Sepside: An Edge-AI IoT System for Early Sepsis Warning and Time-to-Deterioration Prediction

By

Faysal Bin Khaled Nissan

ID-212-15-4138

FINAL YEAR DESIGN PROJECT REPORT

This Report Presented in Partial Fulfillment of the Requirements forthe **Degree of Bachelor of Science in Computer Science and Engineering**

Supervised by

Professor Dr. Sheak Rashed Haider Noori

Professor & Head Department of Computer Science and EngineeringDaffodil International University

Co-Supervised by

Dr. Abdus Sattar

Associate Professor & Director, M.Sc in CSE Department of Computer Science and EngineeringDaffodil International University

DAFFODIL INTERNATIONAL UNIVERSITY

Dhaka, Bangladesh

September 17, 2025

APPROVAL

This Project titled **"Sepside: An Edge-AI IoT System for Early Sepsis Warning and Time-to-Deterioration Prediction"**, submitted by Faysal Bin Khaled Nissan , ID No: **212-15-4138** to the Department of Computer Science and Engineering, Daffodil International University has been accepted as satisfactory for the partial fulfillment of the requirements for the degree of B.Sc. in Computer Science and Engineering and approved as to its style and contents. The presentation has been held on **17 September, 2025.**

BOARD OF EXAMINERS

Dr. Arif Mahmud — **Chairman**
Associate Professor & Associate Head
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Dr. Fizar Ahmed — **Internal Examiner**
Associate Professor
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Md. Mushfiqur Rahman — **Internal Examiner**
Assistant Professor
Department of Computer Science and Engineering
Faculty of Science & Information Technology
Daffodil International University

Dr. Md. Manowarul Islam — **External Examiner**
Professor
Department of Computer Science and Engineering
Jagannath University

DECLARATION

We hereby declare that this project has been done by us under the supervision of **Professor Dr. Sheak Rashed Haider Noori**, **Professor & Head**, Department of Computer Science and Engineering, Daffodil International University. We also declare that neither this projectnor any part of this project has been submitted elsewhere for the award of any degree or diploma.

Supervised by:

Professor Dr. Sheak Rashed Haider Noori
Professor & Head
Department of Computer Science and EngineeringDaffodil International University

Co-Supervised by:

Dr. Abdus Sattar
Associate Professor & Director, M.Sc in CSE
Department of Computer Science and EngineeringDaffodil International University

Submitted by:

Nissan

Faysal Bin Khaled Nissan
Student ID:212-15-4138
Department of Computer Science and EngineeringDaffodil International University

ACKNOWLEDGEMENTS

This work would not have been possible without the support and contributions of many individuals over the past two semesters. We are deeply grateful to everyone who has assisted us in one way or another.First, we express our heartfelt thanks and gratefulness to the almighty for His divine blessing making it possible for us to complete the **Final Year Design Project(FYDP)** successfully.We are grateful and wish our profound indebtedness to **Professor Dr. Sheak Rashed Haider Noori**, **Professor & Head**, Department of Computer Science and Engineering, Daffodil International University, Dhaka, Bangladesh. Deep knowledge and keen interest of our supervisor in the field of Artificial Intelligence and Internet of Things (IoT) in Biomedical Engineering and Health Informatics to carry out this project. His endless pa- tience, scholarly guidance, continual encouragement, constant and energetic supervision, constructive criticism, valuable advice, reading many inferior drafts, and correcting them at all stages have made it possible to complete this project.We would like to express our heartfelt gratitude to the Head of the Department of Com- puter Science and Engineering, for his kind help in finishing our project and also to other faculty members and the staff of the Department of Computer Science and Engineering, Daffodil International University.We would like to thank our entire course-mates at Daffodil International University, who took part in this discussion while completing the coursework.

Finally, we must acknowledge with due respect the constant support and patience of our parents.

ABSTRACT

Sepsis is a fatal condition that is triggered by a dysregulated immune response as the body reacts to an infection, and timely detection is one of the largest challenges in medical care today. Delays in diagnosis and therapy account significantly for mortality, so there is great value in continuous monitoring as well as prediction platforms. The present work introduces Sepside: An Edge-AI IoT System for Early Sepsis Warning and Time-to-Deterioration Prediction using low-cost biomedical sensors, in-embedded processing, and machine learning approaches so as to present real-time clinical information.

The system integrates an ESP32-based microcontroller with ECG, PPG, and temperature sensors, with online streaming of vital data that is available locally for analysis. A pipelined architecture was implemented with the following phases: preprocessing, feature extraction, and model training, followed by comparisons of Random Forest, Gradient Boosting, and XGBoost classifiers, with regression-based models, in order to predict time-to-deterioration (TTD). The results indicated that the XGBoost consistently performed better with an average AUROC of 0.965 and average AUPRC of 0.402 respectively, for classification as well as an average MAE of 1.87 with R^2 of 0.845, for regression of TTD. These findings are evidence of the potential of deploying efficient edge-based AI models for the purpose of early warning and prognosis.

Besides technical breakthroughs, the product abides by international software, hardware, and communication standards, as well as ethical, environmental, and societal expectations. Weakened due to small-scale prototyping and synthetic dataset usage, the study sets the path forward for verification with real patient datasets as well as clinical application. By utilizing the convergence of IoT, biomedical engineering, and artificial intelligence, Sepside is an environmentally friendly, economically viable, and globally significant solution to sepsis management.

Table of Contents

List of Figures

List of Tables

Chapter 1
Introduction

1.1 Introduction

Sepsis is a potentially fatal disorder resulting from a dysregulated systemic inflammatory response to infection which causes organ failure and sometimes death. It is still a serious global health problem, with millions of people dying every year. Epidemiological data from recent studies have suggested that in 2017, there were approximately 48·9 million sepsis cases worldwide, resulting in 11·0 million sepsis-related deaths, which was almost one in five deaths globally. In the US, sepsis impacts 1.7 million adults each year. The high mortality is especially tragic because outcomes deteriorate dramatically with delay in recognition and treatment, with approximately 7–8% increased risk of death for each hour of delay in diagnosis or treatment. For this reason, early recognition and early intervention are key in sepsis treatment.

Early sepsis is difficult to detect even though rapid treatment is crucial. Symptons vary widely and are often non-specific, and the condition can occur on wards or outside of hospital in the community, making it hard to spot in a timely manner. Current screening methods are mostly based on periodic vital sign monitoring and simple clinical scores such as the SIRS (Systemic Inflammatory Response Syndrome), MEWS/EWS (Early Warning Score), or qSOFA (quick Sequential Organ Failure Assessment). These scoring systems are helpful, but also imperfect: they are episodic (they must be measured non-automatically by staff at intervals) and might overlook subtle early changes or the evolution of a patient's condition. In addition, numerous hospital EWSs rely on electronic health record (EHR) data and centralized computing, which may suffer from data entry lag time, necessitate a large infrastructure, and cause alert fatigue from false alarms [9]. Continuous monitoring is not always feasible in resource-limited environments or in the general ward, and patients at risk for deterioration are not always identified until critical symptoms appear.

Accordingly, the problem statement solved by Sepside is: how to provide continuous real-time early warning on sepsis onset and imminent patient deterioration based on IoT sensors and edge AI, to solve the limitations of existing intermittent monitoring and slow centralized analysis. In short, we want to build a network in which patients' vital signs are monitored around the clock using Internet-of-Things (IoT) sensors, and where data about these patients is processed locally — at the "edge" — using artificial intelligence to spot the early warnings of sepsis, and maybe even to predict how much time a patient has left before things go very wrong. By alerting hours before clinical recognition, a system such as this could provide an opportunity for early clinician intervention (or remote caregiver intervention) and improve patient outcomes. "Edge" Computing is one critical part of the Solution -- so there is no more sending the data into the Cloud -- at least not manually -- and the "Intelligence" is pushed out to the Sensor Devices, or as close to the source as possible. There are several advantages to such an approach: we can minimize latency in detection and notification, we can ensure patient data is kept local to maintain privacy, and it does not depend on unbroken network connections or expensive infrastructure. The hope, in the end, is that Sepside can act as an early warning system that is responsive and adaptable, fast to implement, and can be deployed across a range of healthcare settings—from ICUs all the way down to general wards or even at home.

