

AIDEx - An Open-source Platform for Real-Time Forecasting Sepsis and A Case Study on Taking ML Algorithms to Production

Fatemeh Amrollahi^{*1}, Supreeth Prajwal Shashikumar¹, Pradeeban Kathiravelu²,
Ashish Sharma² and Shamim Nemati¹

Abstract—Sepsis, a dysregulated immune response to infection, has been the leading cause of morbidity and mortality in critically ill patients. Multiple studies have demonstrated improved survival outcomes when early treatment is initiated for septic patients. In our previous work, we developed a real-time machine learning algorithm capable of predicting onset of sepsis four to six hours prior to clinical recognition. In this work, we develop AIDEx, an open-source platform that consumes data as FHIR resources. It is capable of consuming live patient data, securely transporting it into a cloud environment, and monitoring patients in real-time. We build AIDEx as an EHR vendor-agnostic open-source platform that can be easily deployed in clinical environments. Finally, the computation of the sepsis risk scores uses a common design pattern that is seen in streaming clinical informatics and predictive analytics applications. AIDEx provides a comprehensive case study in the design and development of a production-ready ML platform that integrates with Healthcare IT systems.

I. INTRODUCTION

A. Sepsis – A Health Crisis

Sepsis is a syndromic, life-threatening condition that occurs when the body exerts an exaggerated response to infection [1]. Sepsis when left untreated progresses to deadly severe sepsis or septic shock that the host triggers injuring internal organs. Nearly 6% of the inpatient hospital population in the United States will carry a diagnosis of sepsis during their stay. 35% of all hospital deaths, in the US, are attributed to sepsis, and it accounts for \$23.7 billion in annual costs[2]. Numerous trials have demonstrated dramatic improvements in survival outcomes for sepsis by early recognition of the condition and rapid treatment [3-6]. While there are effective protocols for treating sepsis once it has been diagnosed, there are several challenges in reliably identifying septic patients early in their course owing to the significant variability in the disease’s presentation. The Sepsis-3 guidelines[1] have narrowed the constellation of signs and symptoms of sepsis into a clinical criterion that can be reliably used by clinicians and researchers to identify this life-threatening condition retrospectively. However, this criterion alone cannot help identify a patient, who is experiencing the effects of sepsis early in the disease’s course.

B. Sepsis Prediction using Machine Learning

In recent years, the increased adoption of electronic medical records (EMR) has spurred the development of machine learning based surveillance tools for detection [7-10] and prediction [9-12] of patients with sepsis or septic shock. However, there has been slow progress in real-time implementation of a high-dimensional machine learning model in an Intensive Care Unit (ICU) environment.

Recently, Nemati et al have developed the Artificial Intelligence Sepsis Expert (AISE), a modified Weibull-Cox model that uses data commonly available in the EMR to predict the onset of sepsis four to six hours in advance with an area under the ROC curve (AUC) of 0.85 [11]. The AISE development cohorts contained over 30,000 patients from multiple hospitals in the Emory Healthcare system and was validated using a cohort of 50,000 patients from the MIMIC-III database [13]. In this work, we present a platform, that is used to deploy the AISE algorithm in a real-world setting, using live clinical data.

Our platform fetches patients records from a real-time EMR database and displays hourly sepsis risk score for each patient. The platform, called Artificial Intelligence Decompensation Expert (AIDEx), is scalable, resilient, open-source and developed using the emerging Fast Healthcare Interoperability Resources (FHIR) standard.

The overarching emphasis of our architecture is to examine integration with healthcare IT systems, significant attention has been given to elements such as software quality control, and tracking feature drifts. The user interface has been designed to minimize false-alarms [12] as well as assist in clinical interpretability and workflow integration requirements necessary for a successful clinical decision support (CDS) system [14-15]. AIDEx sheds some light on the processes and development needed to interface with healthcare IT systems and build deployable applications. Thus, AIDEx¹ also provides a good case-study in taking ML based predictive analytics algorithms, for healthcare applications, to production.

