AI HEALTH SUMMIT

Date: 23 - 24 November 2023 | Venue: W Singapore - Sentosa Cove, 21 Ocean Way, S098374

Abstract Submission Form

The AI Health Summit 2023 will host poster competitions and stand to win generous prizes! All shortlisted presenters will have 5 minutes each to present their posters on stage.

Timeline:

- 30 Sep 2023: Closing date for abstract submission
- 30 Oct 2023: Notify of selected abstract
- 1 Nov - 15 Nov 2023: Submission of posters

Abstract Online Submission Site:
https://form.gov.sg/64a3cbdc9693f400117e275c

Naming Your Abstract File:
Full name of presenting author_Country_Short abstract title

Each abstract should only represent **ONE** theme. Author(s) may submit multiple entries. The committee reserves the right to disqualify unsuitable entries. This submission form for abstracts will be used in shortlisting for poster submissions.

Abstract theme (Please select only **ONE**):

	AI in medical imaging		AI in drug discovery, bioinformatics, and precision medicine
	AI in healthcare operations		AI governance, regulation, ethics and social impact
	AI in medical education	x	AI in health informatics (e.g., electronic health records)
	AI in global health		Natural language processing, large language models, generative AI
	AI in biomedical signal processing (e.g., ECG, EEG)		Emerging AI technologies (e.g., federated learning, blockchain)

Contact details of Presenting Author

Title : ☐ Prof ☐ A/Prof ☒ Dr ☐ Mr ☐ Ms ☐ Mrs

Full name: Cao Zhen

Institution: Institute for Infocomm Research

Designation: Senior Research Engineer

Department: Machine Intellection Department

Country: Singapore

Email: Cao_Zhen@i2r.a-star.edu.sg

Full List of Authors and Affiliations

Name: Cao Zhen[1], Feri Guretno[1], Jeanette Poh Wen Jun[1], Guo Yang[1], Mila Nambiar[1], Chan Yu En[1], Watt Pei Yi[1], Gopinathan Arjunan[1], Yong Mong Bee[2], Pavitra Krishnaswamy[1,*], Savitha Ramasamy[1,*]

** equal contribution*

Institution: [1]Institute for Infocomm Research, A*STAR, Singapore
[2]Singapore General Hospital, SingHealth

Abstract

- Abstracts should be between 250-350 words.
- At most 1 table/figure is allowed.
- Authors are recommended to follow reporting guidelines such as CONSORT-AI, SPIRIT-AI, STARD-AI, and DECIDE-AI. More details are available in the following article: https://doi.org/10.47102/annals-acadmedsg.2022452

Title: A Performance Monitoring and Continual Learning Framework to Improve Reliability of Prediction Models based on Electronic Health Records

Background and Aim

Deep learning approaches have demonstrated tremendous capabilities for a variety of prediction tasks based on Electronic Health Record data. However, the characteristics of EHR data in clinical settings are known to change continually, due to changes in diagnostic criteria, measurement devices, treatment protocols, or patient cohort attributes. These shifts in data distributions can lead to considerable degradation in the performance of deep learning models in deployment, fundamentally limiting reliability of model inferences and deterring adoption in clinical practice.

Here, we introduce an integrated framework to continually monitor model performance and seamlessly update trained models when performance degradation is observed.

Methods

First, we leveraged statistical process control tools to design a module for regular monitoring performance of deep learning-based prediction models, in relation to observed ground truth. We implemented a dashboard to visualize model performance as new data streamed in. Second, to address any significant or consistent performance drops observed, we developed a Bayesian continual learning algorithm that can adapt its neural representations to data distribution changes. Specifically, we used a Bayesian Long Short-Term Memory model (BLSTM) backbone and developed an efficient means to continually update model representations via a combination of architectural pruning, regularization and replay strategies while avoiding catastrophic forgetting. Finally, we integrated the performance monitoring and continual learning capability to demonstrate our framework. We focused on a HbA_{1c} prediction use case, based on EHR data from the Singapore Diabetes Registry (2013-19, 22 sites).

Results

In test scenarios with substantial drops in HbA_{1c} prediction performance across time (2013-18), the continually adapted model showed R^2 improvements of ~4% over the original model trained on data from prior time periods. Moreover, on a prospective 2019 test set, the continual learning model enabled R^2 improvements of ~8% over the original model trained on 2013-18 data.

Conclusion

Our integrated performance monitoring and continual learning framework efficiently and seamlessly addresses performance drops due to drifts in data distribution. As such, it could improve reliability of predictive models deployed in real-world decision support tasks.

Please add table / figure here, if any.