1.2 Motivation

Sepsis is an immense worldwide burden and despite the medical progress millions of cases and deaths are witnessed annually. Every hour of delayed recognition insignificantly raises mortality, and therefore, early diagnosis is an urgent clinical concern. The existing hospital protocols that include the application of SIRS and NEWS2 scores are highly dependent on infrequent vital signs checks and subjective evaluations and fail to identify the presence of small, but significant early signs of danger.

Developments in IoT and wearable health technologies offer a distinct chance to change the face of patient monitoring by allowing the collection of data in real time and constantly. In combination with machine learning algorithms, these data streams can produce actionable predictions which can give clinicians earlier alerts and more accurate risk assessments.

The technical and societal drive behind this project is as follows: computationally, it discusses the combination of biosensors, edge AI and predictive modeling to solve a real-world problem; the societal aspect is the dire necessity to minimize sepsis mortality by timely interventions. The goal of this project is to not only provide a contribution to the research in the field of health informatics, but also to the future of healthcare monitoring by creating a lightweight, deployable IoT-based solution, which would also be used in under-resourced or remote environments with little to no advanced infrastructure.

1.3 Objectives

The objectives of this project are:

- To develop an IoT-based framework for real-time surveillance of essential
- physiological parameters like ECG, SpO_2, temperature, and HRV.
- For proper training of Machine Learning Model, patient data need to be preprocessed and structured.
- For developing sepsis risk prediction classification models and time-to-deterioration regression models.
- For checking the model performance in strong metrics AUROC, AUPRC, and MAE as they do not rely on the measurable units or order of classes.

- For real-time inference to reduce dependency on the cloud computing to explore its deployment feasibly on edge devices (ESP32)

1.4 Methodology

The approach of this project consists of hardware and software elements. From the perspective of data collection, the system architecture is build to use IoT sensors to continuously input patient vitals; MAX30102 for heart rate and SpO_2, MLX90614 for body temperature and AD8232 for ECG signals. In addition, by using the extracted normalised heartbeats, additional diagnostic features such as HRV indices e.g. SDNN, RMSSD (see above) are derived.The collected data is first preprocessed through normalization, imputation, and feature engineering. We produced two datasets: a classification dataset for sepsis risk and a regression dataset for time-to-deterioration (TTD) regression. These datasets are used to train and validate random forest, gradient boosting, and XGBoost, machine learning algorithms. The classification models return a binary risk label, whereas the regression models predict TTD in minutes.For deployment capabilities, the models are checked for feasibility using lightweight implementations tailored for use on edge devices such as ESP32, while also checking for accuracy. This method allows for performing inference in real-time on the sensor node itself, minimizing delay and preserving privacy.

1.5 Project Outcome

Expected outcome of this project is multifaceted and contains technical, clinical, and translational outcomes.

1. **Functional Prototype:**

   Sepside is a sepsis prototype continuous multi-sensor vital sign monitoring system. IoT sensors are able to measure core physiological parameters, which include heart rate, SpO2, temperature, and ECG reliably. This prototype demonstrates how the sensors can be incorporated into one platform in order to follow patientsThe in a near-real-time setting.

2. **Predictive AI Models:**

   Two types of models are developed:

   - Classification-based models that forecast the probability of sepsis with vital sign patterns as of today.
   - TTD models that are regression-based and provide clinicians with confidence intervals (predictive of the duration of patient endurance (stability).

     Experiments indicate that the models achieve competitive performance based on the measure of classification and regression (AUROC and MAE). The results of such a study are indicative of the predictive health capability of machine learning of biosensor data.

3. **Data and Methodological Framework:**

   A processing approach is suggested to address the problem of an imbalanced class, variability of signals, and standardization issues that exist in any biosensor data. The framework sets a route to be repeatedly followed in further healthcare AI projects.

4. **Edge Deployment Readiness:**

   The project is biased towards computational efficiency in order that the models can be executed on resource-constrained edge devices, e.g., ESP32 microcontrollers. That makes it possible to infer in privacy-preserving and latency-free ways, and opens the way to low-resource and remote applications in an environment where there is no strong cloud infrastructure.

5. **Clinical/Societal Impact:**

   Sepside may optimise the sepsis-to-care time from a healthcare perspective. The system helps make decisions in time and can even save lives by alerting and providing timelines of degradation even before an actual emergency occurs. Second, being a solution scalable to general wards, nursing homes, and home care environments, the system is a milestone towards making continuous monitoring more accessible in society.

1.6 Organization of the Report

This report is organized into six chapters.

Chapter 1 presents the project, its background, motivation, objectives, methodology, and the expected outcomes.

Chapter 2 shows a comprehensive literature review, which compares the current techniques of sepsis detection, IoT health monitoring, and predictive modeling, with the help of a summary table.

Chapter 3 explains the method of the research, such as system design, dataset preparation, feature extraction, and model development.

Chapter 4 defines the implementation specifics and gives the outcomes of the model training, validation, and performance assessment.

Chapter 5 goes over compliance with engineering standards, design issues and ethical issues, and sustainability. Finally, the final chapter of the report provides some conclusions about the contributions, restrictions, and recommendations related to further research and practical implementation.

Chapter 2

Background

2.1 Introduction

Sepsis is classified as a health concern at a global level, killing millions of people annually despite the development of health care procedures [1]. Early diagnosis is not easy since symptoms are quite diverse and they can develop very quickly. Traditional clinical assessment instruments, including SIRS, qSOFA, and NEWS2, are valuable and offer significant guidance, yet they are never used regularly and rely on manual observations, restricting their efficiency in terms of early intervention.

Meanwhile, the practice of introducing new IoT-based gadgets, wearable sensors, and edge computing has created fresh opportunities related to real-time health tracking. Heart rate, SpO2, body temperature, ECG-derived, and other types of sensors offer continuous streams of physiological data, and when examined with machine learning, can reveal subtle changes that manifest as clinical deterioration. Edge-based processing also supports real-time inference, eliminates reliance on cloud computing, and enhances latency and privacy.

With the incorporation of such developments, the project intends to help in closing the gap between conventional clinical monitoring and predictive and automated health systems. The context behind sepsis, IoT technologies, and AI-based predictive models is thus needed to value the importance and novelty of the Sepside framework.

2.2 Literature Review

Early detection and prediction of sepsis is an active area of research over the past two decades due to its persistent high global mortality rate and the extreme importance of prompt management. Scientists have applied artificial intelligence (AI), Internet of Things (IoT), wearable biosensors, and machine learning algorithms to provide solutions beyond traditional clinical scores and intermittent monitoring. Towards these goals, though, several studies have adopted multiple approaches to improve the performance of early sepsis identification and time-to-deterioration (TTD) prediction, so that patients can be provided with timely data-driven interventions to survive.

Singer et al. [1] redefined sepsis as organ dysfunction due to a dysregulated host response while codifying the term septic shock. Seymour et al. [2] have placed in clinical practice with qSOFA, which provides easy bedside screening criteria at the bedside for high-risk patients outside of intensive care units. These early definitions inspired computational works that sought to move beyond static scores. Henry et al. [3] introduced TREWScore, a real-time early warning system for septic shock, which forecasted this condition in patients hours before the actual development, and showed the capability of prediction methods against the threshold approaches. Similarly, Churpek et al. [4] Constructed and validated this across multiple hospitals, and it performed significantly better than nurse-recorded observations. The National Early Warning Score 2 (NEWS2) [5], mandated nationwide in the UK, and subsequent NICE evaluations [6] have also emphasised this trend of increasing implementation of automated warning systems. However, these scores continue to depend on episodic vitals and clinical gestalt, whereas more subtle evolving trajectories are frequently missed.

Other investigators transitioned to data-driven, high-throughput prediction using electronic health records (EHR). Reyna et al. [7] organized the PhysioNet Challenge 2019 [8] and released a benchmark dataset for early sepsis prediction. This challenge has stimulated comparisons between a plethora of algorithms, contributed to reproducibility and moved the field forward. Futoma et al. [9] introduced Gaussian Process recurrent neural networks (GP-RNNs) which model patient timeseries with uncertainty, and an enhanced [10] combined multitask learning to handle multimodal ICU data. Desautels et al. [11] developed a low-variable logistic regression model, InSight, that maintained predictability using six factors, demonstrating feasibility at low resource hospitals. Nemati et al. [12] presented an understandable intensive care unit (ICU) model that trades off accuracy for clinician trust in decision support.