II. PLATFORM ARCHITECTURE

ML algorithms are usually developed within controlled environments where researchers can control data, wrangle it into an appropriate format, and once the model has been developed, evaluate, and validate it suitably. The deployment

¹The code for AIDEx platform is available online at <https://github.com/aise-on-fhir>

* Corresponding Author Email: famrolla@eng.ucsd.edu
¹ Division of Biomedical Informatics, University of California San Diego, La Jolla, CA - 92093
² Department of Biomedical Informatics, Emory University Atlanta, GA - 30341

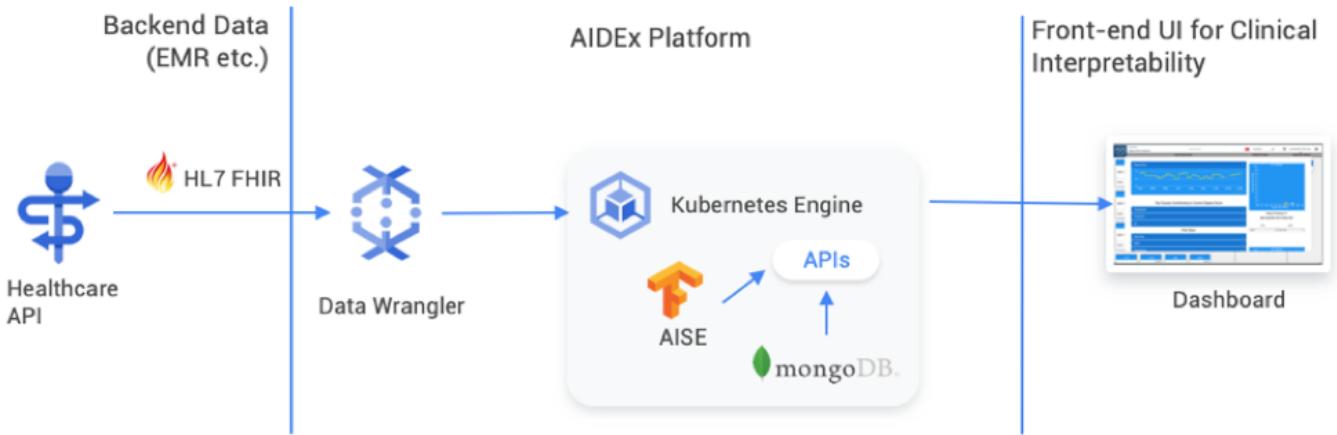


Fig. 1. An overview of the AIDEx platform that illustrates how healthcare data is pulled from the FHIR data store and then the results are streamed into a MongoDB based results services. Clients can use the web-based dashboard to review results

of a trained model, into a real-world setting is a non-trivial activity that can include data-wrangling pipelines that can operate without human intervention; on-demand deployment and scaling of the algorithm; monitoring of the underlying infrastructure; tracking and tuning for performance and latency; user interfaces; and quality control. In most real-world ML systems, the actual ML algorithm or model, is a lot smaller than the infrastructure needed to support it [16]. ML algorithms and systems therefore incur, what is known as the technical debt of ML. This hidden technical debt, results in a highly incomplete view of the field and, by overly simplifying the process, contributes to the hype-cycles. One of the key contributions of AIDEx, is its ability to work with a real-time stream of live clinical data. It required the design and implementation of a real-world system that is scalable, elastic, and fault-tolerant.

AIDEx adopts a modular architecture where each of its core functionality is captured as a microservice. It builds microservices for preprocessing data, executing the prediction algorithm, storing the prediction outcomes, and visualizing outcomes. Figure 1 illustrates the flow of data in AIDEx. It has been deployed on a Cloud Platform, and uses a few managed services and functionalities, including security features such as firewalls and virtual private clouds (VPCs). It is worth noting that while this paper focuses on deployment of AIDEx on the Cloud, AIDEx is also a platform agnostic container-based system that can be deployed on the cloud as well as on-premises.

A. AIDEx Microservices

AIDEx consumes patient data as a series of FHIR resources, computes the risk of developing sepsis in the next four to six hours (sepsis scores), and presents them in an interpretable manner via a web-based dashboard. The environment is secured via a VPC and utilities have been deployed to push data from the institution to the cloud. The use of containerized microservices removes the need to install distinct applications and their associated dependencies on a host machine at various deployment sites. It also allows

us to leverage the inherent scalability and fault tolerance. The AIDEx pipeline is unique in its health system agnostic design and its use of a state-of-the-art machine learning algorithm capable of accurately identifying patients with sepsis early in their clinical course. Though this tool provides population-level surveillance of a large cohort of ICU patients, its real strength lies in its ability to provide clinicians with individual patient vital sign trends and the most relevant features contributing to their risk score. Table 1 presents a summary of the services comprising the pipeline and are described in greater detail in the following sections.

