
ICLR 2024 Workshop on Learning from Time Series for Health

Abstract

This in-person workshop delves into the intricacies of extracting valuable insights from healthcare time series data, aiming to deepen our comprehension of human health. The session will concentrate on two important themes within this domain: Behavioral Health and Foundational Models. Throughout the workshop, we will highlight cutting-edge approaches and solutions tailored for analyzing healthcare time series data, positioning attendees at the forefront of advancements in this field. By participating in this workshop, attendees will also become part of a diverse community committed to driving forward the boundaries of knowledge in human health.

Main Proposal

<https://timeseriesforhealth.github.io/>

Introduction Time series data are pervasive throughout healthcare, arising in electronic health records, wearable devices, mobile health, population health trends, and physiological waveforms. Today, there is an exciting opportunity to advance machine learning methods for extracting actionable insights about human health from time series data [3]. However, machine learning for health presents unique challenges, as does learning from time series. Together, this intersectional research area is rapidly evolving, yet requires immense community focus. Challenges include noisy and missing labels [4], multimodal and extremely high dimensional data [6], irregularly-sampled measurements [1, 7], distributions that shift rapidly over time [10, 5], strict explainability requirements [11], and maintaining deployed models over time [12]. Additionally, health systems have systemic, technological, and legal barriers that prevent uptake of ML algorithms, even those that solve all the aforementioned problems [9]. Healthcare has been notoriously slow in adapting innovation in the digital age [8]. As other application areas of ML have seen substantial growth in generative models, multi-modal algorithms, and foundational models, the question of when and how these methods will apply to health has yet to be determined [2].

These challenges introduce interesting research problems that the community has been actively working on for the last few years, with significant room for contribution still remaining. Learning from time series for health is a uniquely challenging and important area with increasing application. Significant advancements are required to realize the societal benefits of these systems for healthcare. This workshop will gather together machine learning researchers dedicated to advancing the field of time series modeling in healthcare to bring these models closer to deployment.

Our workshop will focus on two themes: 1) Behavioral Health, and 2) Foundation Models. These themes are reflected in our invited speakers, call for participation, research round-tables, interactive panel discussion, and will be prioritized for spotlight talks. As each theme represents a lively research area, our workshop will help researchers congregate and discuss the state of the field and these areas. Our speakers will be informed of the themes and encouraged to share their opinions on challenges and opportunities in these areas.

Workshop Goals On top of fostering collaborations, networking, and mentorship, we aim to accelerate research in our area by:

1. Hosting discussions of how ML systems like foundation models and generative modeling can be applied to time series to improve human and behavioral health.

2. Highlighting recent successes and pitfalls of deployed time series systems, to broaden the benefits of future deployment.
3. Identifying promising approaches that perform well while maintaining privacy and other legal and ethical constraints necessary for deployment in healthcare settings.
4. Hosting mentorship and networking activities that prioritize underrepresented groups and broaden participation.

Anticipated Audience Our team has organized a similar workshop at NeurIPS 2022 and we had roughly 250 attendees. Though we anticipate smaller turnout from North American institutions due to the location, we believe this will be offset by attendees from European, Asian, and African institutes.

To cultivate a vibrant audience for our workshop, we are implementing several strategies: including leveraging social media platforms (Twitter/X), academic networks (Machine Learning for Health Symposium, Association for Health Learning and Inference), and industry partnerships to disseminate information and generate interest. We are committed to reaching underrepresented groups in machine learning and healthcare, ensuring invitations and promotional materials are circulated through affinity groups (WIML, Black in AI, and others).

Invited speakers

1. Dr. Mihaela Van Der Schaar, Professor, University of Cambridge (**confirmed**)
2. Dr. Marinka Zitnik, Assistant Professor, Harvard Medical School (**tentatively confirmed**)
3. Dr. Tim Althoff, Assistant Professor, University of Washington (**confirmed**)
4. Dr. Maxine Mackintosh, Program Lead - Diverse Data Initiative, Genomics England (**invited**)

Timeline Submission deadline: *February 3*. Author notification: *March 3*. Camera-ready papers due: *April 11*. Workshop takes place on *May 11*.