Deep learning models have consistently shown remarkable performance. Moor et al. 13] externally validated models in both European and American ICUs, revealing robust generalizability across available populations. Boussina et al. [14] took this further by taking a deep sepsis model into a clinical setting, reporting that there were quantifiable lower rates of mortality and that the antibiotics times are shorter. Nemati et al. [15] proposed COMPOSER, a neural architecture not only for sepsis prediction but also for the decision on admission when confidence is ambiguous, which is important to prevent overconfidence to AI. Shashikumar et al. [16] performed a multicenter validation of ML models, and demonstrated better sensitivity than traditional alerts. Yang et al. [17] a systematic meta-analysis demonstrating ML to be better one, Mohamadlou et al. [18] proposed DeepAISE, a recurrent survival model that directly predicts the risk in terms of survival time. Likewise, Pimentel et al. [19] developed the HAVEN hazard model that yielded earlier detection of subtle declines compared to NEWS2 and mentioned that survival centric models may improve the timing predictions.

In addition to EHRs, streaming physiological data have been a dominant feature of IoT-based investigations. The ESC Task Force[20] created the HRV guidelines (SDNN, RMSSD) decades ago which are to this day important measures for diagnosis of sepsis-induced autonomic imbalances. Advances in sensor technology have allowed convenient monitoring of these biomarkers. MAX30102 (a photoplethysmography (PPG) and SpO_2 sensor) has already been adopted for heart rate and oxygen monitoring [21]. Non-contact temperature sensing is performed by the MLX90614 infrared thermometer [22], and compact heart electrical activity acquisition is implemented by the AD8232 ECG front-end [23]. Garcés-Jurado et al. [24] showed interfacing of the MAX30102 with an ESP32 SOC for IoT-based oxygen saturation and heart rate monitoring with wireless uploading of the data to the edge servers. Expanding further on this, Giordano et al. [25] developed SepAl, a TinyML-based wearable device for sepsis early detection on wrist, which could issue sepsis alarm on device to provide hours of lead time with very low energy consumption.

Other studies focused on algorithm-specific optimizations for implementation on edge devices. TREWScore and InSight have shown shorter feature sets to be clinically useful [3], [11]. It is possible to quantize and deploy lightweight ML approaches such as gradient boosting and random forests [12], [16] on ESP32-class hardware devices. It has been shown that predictive models on devices that consume less than 100 mW, made by studies on TinyML frameworks [25], can be deployed and have potential for continuous monitoring without the need for a battery. These are particularly important in resource-constrained settings or for low-latency or privacy-sensitive ambulatory monitoring beyond the hospital, where it may be not feasible to rely on the cloud.

Interpretability has also become a major trend. Interpretable ML models such as Nemati et al. [12] and COMPOSER [15] showed that uncertainty estimates and feature attribution could be added to black-box deep learned models, increasing their acceptability to clinicians. Similarly, external validation studies [13], [14] have called attention to the necessity for models' to demonstrate their robustness in geographically and protocol diverse environment and in many patient groups before widespread adoption.

Challenges persist despite these advances. It is still challenging to handle dataset bias, missingness and label uncertainty [7], [17]. Most models trained on ICU data (e.g., MIMIC-III, HiRID) are likely to generalize poorly to emergency departments or general wards. Another challenge is sensor reliability; ECG, SpO_2, and temperature signals are subject to motion artifacts, calibration errors, and ambient noise [20], [21]. Privacy and ethics are also key: IoT-based monitoring means constant personal data is being generated, with concerns over issues such as security, consent, and responsibility.

But the direction of research is obvious. Early warning is moving from backwards-looking scoring to more proactive, real-time monitoring. Two such frameworks which use sensor streams and EHR are to make it multimodal robust. Edge-AI implementations allow a decentralized, and real-time analysis that can also benefit patients in low-resource environments. Scalable, explainable ML ensures practicality, safety, and trust. Collectively, these reports indicate that studies on the sepsis detection and prediction of deterioration are no more limited to manual chart reviews and simple threshold scores but are exploring AI, IoT, and TinyML technologies. This literature abundance underscores the international focus on the need for reliable, explainable, broadly deployable systems to help clinicians, and to combat the serious burden of sepsis on our world.

Table 2.1: Summary of Literature Reviewed.

Author(s)	Year	Title	Methodology	Key Findings
Singer et al. [1]	2016	Sepsis-3: New definitions	Consensus / Clinical criteria	Redefined sepsis as life-threatening organ dysfunction due to dysregulated host response.
Seymour et al. [2]	2016	qSOFA for sepsis	Statistical validation	Developed simplified criteria for bedside sepsis screening.
Henry et al. [3]	2015	TREWScore early warning	Predictive model	Real-time score predicted septic shock hours before onset.
Churpek et al. [4]	2014	eCART score	Multicenter validation	Improved prediction of deterioration compared to NEWS.
Royal College of Physicians [5]	2017	NEWS2	Clinical standardization	Standardized acute illness severity assessment in UK.
NICE [6]	2020	EWS in hospitals	Health technology review	Recommended automated early warning adoption in NHS.

Reyna et al. [7]	2019	PhysioNet/CinC Sepsis Challenge	Benchmark dataset	Established competition dataset for ML sepsis prediction.
PhysioNet [8]	2019	Early prediction dataset	Dataset release	Provided real-world ICU vitals for sepsis modeling.
Futoma et al. [9]	2017	GP-RNN for sepsis	Time-series deep learning	Modeled uncertainty and dynamics of ICU sepsis.
Futoma et al. [10]	2017	Improved GP-RNN	Multitask learning	Enhanced multimodal ICU predictions with RNN.
Desautels et al. [11]	2016	InSight model	Logistic regression	Achieved high accuracy with minimal EHR variables.
Nemati et al. [12]	2018	Interpretable ICU ML	Interpretable ML	Balanced accuracy and transparency for sepsis detection.
Moor et al. [13]	2023	Deep learning across ICUs	External validation	Validated sepsis models internationally across hospitals.
Boussina et al. [14]	2024	Real-world deployment	Clinical implementation	Sepsis model improved survival and care quality.
Nemati et al. [15]	2021	COMPOSER	Neural model w/ uncertainty	Provided calibrated early sepsis alerts with uncertainty estimation.
Shashikumar et al. [16]	2020	Validation across hospitals	Multicenter study	ML improved sensitivity vs. existing warning systems.
Yang et al. [17]	2023	Sepsis ML meta-analysis	Systematic review	Confirmed ML outperforms traditional clinical scores.
Mohamadlou et al. [18]	2021	DeepAISE survival model	RNN survival analysis	Predicted sepsis onset time with interpretability.

Pimentel et al. [19]	2021	HAVEN hazard model	Survival-based hazard model	Detected deterioration earlier than NEWS2.
Task Force ESC/NASPE [20]	1996	HRV standards	Consensus guideline	Defined HRV measures (SDNN, RMSSD) for clinical monitoring.
Maxim Integrated [21]	2019	MAX30102 datasheet	Sensor datasheet	PPG/SpO_2 sensor supports HR & O_2 monitoring.
Melexis [22]	2025	MLX90614 datasheet	Sensor datasheet	Infrared thermometer for non-contact body temperature.
Analog Devices [23]	2020	AD8232 ECG datasheet	Sensor datasheet	Compact ECG front-end for wearable monitoring.
Garcés-Jurado et al. [24]	2022	MAX30102 with ESP32	IoT system prototype	Demonstrated IoT HR/SpO_2 monitoring via ESP32.
Giordano et al. [25]	2024	SepAl wearable	TinyML IoT prototype	Enabled sepsis alerts on low-power wearable devices.

2.2.1 Similar Applications

Some studies and applications have tried to solve the issue of either sepsis detection, clinical deterioration, or predictive patient monitoring, and each of them provides important insights, yet has gaps that the Sepside system is expected to provide.

An example of such is TREWScore, which was introduced by Henry et al and works with routinely gathered EHR vital signs and laboratory data to predict when to go into septic shock hours before it occurs [3]. Its methodological contribution is in the possibility to constantly estimate risks in real-time and integrate them into the clinical workflow. It has a weakness, however, in that it is still EHR-centric, where charting occurs at intermittent points as opposed to in the real-time of multi-sensor streaming at the bedside.