TABLE I
MICROSERVICES THAT MAKE UP THE AIDEX PLATFORM.

Name	Service Objective
Data Wrangler	Retrieves live data streams as FHIR Resources; prepares data; orchestrates predictive algorithm; saves data
Results Store	A time-series data store that stores the data for each patient at each time-point
Sepsis Predictor	Algorithm running inference on data for each patient and forecasts the onset of sepsis
Clinical Dashboard	Presents the outcomes in an interpretable user-friendly interface

1) *Data Wrangler - Clinical Data Harmonization:* Pre-processing the data is a crucial first step in the machine learning applications. Data arriving from an active EMR is not always ready for use by a machine learning algorithm and requires a series of pre-processing steps. AIDEx consists of a Data Wrangler service that pulls real-time patient data from the host EMR's FHIR database to the sepsis predictor service, and finally to the results store. Figure 1 illustrates the data flow in the AIDEx platform. The Data Wrangler service starts its execution by querying a live EMR FHIR database capturing the patient features necessary for sepsis prediction. These features include laboratory results, vital signs, and demographic information for all active patients over the last

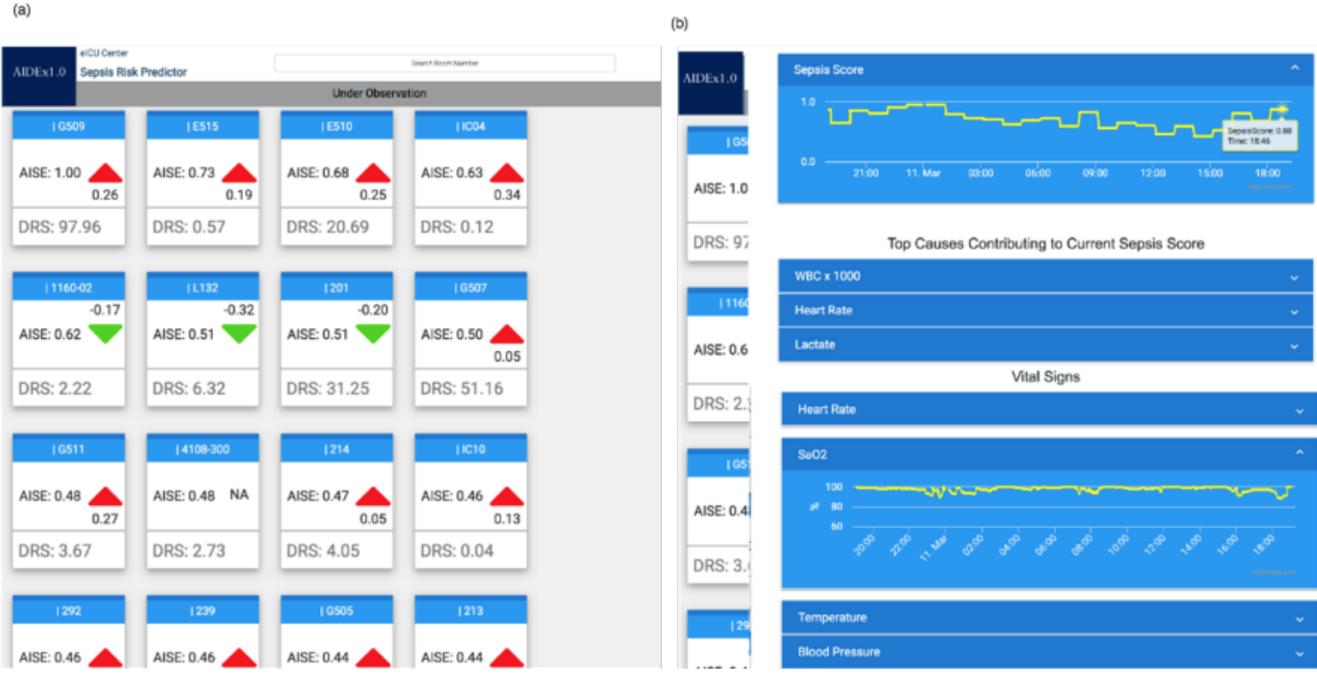


Fig. 2. A population view of all ICU patients. Each patient is represented by a single card that displays: a sepsis risk score (AISE), discharge readiness score (DRS), the increase in patient's sepsis score (delta). Patients are listed descending based on their AISE score (on left). The detailed view for a single patient is displayed. Visible is the 12-hour trajectory of the patient's sepsis risk score, and their vital signs.