Tentative schedule

- 09:00 - 09:15: Opening remarks from organizers
- 09:15 - 09:45: Invited talk 1 (25 min + 5 min Q&A)
- 09:45 - 10:15: Invited talk 2 (25 min + 5 min Q&A)
- 10:15 - 11:15: Coffee break and poster session 1 (*discussion time*)
- 11:15 - 12:00: Thematic Spotlight Talks
- 12:00 - 13:00: Networking/Mentorship Lunch (*discussion time*)
- 13:00 - 13:30: Invited Talk 3 (25 min + 5 min Q&A)
- 13:30 - 14:00: Invited Talk 4 (25 min + 5 min Q&A)
- 14:00 - 15:00: Coffee break and poster session 2 (*discussion time*)
- 15:00 - 16:00: Topical Research Roundtable Discussions (small groups; *discussion time*)
- 16:00 - 16:55: Interactive panel with Roundtable Moderators and Invited Speakers (all attendees; *discussion time*)
- 16:55-17:00: Closing remarks from organizers

Ensuring time for discussion We have carefully curated the schedule to ensure adequate time for discussion that is spread out through the day. There will be two poster sessions (one in the morning and one in the afternoon) to give presenters the opportunity to engage with other attendees, receive feedback, and build new collaborations. We will also select a few top papers for a thematic spotlight talk - where authors who are pursuing excellent research in our highlighted areas of interest can share their findings. These talks will be followed by a Q&A session to further promote discussion around the selected themes for this year's workshop. Lunch time will feature a networking/mentorship lunch for trainees. We have leftover funding from a prior workshop to support this and will seek further sponsorship to supplement. The mentorship lunch will pair trainees with mentors to encourage new

research collaborations and make this a more welcoming venue for underrepresented groups. We will prioritize underrepresented groups to be selected for this opportunity

Lastly, the day will conclude with a research roundtable followed by a panel discussion. There will be 3 roundtables for the attendees to choose from, each focusing discussion on a different thematic area. Our moderators will encourage discussion and engagement from participants and will prepare a list of topics to discuss with the panel. After the roundtable session, the whole workshop will convene for an interactive panel with the roundtable moderators and invited speakers. This final session will synthesize key points of discussion from the day and identify research directions moving forward.

Diversity Commitment We also ensure diversity in our organizers and speakers in terms of gender (55% are women), race (27% are non-white), location/affiliations (10 universities and 1 company across 3 countries) and seniority (including one PhD student, two postdoctoral associate, one research scientist, two assistant professors, and two full professors). We also have a variety of workshop-organizing experience - with organizers that are well-experienced to junior researchers still establishing their academic networks.

Additionally, we will have a mentorship lunch that will prioritize trainees from underrepresented groups. Student mentees will be paired with volunteer mentors with similar research interests. We will advertise recruitment for the mentorship lunch, as well as the call for papers, through affinity groups (WIML, Black in AI, and others). If sponsorship allows, we will provide travel grants to students from underrepresented groups, which we did last year.

Access In alignment with our commitment to fostering a diverse and inclusive environment, we will implement various strategies to ensure broad access to the workshop's content and discussions. Recognizing the global nature of the research community and potential travel constraints, we will record invited talks and spotlight talks for asynchronous viewing. Additionally, we will make accepted papers available on our website (subject to author approval) to engage those unable to attend in person. This approach not only democratizes access to knowledge but also caters to different time zones and geographical locations, ensuring that the invaluable insights and discussions generated during the workshop are accessible to the wider community, regardless of physical presence.

Previous Related Workshops Our organizing team has experience organizing a similar workshop at NeurIPPS 2022. This was an incredibly successful showing with 62 submissions and 46 accepted papers. The accepted papers were 4 pages each. We counted around 250 attendees and note that there was standing-room only throughout the day in the workshop room. We believe that the focus on time series and health is of sufficient interest substantial portion of the community. This intersectional research area is large but has had few opportunities to congregate.

Due to the interest surrounding medical time series research, workshops and tutorials on similar themes have been held in recent years. Specifically, the "Workshop on Time Series Representation Learning for Health" (TSRL4H) was held at ICLR 2023, ECML 2022 hosted the "7th Workshop on Advanced Analytics and Learning on Temporal Data" (AALTD), and AAAI 2022 held the "Time Series in Healthcare: Challenges and Solutions" (TSH:CS) tutorial. TSRL4H differs from our proposed workshop by focusing mainly on topics relating to representation learning, rather than our themes of Behavioral Health and Foundation Models. AALTD differed by not primarily focusing on healthcare data and problems. Instead, it more broadly focused on time-series, where medical time series was only one of several topics covered. Lastly, TSH:CS fostered discussion but as it was a tutorial it did not allow for the timeseries for healthcare research community to disseminate their work.