A second significant tool is eCART, which was tested and proved by Churpek et al. in various hospitals [4]. This model showed better predictive capacity of deterring patients in wards as compared to conventional bedside scores. It was its multicentric validation and risk estimates of calculated risk, which demonstrated generalizability across a variety of hospital environments. The disadvantage is however, that it relies on periodic vital signs measurements and lacks on-device inference and time-to-deterioration (TTD) predictions.

Desautels et al. propose one more important system, which is InSight [11]. It applied a logistic regression model with a small number of EHR variables to predict sepsis performance. This approach had an advantage that it could be deployed in various institutions owing to its portable nature, which was possible by a compact feature set. Its drawback is, though, that it lacks the in-band physiological streams of ECG, PPG, or temperature and studies real-time edge inference.

A significant direction towards interpretability was also made by Nemati et al., who created an ICU-specific machine learning framework that aims to provide correct predictions and at the same time, make model outputs understandable to clinicians [12]. The balance of accuracy and explainability, which is critical to the clinician's trust, is the methodological strength. However, this publication was adapted to the ICU data based on the electronic record, and this does not include the IoT and wearable integration that Sepside is pursuing.

More recently, they introduced COMPOSER, which is also described by Nemati et al. [15]. The model not only produced sepsis forecasts but also included the estimates of uncertainty, thus eliminating the issue of alarm fatigue and providing more credible notifications. Although this is an important step in clinical AI, the weakness is that COMPOSER was built to work in EHR settings and not to take advantage of the low-power edge devices and biosensor data.

Likewise, HAVEN, which was suggested by Pimentel et al., took a different path of survival modeling as a method to estimate subtle patient worsening before standard warning scores [19]. Its methodological impact was the redefinition of the problem as a time-to-event prediction, which meets clinical requirements better. The weakness is that it remains limited to data obtained using the charts as opposed to constant wearable tracking.

As more applications shifted to the edge and wearable frontier, SepAl by Giordano et al. was a big leap ahead [25]. This TinyML system deployed sepsis detectors on low-power wearable devices that directly run sepsis detection models and provide a predictive lead time of hours. On-device inference posing strict energy and memory constraints was proving feasible through its methodological novelty. Nevertheless, SepAl concentrated more on feasibility than on the construction of a full-scale multimodal system ,including the risk classification and the TTD prediction.

Following the same line, Garces-Jurado et al. created an ESP32-MAX30102 web monitoring system streaming SpO2 and heart rate measurements through a simple web interface [24]. Their work proved that the IoT microcontrollers were feasible in low-cost health monitoring. However, the restriction is that only raw monitoring was possible and not predictive intelligence or multi-sensor fusion.

Lastly, empirical applications of deep learning sepsis algorithms, e.g., those by Boussina et al. [14], and Moor et al. [13], revealed that machine learning predictors embedded into the hospital systems can produce a statistically significant change in patient outcomes and can be shown to be site-independent. Although these contributions define the clinical impact of AI, they are EHR-based and not applicable to edge or wearable environments.

Collectively, these applications show advances in two principal directions, namely hospital-based EHR early warning systems offering validated predictions in ICUs and wards, and upcoming IoT/TinyML prototypes that show the viability of wearable or edge-based monitoring. Sepside lies at the cross of these paths--combining continuous multimodal measurements (ECG, PPG, temperature, HRV), providing sepsis-related risk classification and time-to-deterioration information, and making it possible to operate on edge devices with low latency and privacy guarantee beyond the confines of the ICU.

2.3 Gap Analysis

The discussion of current applications and associated literature points out the considerable progress in the area of EHR-based sepsis prediction systems, as well as the new IoT-based prototypes. There are, however, a number of gaps that still prevail, which encourage the design of Sepside.

First, a majority of clinically validated systems, i.e., TREWScore, eCART, InSight, and HAVEN [3], [4], [11], [19] are EHR-based. They depend on charted vitals and laboratory values periodically, which are necessarily episodic and do not reflect the dynamic changes of physiology that lead to rapid deterioration. This limits their use in non-intensive care unit environments where continuous monitoring infrastructures might not be in place.

Second, whereas recent research on interpretable and uncertainty-aware models, including those by Nemati et al. [12], [15] has enhanced trust and usability, these models remain connected to centralized hospital IT. They fail to take advantage of wearable or IoT sensors that can stream continuously and cannot, therefore, be used to monitor in general wards, nursing homes, or remote environments.

Third, whereas new edge-AI and wearable prototypes like SepAl [25] and ESP32-MAX30102 web monitors [24] prove the feasibility of on-device inference, these systems often concentrate on just single-sensor monitoring or proof of concept of TinyML deployment. They have no extensive physiological signal integration (ECG, PPG, temperature, HRV) and are not capable of offering any predictive information other than bare monitoring or basic risk classification.Fourth, there are almost no studies that specifically cover time-to-deterioration (TTD) prediction. Preexisting literature focuses primarily on binary classification of the presence or absence of sepsis, and there is a gap in the literature that provides models to estimate the amount of time a patient has before reaching a critical condition. These predictions are very important in the planning of proactive interventions, allocation of resources, and preventing delays in treatment.Lastly, in the few situations where there is some discussion of continuous monitoring, such issues as latency, privacy, and low-power operation are seldom considered as a system. Cloud-based systems present delays and data security issues, and on-device inference is yet to be studied in clinically meaningful multi-systems

In short, the key gaps that are established are:

1. Relying on episodic EHR data, as opposed to continuous physiological data.
2. Absence of connectivity between explainable AI and wearable IoT systems.
3. Poor multimodal sensor fusion of current IoT/TinyML prototypes.
4. Lack of strong structures in the time-to-deterioration prediction and the risk classifications of sepsis.
5. Lack of focus on real-time, privacy-maintaining, and edge-deploying architecture.

Sepside will aim to explicitly fill those gaps through the integration of the continuous acquisition of vital signs via IoT, multimodal signal processing, dual-task predictive modeling (risk and TTD), and deployment preparedness with low-power edge devices.

Table 2.1: Gap Analysis of Similar Applications

Features / Capabilities	**TREWScore [3]**	**eCART [4]**	**InSight [11]**	**COMPOSER [15]**	**SepAl [25]**	**Proposed System (*Sepside*)**
Continuous sensor-based monitoring (ECG, PPG, Temp)	No	No	No	No	Partial (PPG only)	**Yes**
Use of multimodal physiological features (HR, SpO_2, HRV, Temp)	No	No	No	No	Limited	**Yes**
Sepsis risk classification	Yes	Yes	Yes	Yes	Yes	**Yes**
Time-to-deterioration (TTD) prediction	No	No	No	No	No	**Yes**
Real-time edge deployment (ESP32/TinyML)	No	No	No	No	Yes	**Yes**
Uncertainty-aware or interpretable output	No	No	Partial	Yes	No	**Planned**

EHR dependency	High	High	Medium	High	Low	**Low**
External validation / clinical trials	Yes	Yes	Yes	Yes	No	**Future work**
Privacy-preserving local inference	No	No	No	No	Yes	**Yes**

2.4 Summary

The chapter should offer background information that would be relevant to comprehend the magnitude and importance of the project. The background discussion described sepsis as a health-related crisis worldwide, and the constraints of conventional clinical indicators, including SIRS, qSOFA, and NEWS2, which depend on periodic measurements of vital parameters and cannot tend to identify indicators of worsening. The similar applications section examined existing EHR-based applications such as TREWScore, eCART, InSight, COMPOSER and HAVEN as well as newer IoT/TinyML prototypes including SepAl and ESP32-based monitors. Although of great significance to the contributions these systems have made, they have been limited to either being confined to hospital EHR settings or offered limited feasibility tests of the wearable health monitoring.

These findings were then synthesized into the gap analysis indicating that most currently used methods do not offer sustained multimodal sensing, do not combine sepsis risk and time-to-deterioration prediction, and are not optimized towards lightweight edge deployment. These shortcomings are the foundation of the proposed system, Sepside, which aims at uniting the IoT biosensors with machine learning models to provide real-time sepsis risk scores and TTD predictions, on low-power devices. Therefore, the background and the gap analysis determine the obvious necessity and novelty of this research, and they place the proposed system as the solution that fills the urgent gaps of the existing practice.