hour. Errors in data entry can result in values that are not physiologically plausible. The Data Wrangler minimizes the impact of erroneous data by limiting all extreme values to a maximum and minimum value based upon the 95% confidence interval for each feature obtained from a pre-collected patient cohort. Following this preprocessing step, the Data Wrangler service then makes a stateless API call to the Sepsis Predictor service. An active patient's standardized data is then transmitted from the Data Wrangler service to the Sepsis Predictor service.

2) *Sepsis Predictor*: The core of the AIDEx framework consists of the prediction algorithm, that predicts the onset of sepsis early on. The Sepsis Predictor service runs the AISE algorithm[11]. When deployed, the algorithm can alert clinicians four to six hours before a patient meets the Sepsis-3 criterion. The output of the Sepsis Predictor Service is a sepsis risk score and the top three factors contributing to the sepsis risk score (see Fig. 2). The data returned by the Sepsis Predictor service along with all patient features are combined into JSON documents to be stored in the Results Store Data Warehouse.

The Data Wrangler service's final function is to provide a standardized interface for reading and writing data as JSON documents to the Results Store. Each JSON document contains timestamp and corresponding patient features, sepsis risk score, change in risk score over the last four hours, demographic information, and the three factors contributing most to the sepsis risk score.

3) *Results Store*: A data store is necessary to store the outcomes of the Sepsis Predictor. We developed a Results

Store service as a time-series data store to store data for the patients at each time-point and offer a standard access interface to the stored data. The AIDEx pipeline is designed to be scalable and capable of managing data streams from a large patient population. The patient data streams and the computed sepsis scores are transformed into a timeseries JSON document.

These documents are stored in MongoDB - a well-known NoSQL document store that is highly scalable and has been used in a variety of clinical and research applications. The database is accessed via a REST API that is built using an OSGI based declarative middleware called Bindaas [17-18]. Bindaas is an extensible big data middleware that lets the users create interoperable RESTful interfaces to various data sources. Other services in the AIDEx pipeline, including the user-interface access Results Store via this API.

4) *User Interface*: Graphical representations in a user interface supports the clinicians to interpret the outcomes of the AIDEx services. The clinical dashboard retrieves JSON documents generated by the Sepsis Predictor service and displays data in a graphical user interface (UI) for interpretation by clinical team members. As seen in Fig. 2, the UI includes a command center that gives a high-level overview of the ICU population and detailed view that presents detailed information including sepsis scores, clinical interpretations and vital signs. The default view seen in Fig. 2a demonstrates a population-level view of ICU patients. Each patient is represented by a single card, and the front of each card contains the patient's room number at the top, a sepsis score, a discharge readiness score, and

