
Socials	10:00 - 11:00
A piece of world heritage: The Blast furnaces in the City of Science (Tour)	11:00 - 12:00

Table 1. Workshop Activities and Timeline.

on pressing topics at the intersection of Artificial Intelligence and Society. Expert discussions touched upon crucial subjects such as Generative AI, Bias Mitigation, Transparency and Explainability, Privacy, and Security. The active participation of attendees added depth to the conversation, emphasizing the need for coordinated efforts in every facet of AI. The workshop, provided a unique forum for decision-makers, scientists, and end-users to come together and shape the future of AI.

The workshop also invited young students from Luxembourg Tech School (LTS) to feature their projects during the interactive demo session, giving them the opportunity to receive expert feedback, network with professionals and gain insights into the latest advancements in AI research. HITLAML'23 was sponsored by FRN's [RESCOM](#) grant co-funded with the Interdisciplinary Lab for Intelligent and Adaptive Systems (ILIAS) from the UNiversity of Luxembourg.

6 POST-WORKSHOP PLANS

The goal of human-in-the-loop AI with our main focus areas of interaction, natural Language/speech processing, computer vision and responsible AI presents a wealth of interdisciplinary opportunities to researchers and practitioners alike. We aspire to build and foster a community centered around these pivotal topics, uniting diverse voices and expertise. To that end, we plan to publish a report of the workshop highlights and outcomes. Workshop papers along with outcomes will be made available after the workshop on the workshop website and published in workshop proceedings. We also aim to keep fostering communication and interest through mailing lists and blog posts on social media. Except

for the scientific community, we are planning to disseminate selected findings to the general public in plain language and accessible format, to raise awareness on ML topics that may affect everyday lives. Additionally, participants and organizers will discuss next steps at the workshop to drive continued progress, solidify connections, and sustain the community's enthusiasm beyond the workshop's conclusion.

REFERENCES

- [1] Sahar Abdelnabi, Amr Gomaa, Sarath Sivaprasad, Lea Schönerr, and Mario Fritz. 2023. LLM-Deliberation: Evaluating LLMs with Interactive Multi-Agent Negotiation Games. *arXiv preprint arXiv:2309.17234* (2023).
- [2] Joanna J. Bryson and Andreas Theodorou. 2019. *How Society Can Maintain Human-Centric Artificial Intelligence*. Springer Singapore, Singapore, 305–323.
- [3] Leonardo C da Cruz, César A Sierra-Franco, Greis Francy M Silva-Calpa, and Alberto Barbosa Raposo. 2021. Enabling autonomous medical image data annotation: A human-in-the-loop reinforcement learning approach. In *2021 16th Conference on Computer Science and Intelligence Systems (FedCSIS)*. IEEE, 271–279.
- [4] Joseph DelPreto, Jeffrey I Lipton, Lindsay Sanneman, Aidan J Fay, Christopher Fourie, Changhyun Choi, and Daniela Rus. 2020. Helping robots learn: a human-robot master-apprentice model using demonstrations via virtual reality teleoperation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 10226–10233.
- [5] Gianluca Demartini, Stefano Mizzaro, and Damiano Spina. 2020. Human-in-the-loop Artificial Intelligence for Fighting Online Misinformation: Challenges and Opportunities. *The Bulletin of the Technical Committee on Data Engineering* (2020), 65–74.
- [6] Mateusz Dubiel, Bereket Abera Yilma, Kayhan Latifzadeh, and Luis A Leiva. 2022. A Contextual Framework for Adaptive User Interfaces: Modelling the Interaction Environment. *arXiv preprint arXiv:2203.16882* (2022).
- [7] Bhavya Ghai and Klaus Mueller. 2022. D-BIAS: a causality-based human-in-the-loop system for tackling algorithmic bias. *IEEE Transactions on Visualization and Computer Graphics* 29, 1 (2022), 473–482.
- [8] Milan Gnjatović, Jovica Tasevski, Milutin Nikolić, Dragiša Mišković, Branislav Borovac, and Vlado Delić. 2012. Adaptive multimodal interaction with industrial robot. In *2012 IEEE 10th Jubilee International Symposium on Intelligent Systems and Informatics*. IEEE, 329–333.
- [9] Amr Gomaa, Alexandra Alles, Elena Meiser, Lydia Helene Rupp, Marco Molz, and Guillermo Reyes. 2022. What's on your mind? A Mental and Perceptual Load Estimation Framework towards Adaptive In-vehicle Interaction while Driving. In *Proceedings of the 14th International Conference on Automotive User Interfaces and Interactive Vehicular Applications*. 215–225.
- [10] Amr Gomaa, Bilal Mahdy, Niko Kleer, Michael Feld, Frank Kirchner, and Antonio Krüger. 2023. Teach Me How to Learn: A Perspective Review towards User-centered Neuro-symbolic Learning for Robotic Surgical Systems. *arXiv preprint arXiv:2307.03853* (2023).
- [11] Amr Gomaa, Guillermo Reyes, and Michael Feld. 2021. ML-PersRef: A Machine Learning-Based Personalized Multimodal Fusion Approach for Referencing Outside Objects From a Moving Vehicle. In *Proceedings of the 23rd International Conference on Multimodal Interaction*. ACM, New York, NY, USA, 318–327.
- [12] Lisa Hassel and Eli Hagen. 2005. Adaptation of an automotive dialogue system to users' expertise. In *Proceedings of the 6th SIGdial Workshop on Discourse and Dialogue*. 222–226.
- [13] Sungsoo Ray Hong, Jessica Hullman, and Enrico Bertini. 2020. Human factors in model interpretability: Industry practices, challenges, and needs. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1 (2020), 1–26.
- [14] Srinivasan Janarthanam and Oliver Lemon. 2014. Adaptive generation in dialogue systems using dynamic user modeling. *Computational Linguistics* 40, 4 (2014), 883–920.
- [15] Isaac Lage, Andrew Ross, Samuel J Gershman, Been Kim, and Finale Doshi-Velez. 2018. Human-in-the-loop interpretability prior. *Advances in neural information processing systems* 31 (2018).
- [16] Zhe Liu, Yufan Guo, and Jalal Mahmud. 2021. When and Why a Model Fails? A Human-in-the-loop Error Detection Framework for Sentiment Analysis. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Papers*. Association for Computational Linguistics, Online, 170–177. <https://doi.org/10.18653/v1/2021.naacl-industry.22>
- [17] Zimo Liu, Jingya Wang, Shaogang Gong, Huchuan Lu, and Dacheng Tao. 2019. Deep reinforcement active learning for human-in-the-loop person re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision*. 6122–6131.
- [18] Udara E Manawadu, Mitsuhiro Kamezaki, Masaaki Ishikawa, Takahiro Kawano, and Shigeki Sugano. 2017. A multimodal human-machine interface enabling situation-Adaptive control inputs for highly automated vehicles. In *2017 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 1195–1200.
- [19] Menaka Narayanan, Emily Chen, Jeffrey He, Been Kim, Sam Gershman, and Finale Doshi-Velez. 2018. How do humans understand explanations from machine learning systems? an evaluation of the human-interpretability of explanation. *arXiv preprint arXiv:1802.00682* (2018).
- [20] Natalia Neverova, Christian Wolf, Graham Taylor, and Florian Nebout. 2015. Moddrop: adaptive multi-modal gesture recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38, 8 (2015), 1692–1706.
- [21] Andrzej Nowak, Paul Lukowicz, and Pawel Horodecki. 2018. Assessing artificial intelligence for humanity: Will AI be the our biggest ever advance? Or the biggest threat [Opinion]. *IEEE Technology and Society Magazine* 37, 4 (2018), 26–34.