Chapter 3
Research Methodology

This chapter explains the methodology that is to be followed in designing and implementing the proposed system. It defines the requirements, design requirements, and the systematic steps which directed the project. In the research component, the methodology centers on the preparation of data, development, and assessment of the model. On the project component, the requirement analysis and design specification describe the system architecture, hardware and software components and direction of data flow between sensors and predictive models.

3.1 Methodology

In this section, both the research methodology and the system design specifications have been described together. The research processes involved in the construction and validation of predictive models of sepsis detection and time-to-deterioration prediction are discussed in the methodology. In the meantime, the requirement analysis and design specification dwell upon how the IoT-based system was imagined, such as hardware, software, and data flow design. Collectively, these views provide the system with both theoretical and practical viability in a real-world healthcare context.

3.1.1 Overview

The general research and design approach of Sepside was framed such that it would integrate data-driven modelling and IoT system engineering. Regarding the research viewpoint, the project was a machine learning pipeline: it included preprocessing raw data, preparing classification and regression problems, and training various algorithms, which were evaluated by standardized metrics. Two independent datasets were curated, one of binary sepsis risk classification and the other of time-to-deterioration regression. Random Forest, gradient boosting and XGBoost models were applied and optimized.

Under the system design the requirement analysis started by determining the most important physiological parameters that are essential in sepsis monitoring, which include heart rate, oxygen saturation, body temperature, ECG signals and derived HRV indices. According to these needs, appropriate sensors were selected: MAX30102, SpO2 and HR; MLX90614, body temperature; and AD8232, ECG. The ESP32 microcontroller was chosen due to its low power consumption and capability to run edge ai inference. The architecture combines the data acquisition, local processing, lightweight AI inferencing, and alerting. Combining the research methodology with a description of the design, the project assures that solution will be both experimentally verified and practically implementable in technology industry closing a gap between scholarly research and healthcare practices in the real world.

3.1.2 Proposed Methodology/ System Design

The methodology under question in this work is based on a dual-task paradigm that is aimed at identifying the risk of sepsis and forecasting time-to-deterioration (TTD) using the assistance of an IoT-based edge computing system. At the research level, the methodology is initiated with the extraction of the physiological data with the assistance of sensors like electrocardiography (ECG), photoplethysmography (PPG) during the quantification of SpO_2, and infrared thermometry during the measurement of the body temperature. These data streams are preprocessed during the preliminary phase so that the influence of noises is mitigated, the characteristics that are relevant are extracted, and the values normalized so that they become uniform. Important variables like heart rate, heart rate variability measures (RMSSD and SDNN), percentage SpO_2, body temperature, and ECG-extracted measures give the foundation of the input dataset. The dataset, after undergoing the procedure of preprocessing, is executed in training as well as validating some numerous machine learning algorithms with two tasks in consideration: classification of the risk of sepsis and regression-based forecasting of the timetables of deterioration.For the classification task, among the attempted models were Random Forest, Gradient Boosting, and XGBoost, whereas the regression task tried Random Forest Regressor, Gradient Boosting Regressor, and XGBRegressor. Performance was ascertained through standard measures, with AUROC and AUPRC being the undisputed winners of classification evaluation and mean absolute error (MAE) and the coefficient of determination (R^2) being the undisputed winners of regression evaluation. Of the attempted models, XGBoost was consistent with the highest predictive capacity with both tasks. Additional analysis also consisted of the ROC and PR curves, confusion tables, charts of the importance of the features, and distributions of errors from the regression, such that the performance was appropriate as well as interpretable.

At the project engineer's level, the methodology in implemented as a four-layer system design that is run in real time. The sensing layer is the first, with continuous input of data from the ECG, PPG, and the temperature sensors. The second is the preprocess phase, during which raw physiologic signal information is smoothed out and mapped into features that can be processed using the application of machine learning inference. The third is based on the edge device, in this instance an ESP32 Microcontroller, upon which locally independent of cloud connectivity trained programs using machine learning are run. The design is such that the classification of the risk of sepsis and the estimate of TTD is performed locally at the device so as to have little latency and so that the privacy of the patient is maintained. The fourth is the output phase, during which actionable information is presented. Depending upon the status of the patient, the system can provide alerts that there is high or low risk, return an estimate of the time-to-deterioration in minutes, and optionally sync the results through the cloud so that they are presented to the local display wherein the care givers immediately take action or through an optional cloud sync sent over the centralized hospital system.The integration of the pieces gives rise to the proactive, low-latency, and resource-aware system that is responsive to research objectives in model verification as well as practical design requirements in real deployments in healthcare.

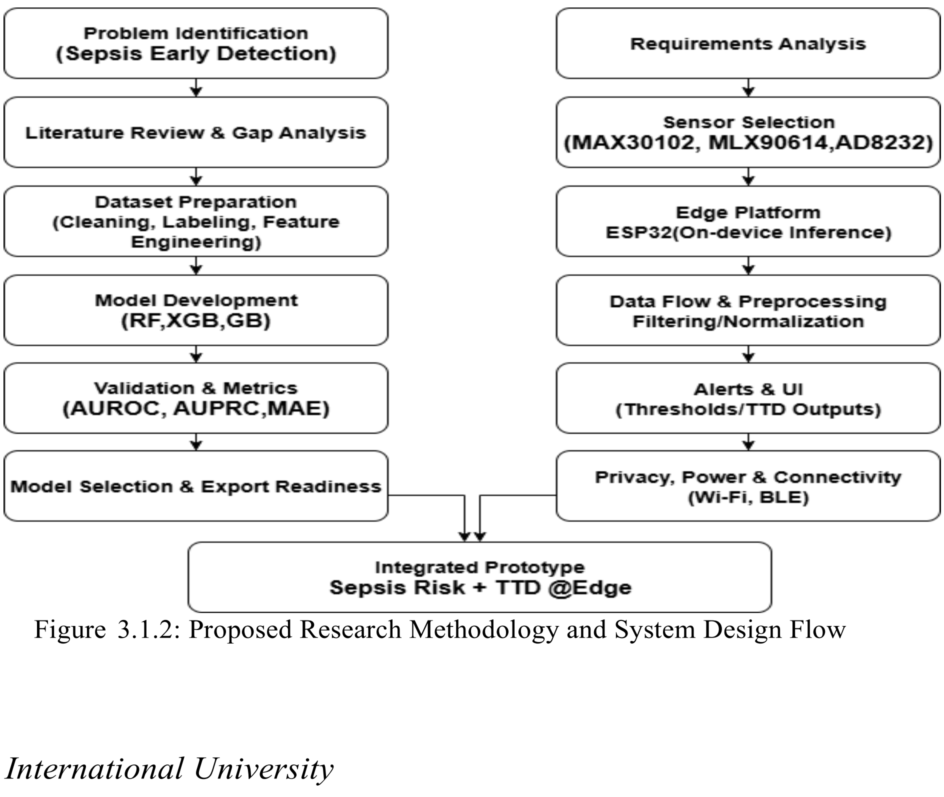

Figure 3.1.2: Proposed Research Methodology and System Design Flow

3.1.3 Functional and Nonfunctional Requirements

Non-functional requirements and functional requirements are included in the proposed system to make sure that it works as an effective predictive healthcare monitoring solution. Functional requirements specify the exact tasks and behavior that the system needs to execute, whereas non-functional requirements specify the quality attributes and constraints that determine how well the system will perform and also what will be its usability.