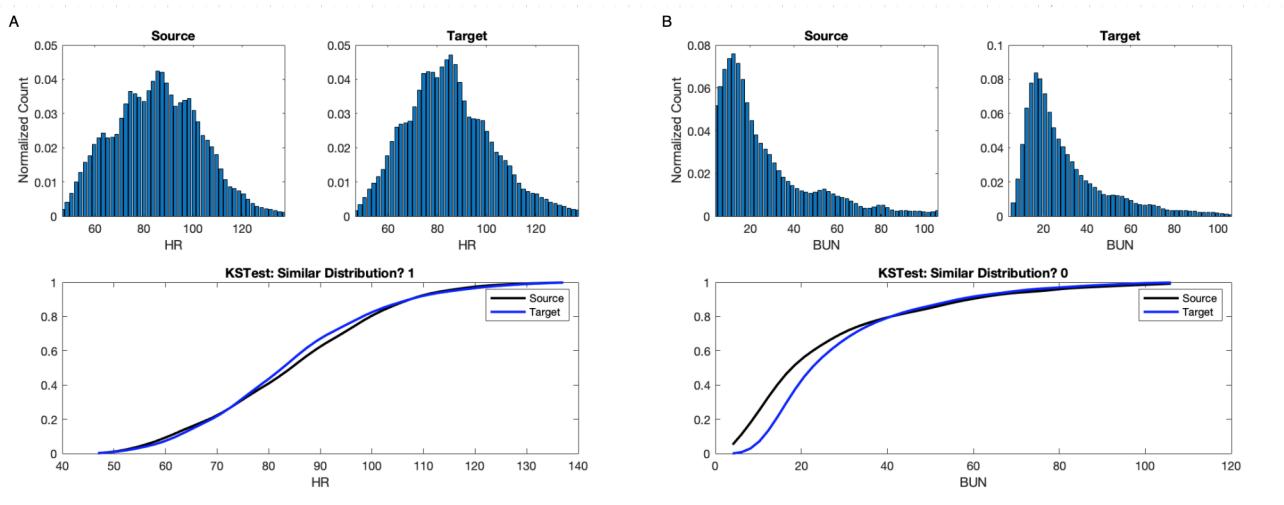


Fig. 3. An illustrative example of assessing the feature distribution over the source and the target ICU patient population. A: the upper panels show the histogram of the heart rate, recorded every 5 minutes, from source and target population. The lower panel shows the CDF for the same population cohorts. The KS test reveals that the heart rate real-values distribution between the source and target ICU population are similar. B: the upper panels represent the histogram of the BUN from the source and the target population. The lower shows the CDF for the same population cohorts. The KS test reveals that the BUN real-values distribution between the source and target ICU population are dissimilar.

finally a directional arrow with magnitude representing the acceleration (i.e. delta) of a patient's sepsis risk score over the last four hours.

The patient list is ranked according to the sepsis risk score with the most acute patients at the top of the list and a second UI see in Figure 2b is revealed. This patient centric view reveals the top three factors contributing to the risk score in addition to the vital sign trends for the patient over the last 24 hours. As previously described the Data Warehouse stores patient features in addition to the AISE Algorithm outputs in JSON files inside the Data Warehouse. This approach to data storage makes it simple for the UI to obtain patient data from the Data Warehouse for display in the user interface.

B. Security

We enforce security measures in development and deployment, to ensure the code satisfies the test requirements while the deployed AIDEx platform can be accessed only by the intended users. To prevent unauthorized access to sensitive data and APIs, we must ensure proper access policies and authentication mechanisms are in place. In AIDEx we configured secured access to the services to ensure proper authentication and authorization. We further configured firewall policies at the cloud instances and on-premise servers to ensure only the specified IP addresses can access the services, and only through the explicitly specified ports. In addition to securing against unauthorized accesses, we note that such protected network also minimizes the potential for denial of service attacks, by ignoring the service invocations from the unknown sources.

C. Testing Reliability and Quality Control for Model and Features

Deploying a machine learning model in a real-time environment poses challenges not common in offline experiments

[19]. Assessing production readiness level, monitoring and testing the system automatically are key consideration for a real-world ML software system. Real-time ML platforms are greatly dependent on the nature of data, more precisely the features. We have developed a series of automatic tests, specific to real-time ML platform, that run alongside traditional software engineering regression tests. In addition to the unit, integration, and system level tests to evaluate the functionality of the pipeline, AIDEx includes complement set of tests to assess, monitor and track the features and the data[19]. The underlying data distribution may change over time. ML platforms rely on the hidden representation of the feature, so changes in underlying data distribution will affect the model performance. This is a well-known concept of feature drift wherein a model built on stale data becomes inconsistent with newer data [19]. Feature drift affects Model reliability, which is crucial in the clinical environment so monitoring the data distribution and considering the feature drift is the important key in the real-time production ML system.

In AIDEx we have developed a set of automated tests to evaluate the difference in the distribution of the features. We performed the Kolmogorov-Smirnov test (KS test) to quantify the distance between the distribution of features from source and target cohorts. Target features are weekly sampled from our ICU population and the ML algorithm is developed based on source features. We store the p-value and the power of each test for our monitoring purposes and for updating the sepsis predictor algorithm in cases of feature drift. Further, we assess the difference between the predicted sepsis score for our ICU population with a 99% confidence interval every week.

