BLOOD STREAM INFECTION DETECTION USING DEEP LEARNING

MEDICAL AI: SLOW PACED REGULATION MEETS FAST PACED INNOVATION

An Undergraduate Thesis Portfolio
Presented to the Faculty of the
School of Engineering and Applied Science
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By

Matthew Pillari

August 11, 2022

EXECUTIVE SUMMARY

One out of every seven diagnoses are incorrect. Diagnostic errors kill approximately 60,000 people each year. Meanwhile, deep learning technology promises to eliminate a large portion of these errors. In the case of bloodstream infection detection, even the most skilled intensive care unit doctors miss about half of sepsis cases. This problem is the basis of the technical research paper. An ongoing research effort on automated diagnosis of sepsis is led by University of Virginia (UVA) Professor of Computer Science Rich Nguyen. This team tested a battery of time-series prediction deep learning models on bedside medical data in order to predict when a bloodstream infection might be taking hold. Beyond bloodstream infection detection, medical data has been used to train deep learning models to read X-rays with higher accuracy and speed than human radiologists. The STS research paper investigates diagnostic error, computer aided diagnosis (CAD), and the U.S. Food and Drug Administration (FDA) process will lead into an Ethics of Care based STS discussion of how a CAD-specific FDA approval pathway could improve the quality of care for all patients.

The bloodstream infection detection team is led by UVA Professor of Computer Science Rich Nguyen, accompanied by UVA Professor of Medicine Dr. Randal Moorman, UVA Doctor of Infectious Diseases Dr. Chris Moore, Director of UVA's Intensive Care Unit (and 2022 graduation speaker) Dr. Taison Bell, University of Pittsburgh Professor of Critical Care Medicine Dr. Gilles Clermont, assisted by UVA Graduate Students Jackson Brandberg and Joy Qiu, as well as my fellow Undergraduate Researchers Louisa Edwards and Zach Boner. The goal of this team was to assist doctors in predicting when an ICU patient might develop a bloodstream infection, just based on bedside vital signs. Currently, doctors look for abnormal, low or high, body temperature plus low blood oxygen and increased heart rate before suspicion is raised and a sepsis test is administered. However, these warning signs are extremely common in other

conditions, especially in the ICU environment. The aim of using deep learning for this problem is to extract more subtle patterns that may arise when considering bedside vital signs learning up to a positive sepsis test.

In order to create a deep learning model to predict sepsis, first a dataset must be compiled. The training data used in this case included UVA ICU data, University of Pittsburgh ICU data, and the publicly available MIMIC dataset. All of these datasets provided the bedside vitals of patients at regular intervals, as well as an indication of when a blood culture was taken, and the results of that test. A survey of deep learning literature led the team to a shortlist of potential deep learning models that operate on multivariate time-series data—data with multiple variables that change over time. After preliminary tests with these models, our model showed some predictive power on the dataset. These results show promise that a well designed algorithm could predict sepsis more accurately, and without input from the doctor.

The U.S. Food and Drug Administration (FDA) has consistently failed to accommodate the specific needs of software medical devices that use deep learning. The STS part of this portfolio investigates the practical causes as well as the social and moral implications of the FDA's poor support for software medical devices. Scientific studies and clinical trials of CAD devices have shown that they consistently outperform doctors and radiologists on a wide range of tasks. While the FDA has indicated that it intends to approve more deep learning powered medical devices, no tangible changes have been made to the 510(k) or De Novo pathways to acceptance, which often take a year or more to complete. In a fast moving and constantly changing field such as deep learning, these restrictions severely limit the performance of devices, even if the FDA approves them.

The Ethics of Care framework is a normative, virtue ethics-based framework that views "true care" and benevolence as virtues. Doctors feel the direct responsibility for the care of the patients, and are held to the highest standard by licensing boards and the threat of legal action. Meanwhile, the engineers who develop computer aided diagnostic software are far removed from the effects. As such, it is important for engineers working on CAD to adopt an ethics of care perspective. If an engineer does not truly care about the impact of the software they are creating, then it is possible for neglectful or outright incorrect code to cause a patient to receive incorrect medical care and die as a result. Minimizing this risk is partially the job of the FDA, but no one knows the product as well as the engineers who build it. These engineers must practice the same level of care and attention to the creation of the software as a doctor who is attending to a sick patient.

Computer aided diagnosis is inarguably the future of radiology and diagnostic medicine. However, roadblocks including the FDA process and community trust in CAD stand in the way of widespread impact.

TABLE OF CONTENTS