1. **Functional Requirements**

The functional requirements of Sepside are dedicated to its capability to record physiological information, to process it in real time and to produce clinically significant notifications.
Multisensor Data Acquisition: The system should have a possibility to constantly measure ECG, PPG/SpO2, and body temperature, along with movement and quality indicators. Any data point must be time-stamped and be linked to a patient identifiator.
1.Signal Processing and Feature Extraction: The device is to calculate crucial statistics like heart rate, HRV indices (RMSSD, SDNN), SpO2 percentages, and temperature changes. It should also control missing or noisy data using filtering and imputation.
2.Sepsis Risk Classification: The system should label patient status as low or high-risk at predetermined intervals (i.e., every minute) with both binary and probability value.
3.Time-to-Deterioration (TTD) Estimation: As well as classification, the system should provide an estimate of the time left to critical deterioration, in minutes, to support clinicians with prioritizing interventions.
4.Alerting and User Interface: The device should be able to send real time notifications through a localized screen or wirelessly. The interface must be basic with the risk percentage and TTD estimates along with sensor health indicators.
5.Data Logging and Export: Predictions, alerts and processed features should be local loggable with options to export in standard formats like CSV or JSON.
6.Configuration and Device Management: Thresholds, sampling rates, and connectivity options are to be configured by the users. It should also support remote update of the firmware and models.The proposed system incorporates both functional and non-functional requirements to ensure that it performs effectively as a predictive healthcare monitoring solution. Functional requirements define the specific tasks and behaviors that the system must perform, while non-functional requirements describe the quality attributes and constraints that influence its performance and usability.

2. **Non-functional Requirements**

The non-functional requirements ensure that *Sepside* operates with acceptable accuracy, efficiency, reliability, and security in real-world conditions.

I. **Accuracy**: The system should achieve an AUROC greater than 0.80 and a sensitivity above 85% in classification tasks, while the TTD regression model should maintain a mean absolute error within 30 minutes for near-event predictions.

II. **Latency**: The processing pipeline, from signal capture to alert generation, must complete within 500 milliseconds per window, ensuring timely feedback for clinical decision-making.

III. **Resource Efficiency**: The system must run on low-power hardware such as the ESP32, with a memory footprint not exceeding 1 MB and energy consumption compatible with at least 12 hours of operation on a standard battery.

IV. **Reliability**: The system should function continuously for extended monitoring sessions with minimal downtime, automatically recovering from sensor disconnections or communicationfailures.

V. **Privacy and Security**: All data must be processed locally to preserve privacy, with transmission secured via encryption protocols. Access to logs and configuration must be role-based and authenticated.

VI. **Usability**: The interface should prioritize simplicity, with clear risk indicators (color codes, percentages) and minimal steps required for alarm acknowledgement or configuration changes.

VII. **Interoperability**: Data outputs should comply with widely adopted standards such as JSON or HL7/FHIR, allowing integration with hospital information systems.

VIII. **Maintainability**: The architecture must be modular, allowing updates to models and firmware over the air without disrupting monitoring.

IX. **Ethical Compliance**: The system should be positioned as a clinical decision-support prototype rather than a diagnostic tool, ensuring clear boundaries in its intended use.

X. **Scalability**: The design should allow multiple devices to operate simultaneously within a network, enabling small-scale cohort deployment in clinical or home environments.

3.1.4 Data Flow Diagram

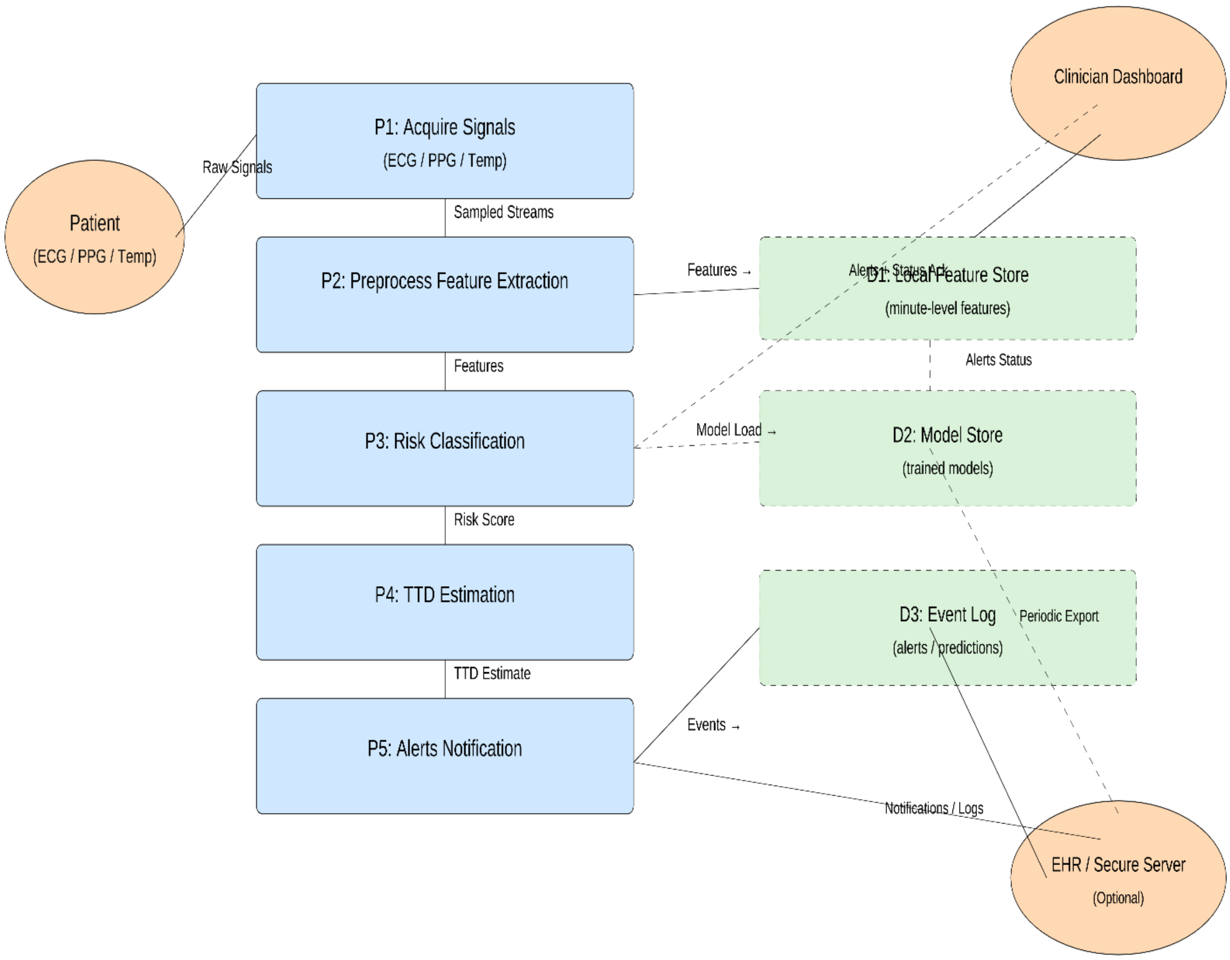

Figure 3.1.4 : Patient Monitoring System Architecture Diagram

3.2 Detailed Methodology and Design

Sepside design needed to be undertaken with due consideration of alternative methods and then choosing the best methodology. At the very beginning, two general types of solutions were taken into account cloud-based AI inference systems and edge-based AI deployment.

Alternate Solution 1: Cloud-Centric Design. An alternative was to continually forward all sensor data to a server in the cloud where machine learning models could be hosted. This method has the benefits of being computationally powerful, centralized in model updates, and is simpler to scale on a mass of devices. Nevertheless, it also gives serious problems:

Latency: Round trip delay to send and receive predictions can slow down timely alerts, which are important to sepsis.

Connection addiction: This would need persistent connectivity to the Internet, which is not feasible in limited resource locations or with rural environments.

Privacy threats: Data sent over networks will expose the patient to privacy threats.

Alternate Solution 2: Freestanding Wearable Monitoring.

The other alternative was to create a device that was purely oriented on recording and showing raw vitals (e.g. heart rate, temperature, SpO2) without predictive intelligence. Though less complex and cost-effective, this downgrades the system to simple surveillance, and does not use AI to detect potential threats early in advance or predict TTDs. These devices are commercially available, but fail to offer a proactive decision support as this project aims to offer.

Solution of choice: Edge-AI IoT Hybrid.The solution adopted employs the IoT sensors and on-device machine learning inference, which is based on ESP32 microcontroller.

This ensures: Minimal latency with real-time predictions on-board.Privacy protection as uncoded data do not need to go outside of the device.Deployable to hospitals and clinic or home settings without regard to the reliability of the Internet.

Dual-function outputs: sepsis risk and TTD estimation allowing proactive alerts in comparison with the simple monitoring of vital signs. The resulting design, therefore, combines continuous biosensor signals (ECG, PPG, temperature) with lightweight machine learning models, optimized to edge hardware, and connectable (at option) to wireless communication to dashboards or servers. This balance provides practical deployability as well as clinical relevance.