Figure 3 illustrates two examples of assessing the difference between the distribution of heart rate and blood urea nitrogen (BUN) real-values from source and target ICU

patient population. Fig.3A illustrates the assessment of the difference of the heart rate distribution between source and target ICU patient population. The upper panels represent the heart rate distribution from the source and target population. The lower panel shows the cumulative distribution function (CDF) of the heart rate. The KS test reveals that for heart rate the two distribution are similar with the confidence level of 0.99. Fig.3B illustrates the assessment of the difference of the BUN real-values distribution between source and target ICU patient population. The upper panels represent the BUN distribution from the source and target population. The lower panel shows the cumulative distribution function (CDF) of the BUN. The KS test reveals that for the BUN the two distribution are dissimilar with the same confidence level.

III. CONCLUSION

Existing literature on the application of machine learning and deep learning techniques to healthcare applications are narrowly focused on novel algorithms, while in practice it takes coordinated efforts of many teams, including machine learning experts, software engineers, implementation scientists, and hospital IT teams, to bring such systems to the bedside. This is no trivial amount of work and every design and implementation choice makes a difference. AIDEx captures many of the critical elements necessary to take a well-tested and validated machine learning algorithm to production. It adopts a robust testing and quality control methodology that spans the software and the data. This has allowed us to tackle the major healthcare problem of sepsis. Early detection and treatment of sepsis is categorically one of the most important interventions that can be taken in a modern ICU. In this work, we have developed the AIDEx platform, and make it open-source available as a comprehensive way to detect, triage, and inform clinicians of a patient's risk for developing sepsis.

We are currently undertaking extensive external validations of the algorithm as well as integrations of AISE and gathering data for regulatory approval. AIDEx has allowed us to begin planning a multi-center clinical trial to examine the interventional use and utility of our sepsis prediction algorithms.