The 1st International Workshop on Human-in-the-Loop Applied Machine Learning (HITLAML)

- [22] L. Ouyang, J. Wu, X. Jiang, D. Almeida, C. Wainwright, P. Mishkin, C. Zhang, S. Agarwal, K. Slama, J. Schulman, A. Ray, J. Hilton, F. Kelton, L. Miller, M. Simens, A. Askell, P. Welinder, P. F. Christiano, J. Leike, and R. Lowe. 2022. Training language models to follow instructions with human feedback. In *Proceedings of NIPS*. 27730–27744.
- [23] Seth Rogers, C-N Fiechter, and Cynthia Thompson. 2000. Adaptive user interfaces for automotive environments. In *Proceedings of the IEEE Intelligent Vehicles Symposium 2000 (Cat. No. 00TH8511)*. IEEE, 662–667.
- [24] Ben Shneiderman. 2020. Human-centered artificial intelligence: Reliable, safe & trustworthy. *International Journal of Human–Computer Interaction* 36, 6 (2020), 495–504.
- [25] SIGACCESS. 2023. *Accessible Presentation Guide*. <https://www.sigaccess.org/welcome-to-sigaccess/resources/accessible-presentation-guide/>
- [26] Eric Wallace, Pedro Rodriguez, Shi Feng, Ikuya Yamada, and Jordan Boyd-Graber. 2019. Trick me if you can: Human-in-the-loop generation of adversarial examples for question answering. *Transactions of the Association for Computational Linguistics* 7 (2019), 387–401.
- [27] Charles Welch, Chenxi Gu, Jonathan K. Kummerfeld, Veronica Perez-Rosas, and Rada Mihalcea. 2022. Leveraging Similar Users for Personalized Language Modeling with Limited Data. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Dublin, Ireland, 1742–1752. <https://doi.org/10.18653/v1/2022.acl-long.122>
- [28] Wei Xu. 2019. Toward human-centered AI: a perspective from human-computer interaction. *interactions* 26, 4 (2019), 42–46.
- [29] Bereket Abera Yilma, Najib Aghenda, Marcelo Romero, Yannick Naudet, and Hervé Panetto. 2020. Personalised visual art recommendation by learning latent semantic representations. In *2020 15th International Workshop on Semantic and Social Media Adaptation and Personalization (SMA)*. IEEE, 1–6.
- [30] Bereket A Yilma and Luis A Leiva. 2023. The Elements of Visual Art Recommendation: Learning Latent Semantic Representations of Paintings. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–17.
- [31] Bereket A Yilma and Luis A Leiva. 2023. Together Yet Apart: Multimodal Representation Learning for Personalised Visual Art Recommendation. In *Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization*. 204–214.
- [32] Bereket Abera Yilma, Yannick Naudet, and Hervé Panetto. 2021. Personalisation in cyber-physical-social systems: A multi-stakeholder aware recommendation and guidance. In *Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization*. 251–255.
- [33] Yanxia Zhang, Sophie Stellmach, Abigail Sellen, and Andrew Blake. 2015. The costs and benefits of combining gaze and hand gestures for remote interaction. In *Human-Computer Interaction – INTERACT 2015*. Springer, 570–577.
- [34] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. 2020. Differentiable augmentation for data-efficient gan training. *Advances in neural information processing systems* 33 (2020), 7559–7570.