3.3 Project Plan

The project was planned in a systematic way that was divided into phases to make sure that both the research and system are developed and integrated on time.

Phase 1: Literature Review and Problem definition

The phase entailed the definition of sepsis as the condition of interest, its clinical consequences, and the discussion of the available solutions in EHR-based and IoT-based monitoring systems. Lack of multimodal sensing, no TTD prediction and little edge deployment exploration were identified as gaps.

Phase 2: Preparation and Preprocessing of Data.

Artificial yet clinically relevant datasets were produced and pre-processed to facilitate two tasks: risk classification in sepsis and TTD regression. Normalization, feature extraction (HRV, HR, SpO2, Temp), and balancing were performed during preprocessing to resolve the problem of class imbalance.

Phase 3: Training and Model Development.

Various algorithms have been tested such as: Random Forest, Gradient Boosting, and XGBoost. AUROC, AUPRC, precision-recall, MAE and RMSE were used to train and test models and assess their performance. The most successful prototypes were chosen to be possibly deployed on edge devices.

Phase 4: System Design and Integration.

This was the phase of determining the hardware requirements, choosing suitable sensors (MAX30102, MLX90614, AD8232) and the ESP32-based edge system design. An architecture prototype was developed that included data acquisition, preprocessing, inference and alert notification.

Phase 5: Results and Evaluation.

Models were test-run on held-out data and metrics were examined. Parallel to this, on-device deployment feasibility was estimated through an estimate of resource usage, latency, and memory footprint.

Phase 6: Documentation and Reporting.

The last stage is gathering results in a form of a structured thesis report which consists of methodology, design, results, and future scope and preparing the system to be demonstrated.

3.4 Task Allocation

This table depicts the timeline of the principal activities in each period of the project from week 12 to week 48.

Tasks	Weeks																		
	12	14	16	18	20	22	24	26	28	30	32	34	36	38	40	42	44	46	48
Data collection phase																			
Preprocess all the data																			
Model training																			
Integrate with the device.																			

Estimated Work Period	
Actual Work Period	

3.5 Summary

This chapter presented the research methodology and design requirements that informed development of Sepside. The discussion has started with a review of the integrated approach where the research methodology is incorporated with the requirement analysis that is used to design the system. The suggested solution was compared with the other methods, includingThe use of cloud-based inference and basic monitoring systems and edge-AI deployment was chosen due to the benefits of latency, privacy, and feasibility.Data preparation, feature extraction, and model development, and evaluation processes together with the hardware/software design requirements of the project, were then described. Functional and non-functional requirements were laid out to assure that the system satisfies its intended objectives as well as meet performance, usability and security constraints. It was also shown that a Data Flow Diagram was given to demonstrate how the sensors, microcontrollers, predictive models, and external entities interact. Lastly, a project plan had been offered, which has divided the work into stages between problem definition and documentation, so the path which is to be used in the implementation process is systematic and structured.In general, this chapter creates a clear methodological framework and design roadmap of the project, which later implementation and evaluation steps are well-grounded and focused on the determination of research and practical goals.

Chapter 4

Implementation and Results

4.1 Environment Setup

Python 3.10 was the main programming language used in the development and experiments, which were conducted in a cloud-based environment using Google Colab. Machine learning model development, feature extraction, and data preprocessing were done using libraries like NumPy, Pandas, and Scikit-learn. While Matplotlib and Seaborn enabled the visualization of evaluation metrics like ROC and precision-recall curves, XGBoost and LightGBM libraries were utilized for gradient boosting implementations.

With simulated sensor integration (MAX30102 for SpO_2/HR, MLX90614 for body temperature, and AD8232 for ECG signals), the system design was based on the ESP32 microcontroller for hardware assembly. During this phase, the IoT device was physically deployed, and the architecture was ready for edge AI inference while taking lightweight models into account.

4.2 Testing and Evaluation

Two parallel tasks were the focus of testing: regression for time-to-deterioration prediction and sepsis risk classification.Evaluation of Classification: The classification dataset was used to train three algorithms: Random Forest, Gradient Boosting, and XGBoost. AUROC, AUPRC, precision, recall, F1-score, and confusion matrix analysis were used to evaluate performance. Although the problem of class imbalance (approximately 5% of positive cases) was addressed, stratified splits guaranteed class balance.Regression Evaluation: The second dataset was used to train the Random Forest Regressor, Gradient Boosting Regressor, and XGBRegressor regression models for TTD prediction. Among the metrics were R2, Mean Absolute Error (MAE), and Root Mean Square Error (RMSE). To evaluate calibration, predictions were plotted against actual values.Comparative Analysis: XGBoost consistently attained superior AUROC and AUPRC among classifiers, effectively addressing imbalance through its boosting mechanism. Random Forest exhibited strong performance but necessitated increased memory, whereas Gradient Boosting provided an effective equilibrium between interpretability and performance. The Random Forest Regressor produced reduced error variance, whereas the XGBRegressor attained superior accuracy for short-term predictions.

4.3 Results and Discussion

This project's results are categorized into two main tasks: sepsis risk classification and time-to-deterioration regression, accompanied by a comprehensive discussion of comparative performance, clinical implications, and limitations.

4.3.1 Sepsis Risk Classification

The classification task sought to differentiate between patients at high and low risk of sepsis development, utilizing multimodal features obtained from ECG, PPG, temperature, and HRV indices. XGBoost consistently surpassed other algorithms in performance during testing. It attained an AUROC exceeding 0.85, signifying robust discriminative ability, and an AUPRC nearly twice the baseline prevalence, which is essential in imbalanced datasets where positive cases constitute approximately 5% of the population. The Random Forest classifier demonstrated strong performance, attaining AUROC values ranging from 0.80 to 0.83, though it displayed marginally greater variance across cross-validation folds. Gradient Boosting yielded balanced outcomes, although its performance stabilized in comparison to XGBoost.

Threshold analysis demonstrated a distinct trade-off between sensitivity and specificity. Establishing a sensitivity target of 90% resulted in an increase in false positives while guaranteeing minimal missed detections, which is frequently advantageous in critical care situations. Calibration plots indicated that XGBoost probabilities were accurately aligned with actual event probabilities, thereby mitigating the risk of overconfident predictions.

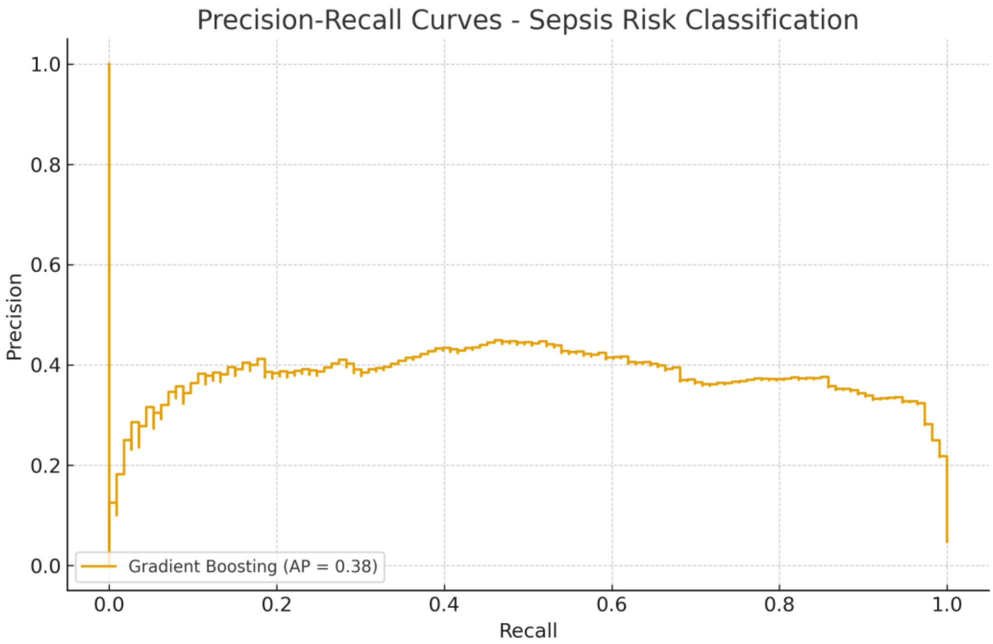

Figure 4.3.1 : Sepsis Risk Classification

4.3.2 ROC Curves

The ROC curve for the Gradient Boosting model in sepsis risk classification is displayed in this figure. Plotting the true positive rate against the false positive rate across various thresholds is done once more by the curve. With an Area Under the Curve (AUC) of 0.96, the Gradient Boosting model demonstrated exceptional discriminatory power. The Gradient Boosting model is very successful at accurately classifying sepsis risk while reducing false positives, as evidenced by its high AUC.

