AMIA 2018 Annual Symposium

View Abstract

Title: DeepAISE on FHIR — An Interoperable Real-Time Predictive Analytic Platform for Early Prediction of Sepsis

Abstract Body:

Project abstract (1050 Characters): Sepsis, a dysregulated immune-mediated host response to infection, is lethal, prevalent, and costly. It's early detection has the potential to drastically reduce morbidity/mortality. We have developed a real-time cloud-based application that predicts onset-time of sepsis based on live ICU data and provides clinicians with actionable visual alerts. Clinicians and nurses can examine these alerts and initiate appropriate interventions. The prediction engine (DeepAISE) is a Deep Learning-based algorithm trained to reliably predict sepsis 4-6 hours in advance of clinical recognition. A scalable, cloud-based, system continuously streams bedside data and uses the prediction engine to generate hourly scores and displays these to clinicians. Interoperability is achieved through the use of FHIR resources and APIs. This system is monitoring ~100 patients on a daily basis at the Emory Tele-ICU center, and has been shown to reliably predict onset of sepsis with an AUC of 0.9.

Project rationale, impact and innovation (3500 Characters): Sepsis, an autoimmune response to infection, is one of the deadliest and most expensive conditions treated in U.S. hospitals, affecting approximately 750,000 Americans each year. Early recognition of septic patients has substantially improved the care they receive, but by the time sepsis is clinically apparent it is often too late to prevent significant morbidity and mortality. Our multidisciplinary team of machine learning and critical care experts has recently demonstrated that an Artificial Intelligence Sepsis Expert (AISE) algorithm trained on electronic medical record (EMR) data from over 30,000 patients from the Emory Healthcare system, and validated on an independent cohort of 50,000 patients from the MIMIC-III ICU database, can reliably predict new sepsis in the ICU (defined using an automated approach based on sepsis-3, the most recent consensus of sepsis [Singer et al., JAMA 2016]), 4-6 hours in advance, with an area under the curve (ROC) of 0.85 [Nemati et al., PMID 29286945].

In addition to accurate prediction of sepsis, the algorithm provides a list of the most important clinical variables contributing to the risk score at any given point in time, thus ensuring interpretability. A clinically viable version of the algorithm has been developed that utilizes Emory's Tele-ICU (eICU) real-time data feeds (including vitals, labs, and demographics), harmonized in a Fast Healthcare Interoperability Resource (FHIR) compatible database, to generate hourly sepsis risk scores for over 100 patients at any given time. We have recently demonstrated that a deep learning version of this algorithm (DeepAISE-on-FHIR) achieves a testing AUC of 0.90. Since early detection is key, clinicians can see the score and its rationale when treatment with antibiotics is most effective.

For short terms effects, such a system would impact how early doctors are able to detect the condition thereby reducing sepsis mortality rates and intervening before the situation gets worse. Given that sepsis is one of the most expensive conditions treated in hospitals, such a system would have monetary benefit for hospitals. Long term effects would be a better understanding of the infection itself and various factors contributing to it and possibly directions for predicting other conditions in the ICU.

Project design and implementation (7000 characters): The system is architected as data captured from patient bedsides; streamed into a predictive engine that runs Deep-AISE, and produces a composite Sepsis Risk Score displayed on a dashboard that is designed to be easy for clinicians to read at a glance. An visual alarm system notifies clinicians when any patient reaches a threshold likely to develop sepsis, making it easier for busy caregivers to respond quickly. The system was adapted for and deployed on the Google Cloud Platform (GCP) for scalability and redundancy. It uses three GCP managed services namely: 1) the Google Healthcare API as the underlying FHIR server; 2) Pub/Sub as the data streaming system; 3) Dataflow as the data transformation engine; 4) ML-Engine and TensorFlow as the execution environment for the prediction engine; and 5) Kubernetes and the Google Container Engine for scalability, redundancy, and easy deployment.

This cloud-based version is currently available in demo-mode (via http://sepsis.app:3001/). De-identified data from the Emory eICU database is sent to a HIPAA compliant cloud storage from the Emory Tele-ICU systems. DeepAISE-on-FHIR utilizes the Google Healthcare APIs, and the Google Pub/Sub data streaming services to stream the relevant FHIR resources (i.e., data elements such as patients' latest laboratory values, heart rate, respiratory rate, etc., typically at 15 minute to 1 hour resolution) to our prediction engine. Our design

philosophy was to avoid 'boiling the ocean' by focusing on mapping of a limited set of EMR data elements required by the DeepAISE into FHIR resources, thus ensuring interoperability.

A major challenge was being able to harmonize data with FHIR APIs and ensuring security of medical data. There was a lot of research done on the current state of data integration across EMR vendors and the upcoming technologies and specifications built. We also built modules to convert the observational data into FHIR resources which passed through data pipelines. The system uses HIPAA compliant storage and services. Along with this, our collaboration with Smartronix, Inc. ensured compliance and regulatory requirements are met. By capturing logs related to service usage, data movements across the platform, network and infrastructure, we receive alerts for security events that carry risks. Notifications are given so that response action can be immediate.

Project evaluation and sustainability (3500 characters): Evaluation

Design requirements were collected from our ICU nurses in early 2017. The main objectives were workflow awareness, actionability, and data integrity, and application stability. We conducted a series of interviews with our nurses to ensure that our design aligned with their day-to-day tasks involving evaluation of patients and assessment of risk for sepsis. The interface allowed for direct incorporation of clinical actions via drag-and-drop operations into appropriate treatment categories. Data integrity assessment was carried out via manual chart-review of patient records (vitals, labs, etc) to ensure that the displayed patient data on the user interface matches the corresponding values within the clinician facing EHR application (Cerner's PowerChart). Application stability assessment was performed by keep track of up-time of each containerized application. We continue to keep track of these evaluation metrics in preparation for an FDA submission.

Sustainability

The project is currently funded through a combination of NIH and DoD grants. We have built a collaborative network of academic users and industry partners for this project. The software is continuously developed and maintained by a team of researchers and developers.
The individual components deployed on Google Cloud Platform are well documented and maintainable. The infrastructure effort is low since most of the components are run on Google's managed services like Dataflow, PubSub and Machine Learning Engine . The project also ensures interoperability. By building a Fast Healthcare Interoperability Resources (FHIR) database on GCP, we ensure that the engine can scale and operate across institutions on a reliable, secure, and private platform.

Twitter project summary (140 characters): DeepAISE: a scalable cloud-based AI app predicts onset of Sepsis in ICU patients.It uses FHIR for interoperability across EMRs and hospitals

Customers: No

Solution Date: 06/01/2017

Implementation Date: 07/01/2018

How many users does your solution have or how many patients have been impacted by it (please indicate time frame)?: 2663

Authors: Vidyashankar Lakshman[1], Fatemeh Amrollahi[1], Veera Supraja Koppisetty[1], Supreeth P. Shashikumar[1], Ashish Sharma[1], Shamim Nemati[1]

Institutions: 1. Emory University, Irvine, CA, United States.

This study was partially funded by the National Science Foundation Grant # 1636933 (BD Spokes: SPOKE: SOUTH: Large-Scale Medical Informatics for Patient Care Coordination and Engagement).