

Results from the first EHR DREAM Challenge: Patient Mortality Prediction

Timothy Bergquist, BS¹, Yao Yan, BS¹, Thomas Schaffter, PhD², Thomas Yu, BS², Vikas Pejaver, PhD¹, Noah Hammarlund, PhD¹, Justin Prosser, BS¹, Justin Guinney, PhD², Sean Mooney, PhD¹

¹University of Washington, Seattle, WA, USA; ²Sage Bionetwork, Seattle, WA, USA

Implementing the ‘Model to Data’ framework to enable sandboxing with EHR data

Attendees of this talk will see how our ‘Model to Data’ evaluation infrastructure enables healthcare organizations to crowdsource predictive analytics solutions using protected health data without exposing patient data and will be given the resources to implement this method within their own institution.

Balancing the Security, Usability, and Privacy of Electronic Health Records

The past decade has seen a substantial increase in the use of Electronic Health Record (EHR) systems.¹ Given the size and complexity of this data, machine learning approaches have been used to provide insights in a more automated and scalable manner. Regulatory policies such as HIPAA and HITECH place the onus and financial burden of ensuring the security and privacy of the patient records on the healthcare institutions hosting the data. An unintended consequence of these regulations is the difficulty of sharing clinical data within the research community, slowing the rate at which studies can be conducted using EHR data. The ‘Model to Data’ (MTD) framework for enabling machine learning research on private biomedical data was described by Guinney et al. as an alternative to traditional data sharing methods.² Instead of granting researchers access to the data, the researcher sends a containerized model to the data hosts who are then responsible for running the model on the researcher’s behalf. Using the ‘Model to Data’ framework, we carried out a community challenge (a crowdsourcing competition), the EHR DREAM Challenge: Patient Mortality Prediction where we asked participants to build mortality prediction models that were evaluated against electronic health record (EHR) data from the University of Washington.

Methods

For this study, we leveraged a large EHR dataset from the University of Washington Clinical Data Warehouse, which represented 1.3 million patients. Challenge participants were asked to develop a prediction model to answer the following question: *Given the past electronic health records of each patient, predict the probability that he/she will pass away within the next six months following his/her last visit.* The participants built and tested mortality prediction models using a synthetic dataset that had the same format and structure as the UW data. The participants containerized these models, using the Docker software designed to facilitate the sharing of scripts and dependencies in a single unit called a Docker image,³ and submitted them for evaluation through Synapse, an open-source software platform developed by Sage Bionetworks for researchers to share data, collaborate, and submit to challenges.⁴ A Common Workflow Language pipeline, located in the UW secure environment, detected these submissions and pulled them into the UW computing environment. The submitted Docker containers were run on the UW data, predictions were generated for patients in an evaluation dataset (a holdout set from the full EHR dataset), and an Area Under the Receiver Operator Curve (AUROC) score was generated by comparing the predictions to the true six-month mortality status of those patients. This score was returned to the participants through Synapse. (See workflow in Figure 1a.)

Results

The EHR DREAM Challenge: Patient Mortality Prediction was run from May 9, 2019 to January 9, 2020. As of December 3 2019, 17 teams submitted 39 models without having any access to the data. These submitted models were evaluated and scored against UW EHR data by comparing the output predictions

of the models to the true six-month mortality status of all the patients in the evaluation dataset (~390,000). The top six highest performing teams all built models that had AUROC > 0.90 (Figure 1b).

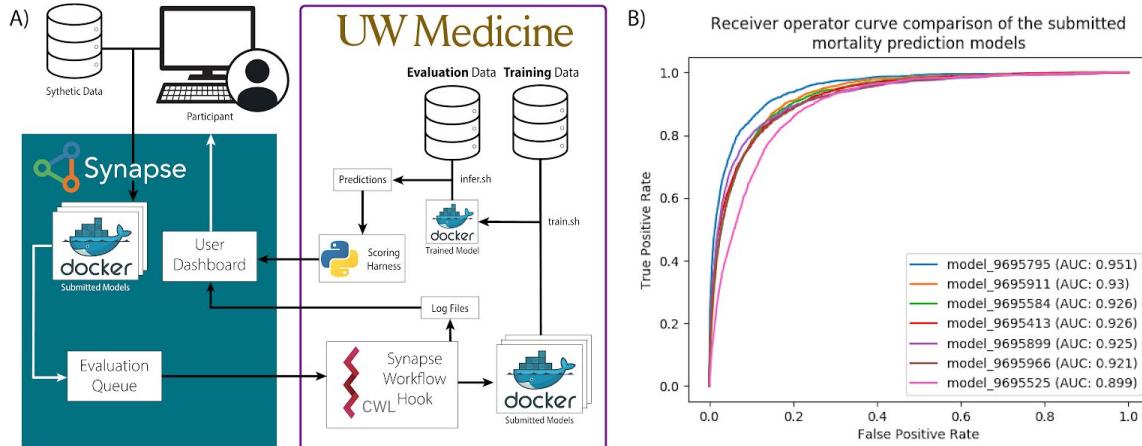


Figure 1. (A) Diagram for submitting and distributing containerized prediction models in a protected environment. Dockerized models were submitted to Synapse by the participants. The CWL Synapse Workflow Hook pulled in the submitted Docker image which trained on the available training EHR data, and then made inferences on the evaluation dataset, outputting a prediction file. The prediction file was compared to a gold standard benchmark and the scores were returned to the participant. (B) A comparison of the receiver operator curves for the top seven mortality prediction models submitted, trained, and evaluated against UW EHR data using the MTD framework.

Discussion

In this study, we implemented the ‘Model to Data’ framework in the context of an institutional clinical data warehouse and demonstrated how model developers can develop clinical predictive models without having direct access to patient data. While we focused on the specific task of mortality status prediction in this community challenge, our platform would naturally be generalizable to other statistical or predictive questions.

Conclusion

We demonstrate the potential impact of the ‘Model To Data’ (MTD) framework to bring clinical predictive models to private data by operationalizing this framework to enable a community challenge using protected UW Medicine EHR data while protecting patient privacy. This work serves as a demonstration of the MTD approach in a real-world clinical analytics environment. We believe this enables future predictive analytics sandboxing activities and the development of new clinical-predictive methods safely.

References

1. Charles D, Gabriel M, Ma TSM. Adoption of Electronic Health Record Systems among U.S. Non-Federal Acute Care Hospitals: 2008-2014. Available from: <https://www.healthit.gov/sites/default/files/data-brief/2014HospitalAdoptionDataBrief.pdf>
2. Guinney J, Saez-Rodriguez J. Alternative models for sharing confidential biomedical data. Nat Biotechnol [Internet]. 2018 May 9;36(5):391–2. Available from: <http://dx.doi.org/10.1038/nbt.4128>
3. Enterprise Container Platform | Docker [Internet]. Docker. [cited 2019 Aug 9]. Available from: <https://www.docker.com/>
4. Saez-Rodriguez J, Costello JC, Friend SH, Kellen MR, Mangravite L, Meyer P, et al. Crowdsourcing biomedical research: leveraging communities as innovation engines. Nat Rev Genet [Internet]. 2016 Jul 15;17(8):470–86. Available from: <http://dx.doi.org/10.1038/nrg.2016.69>