Reviewing We will use the OpenReview system to enforce the following policy: any reviewer (including organizers) cannot assess any submission from someone (1) who has been a colleague with the reviewer within the same organization in the past 3 years; (2) who has co-authored publications with the reviewer in the past 3 years; (3) is currently from the same institution as the submitting authors.

We will make sure to recruit reviewers from a wide variety of institutions and training levels to ensure the review process is unbiased. Lastly, only unpublished work will be accepted for workshop proceedings. Anything that has already been published elsewhere will be desk-rejected.

Organizers

Organization Team:

Sujay Nagaraj is an MD/PhD student in the Department of Computer Science at the University of Toronto and the Vector Institute. His research is on medical time-series data with a focus on mobile health / wearable data. Sujay mainly works on making sense of multi-variate time series from wearable devices, and is also studying how to learn robust classifiers in the presence of label noise. He has served on the program committee for NeurIPS in 2022, has been a moderator for the 2022 Machine Learning for Health (ML4H) Symposium, and a sub-chair for the 2023 ML4H Symposium. Email: s.nagaraj@mail.utoronto.ca

Walter Gerych is a Postdoctoral Associate at MIT CSAIL. He received his PhD in Data Science at Worcester Polytechnic Institute. His research is on bias identification and correction, primarily motivated by human activity recognition using time series of mobile sensor data. Email: wgerych@mit.edu

Sana Tonekaboni is a postdoc fellow at the Eric and Wendy Schmidt Center of MIT and Harvard. She completed her PhD in Computer Science at the University of Toronto. Her research focus is on challenges in modeling time series and barriers to adoption of ML in healthcare environments. She mainly works on representation learning and Bayesian modelling of time series in healthcare, and also approaches for explaining complex temporal models. Sana has co-chaired the NeurIPS 2022 workshop on Learning from Time Series for Health, and has previously served on the program committee for NeurIPS, ICLR, ICML, AISTATS, MLHC, and CHIL. Email: stonekaboni@cs.toronto.edu

Tom Hartvigsen is an incoming assistant professor at the University of Virginia and currently a postdoc at MIT CSAIL. His research is on deployable machine learning for healthcare. Tom is the General Chair of the 2023 Machine Learning for Health Symposium, is an organizer for the ICML 2023 Workshop on Deployable Generative AI, is an organizer for the 2022 Conference on Health, Informatics, and Learning, and co-chaired the NeurIPS 2022 workshop on Learning from Time Series for Health. E-mail: tomh@mit.edu

Ahmed Alaa is an Assistant Professor of Computational Precision Health at UC Berkeley and UCSF, with affiliations in the EECS and Statistics departments at UC Berkeley. Previously, he was a postdoctoral associate at Massachusetts Institute of Technology (MIT CSAIL and IMES) and the Broad Institute of MIT and Harvard University. He was also a joint postdoctoral scholar at Cambridge University, Cambridge Center for AI in Medicine and the University of California, Los Angeles (UCLA). He obtained his Ph.D. in Electrical and Computer Engineering from UCLA, where he received the 2021 Edward K. Rice Outstanding Doctoral Student Award from the UCLA Samueli School of Engineering. His research interests include machine learning for healthcare, computer vision for medical imaging, clinical informatics, statistics, and causal inference.

Emily Fox is a Professor in the Department of Statistics and Department of Computer Science at Stanford University. Prior to Stanford, she was the Amazon Professor of Machine Learning in the Paul G. Allen School of Computer Science & Engineering and Department of Statistics at the University of Washington. From 2018-2021, Emily led the Health AI team at Apple, where she was a Distinguished Engineer. Emily has been awarded a CZ Biohub Investigator award, Presidential Early Career Award for Scientists and Engineers (PECASE), Sloan Research Fellowship, ONR Young Investigator award, and NSF CAREER award. Her research interests are in modeling complex time series arising in health, particularly from health wearables and neuroimaging modalities. E-mail: ebfox@stanford.edu