REFERENCES

- [1] M. Singer, C. S. Deutschman, C. W. Seymour, M. Shankar-Hari, D. Annane, M. Bauer, R. Bellomo, G. R. Bernard, J.-D. Chiche, C. M. Coopersmith, R. S. Hotchkiss, M. M. Levy, J. C. Marshall, G. S. Martin, S. M. Opal, G. D. Rubenfeld, T. V. D. Poll, J.-L. Vincent, and D. C. Angus, "The Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3)," *Jama*, vol. 315, no. 8, p. 801, 2016.
- [2] C. Rhee, R. Dantes, L. Epstein, D. J. Murphy, C. W. Seymour, T. J. Iwashyna, S. S. Kadri, D. C. Angus, R. L. Danner, A. E. Fiore, J. A. Jernigan, G. S. Martin, E. Septimus, D. K. Warren, A. Karcz, C. Chan, J. T. Menchaca, R. Wang, S. Gruber, and M. Klompas, "Incidence and Trends of Sepsis in US Hospitals Using Clinical vs Claims Data, 2009-2014," *Jama*, vol. 318, no. 13, p. 1241, Mar. 2017.
- [3] R. Ferrer, I. Martin-Loeches, G. Phillips, T. M. Osborn, S. Townsend, R. P. Dellinger, A. Artigas, C. Schorr, and M. M. Levy, "Empiric Antibiotic Treatment Reduces Mortality in Severe Sepsis and Septic Shock From the First Hour," *Critical Care Medicine*, vol. 42, no. 8, pp. 1749–1755, 2014.
- [4] A. Rhodes, G. Phillips, R. Beale, M. Cecconi, J. D. Chiche, D. D. Backer, J. Divatia, B. Du, L. Evans, R. Ferrer, M. Girardis, D. Koulenti, F. Machado, S. Q. Simpson, C. C. Tan, X. Wittebole, and M. Levy, "The Surviving Sepsis Campaign bundles and outcome: results from the International Multicentre Prevalence Study on Sepsis (the IMPRESS study)," *Intensive Care Medicine*, vol. 41, no. 9, pp. 1620–1628, 2015.
- [5] S. A. Sterling, W. R. Miller, J. Pryor, M. A. Puskarich, and A. E. Jones, "The Impact of Timing of Antibiotics on Outcomes in Severe Sepsis and Septic Shock," *Critical Care Medicine*, vol. 43, no. 9, pp. 1907–1915, 2015.
- [6] A. K. Venkatesh, U. Avula, H. Bartimus, J. Reif, M. J. Schmidt, and E. S. Powell, "Time to antibiotics for septic shock: evaluating a proposed performance measure," *The American Journal of Emergency Medicine*, vol. 31, no. 4, pp. 680–683, 2013.
- [7] S. P. Shashikumar, Q. Li, G. D. Clifford, and S. Nemati, "Multiscale network representation of physiological time series for early prediction of sepsis," *Physiological Measurement*, vol. 38, no. 12, pp. 2235–2248, 2017.
- [8] S. Horng, D. A. Sontag, Y. Halpern, Y. Jernite, N. I. Shapiro, and L. A. Nathanson, "Creating an automated trigger for sepsis clinical decision support at emergency department triage using machine learning," *Plos One*, vol. 12, no. 4, Jun. 2017.
- [9] T. Desautels, J. Calvert, J. Hoffman, M. Jay, Y. Kerem, L. Shieh, D. Shimabukuro, U. Chettipally, M. D. Feldman, C. Barton, D. J. Wales, and R. Das, "Prediction of Sepsis in the Intensive Care Unit With Minimal Electronic Health Record Data: A Machine Learning Approach," *JMIR Medical Informatics*, vol. 4, no. 3, 2016.
- [10] S. M. Brown, J. Jones, K. G. Kuttler, R. K. Keddington, T. L. Allen, and P. Haug, "Prospective evaluation of an automated method to identify patients with severe sepsis or septic shock in the emergency department," *BMC Emergency Medicine*, vol. 16, no. 1, 2016.
- [11] S. Nemati, A. Holder, F. Razmi, M. D. Stanley, G. D. Clifford, and T. G. Buchman, "An Interpretable Machine Learning Model for Accurate Prediction of Sepsis in the ICU," *Critical Care Medicine*, vol. 46, no. 4, pp. 547–553, 2018.
- [12] K. E. Henry, D. N. Hager, P. J. Pronovost, and S. Saria, "A targeted real-time early warning score (TREWScore) for septic shock," *Science Translational Medicine*, vol. 7, no. 299, May 2015.
- [13] A. E. Johnson, T. J. Pollard, L. Shen, L.-W. H. Lehman, M. Feng, M. Ghassemi, B. Moody, P. Szolovits, L. A. Celi, and R. G. Mark, "MIMIC-III, a freely accessible critical care database," *Scientific Data*, vol. 3, no. 1, 2016.
- [14] R. Postelnicu, S. M. Pastores, D. H. Chong, and L. Evans, "Sepsis early warning scoring systems: The ideal tool remains elusive!," *Journal of Critical Care*, vol. 52, pp. 251–253, 2019.
- [15] J. Norrie, "The challenge of implementing AI models in the ICU," *The Lancet Respiratory Medicine*, vol. 6, no. 12, pp. 886–888, 2018.
- [16] D. N. Sculley, G. D. Holt, D. Golovin, E. Davydov, T. Phillips, D. Ebner, V. Chaudhary, M. Young, J.-F. Crespo, D. Dennison, C. Cortes, D. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, "Hidden Technical Debt in Machine Learning Systems," *Advances in Neural Information Processing Systems* 28, pp. 2503–2511, 2015.
- [17] A. Sharma, A. Kazerouni, N. Saghar, P. Commean, L. Tarbox, and F. Prior, "Framework for Data Management and Visualization of The National Lung Screening Trial Pathology Images," *Pathology Informatics Summit*, pp. 13–16, 2014.
- [18] J. Saltz, A. Sharma, G. Iyer, E. Bremer, F. Wang, A. Jasniewski, T. Diprima, J. S. Almeida, Y. Gao, T. Zhao, M. Saltz, and T. Kure, "A Containerized Software System for Generation, Management, and Exploration of Features from Whole Slide Tissue Images," *Cancer Research*, vol. 77, no. 21, 2017.
- [19] E. Breck, S. Cai, E. Nielsen, M. Salib, and D. Sculley, "The ML test score: A rubric for ML production readiness and technical debt reduction," *2017 IEEE International Conference on Big Data (Big Data)*, 2017.