Anna Goldenberg is a Varma Family Chair in Biomedical Informatics and Artificial Intelligence at SickKids Research Institute. She is an Associate Professor in the Department of Computer Science at the University of Toronto, faculty member and an Associate Research Director, Health at Vector Institute and Canadian AI Chair at CIFAR, as well as a member of CIFAR's Child and Brain Development group. Dr Goldenberg trained in machine learning at Carnegie Mellon University, with a postdoctoral focus in computational biology and medicine. The current focus of her lab is on developing and deploying machine learning models to healthcare. Dr Goldenberg is strongly committed to creating responsible AI to benefit patients across a variety of conditions. Email: anna.goldenberg@utoronto.ca

Program Committee: We have a large pool of reviewers to draw from previous workshops we have organized:

Bret Nestor (UofT), Achille Nazaret (Columbia), Qixuan Jin (MIT), Robert Logan (UCI), Siddharth Biswal (Georgia Tech), Alex George Adam (UofT), Karine Tung (UMass Amherst), Lauren Erdman (UofT), Kingsley Chang (Meta), Hugo Yeche (ETH Zurich), Kimia Hamidieh (UofT), Lovedeep Gondara (SFU), Hyewon Jeong (MIT), Junyi Gao (HDRUK), Divya Shanmugam (MIT), Aslesha Pokhrel (UofT), Aya Abdelsalam Ismail (Maryland), Abhishek Moturu (UofT), Hui Wei (UMass Amherst), Taylor Killian (UofT), Aakash Kaku (NYU), Mai Ali (UofT), Jenny Yu (UofT), Tina Behrouzi (UofT)

References

- [1] Roohallah Alizadehsani, Mohamad Roshanzamir, Sadiq Hussain, Abbas Khosravi, Afsaneh Koohestani, Mohammad Hossein Zangoeei, Moloud Abdar, Adham Beykikhoshk, Afshin Shoeibi, Assef Zare, et al. Handling of uncertainty in medical data using machine learning and probability theory techniques: A review of 30 years (1991–2020). *Annals of Operations Research*, pages 1–42, 2021.
- [2] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- [3] Christina Catley, Heidi Stratti, and Carolyn McGregor. Multi-dimensional temporal abstraction and data mining of medical time series data: Trends and challenges. In *2008 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, pages 4322–4325. IEEE, 2008.
- [4] Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. Recurrent neural networks for multivariate time series with missing values. *Scientific reports*, 8(1):6085, 2018.
- [5] Taesung Kim, Jinhee Kim, Yunwon Tae, Cheonbok Park, Jang-Ho Choi, and Jaegul Choo. Reversible instance normalization for accurate time-series forecasting against distribution shift. In *International Conference on Learning Representations*, 2021.
- [6] Rahul Krishnan, Dawen Liang, and Matthew Hoffman. On the challenges of learning with inference networks on sparse, high-dimensional data. In *International conference on artificial intelligence and statistics*, pages 143–151. PMLR, 2018.
- [7] Steven Cheng-Xian Li and Benjamin Marlin. Learning from irregularly-sampled time series: A missing data perspective. In *International Conference on Machine Learning*, pages 5937–5946. PMLR, 2020.
- [8] Bertalan Meskó, Zsófia Drobni, Éva Bényei, Bence Gergely, and Zsuzsanna Györfly. Digital health is a cultural transformation of traditional healthcare. *Mhealth*, 3, 2017.
- [9] Iredia M Olaye and Azizi A Seixas. The gap between ai and bedside: Participatory workshop on the barriers to the integration, translation, and adoption of digital health care and ai startup technology into clinical practice. *Journal of Medical Internet Research*, 25:e32962, 2023.
- [10] Adarsh Subbaswamy and Suchi Saria. From development to deployment: dataset shift, causality, and shift-stable models in health ai. *Biostatistics*, 21(2):345–352, 2020.
- [11] Sana Tonekaboni, Shalmali Joshi, Melissa D McCradden, and Anna Goldenberg. What clinicians want: contextualizing explainable machine learning for clinical end use. In *Machine learning for healthcare conference*, pages 359–380. PMLR, 2019.
- [12] Jenna Wiens, Suchi Saria, Mark Sendak, Marzyeh Ghassemi, Vincent X Liu, Finale Doshi-Velez, Kenneth Jung, Katherine Heller, David Kale, Mohammed Saeed, et al. Do no harm: a roadmap for responsible machine learning for health care. *Nature medicine*, 25(9):1337–1340, 2019.