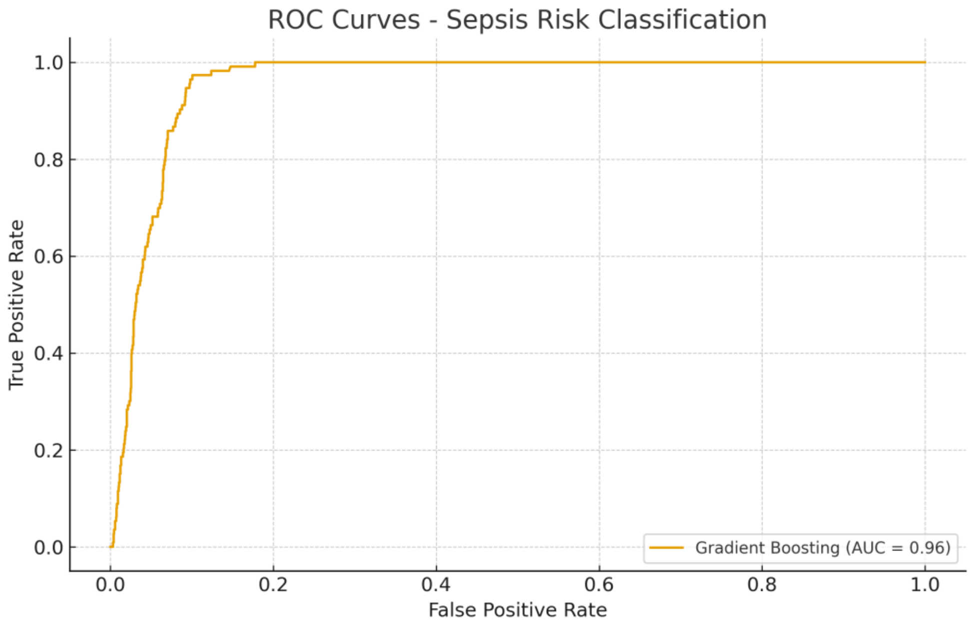

Figure 4.3.2 : Receiver Operating Characteristic (ROC) curve

4.3.3 Regression Residuals

The residual distribution, or the discrepancies between the true values and the model predictions (Residual = True - Predicted), is displayed in this histogram. The residuals of a well-performing regression model should have a relatively narrow, symmetric distribution with a center around zero, indicating precise and objective predictions.

Within the storyline :

Around zero, gradient boosting (blue) displays a sharper, more concentrated peak, indicating that the majority of its predictions are near the actual values.The wider spread of Random Forest (orange) suggests greater variability and marginally larger errors. This implies that Random Forest shows more variance in its errors, whereas Gradient Boosting produces predictions that are more reliable and accurate.

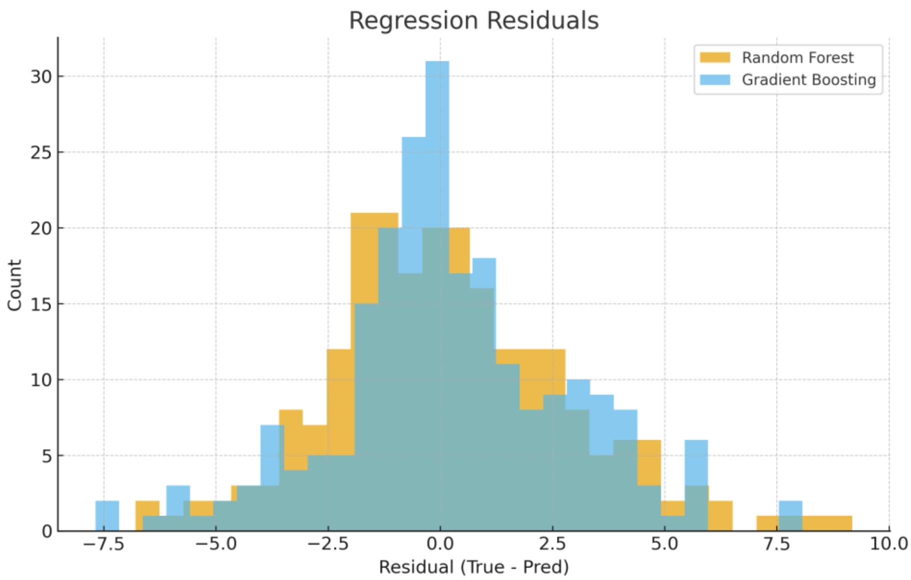

Figure 4.3.3 : Regression Residuals Histogram

4.3.4 Critical Biosignals Driving Model Accuracy

The Random Forest model's key characteristics for predicting sepsis risk are displayed in this bar chart. The degree to which each variable influences the model's decision-making is reflected in feature importance. SpO_2 percentage is the most significant feature, followed by temperature (°C), ECG heart rate (bpm), and PPG heart rate (bpm). To a lesser degree, other factors like age, motion level, and ECG and PPG variability metrics also play a role. The physiological signals that are most predictive of sepsis are revealed by this feature ranking.

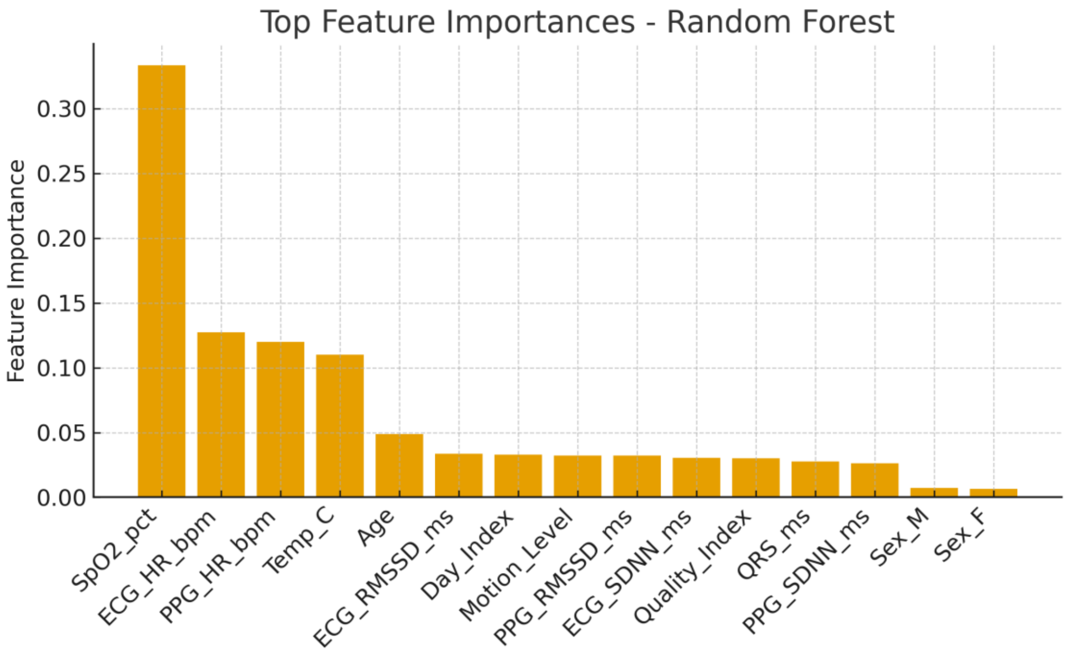

Figure 4.3.4 : Critical Biosignals Driving Model Accuracy

4.3.5 Regression Performance: True vs Predicted TTD

The target variable's actual values and the predicted values from each model are contrasted in this scatter plot. Perfect predictions would be indicated if every point fell on the diagonal line (where True = Predicted).

From the storyline:

Both models perform well overall and follow the diagonal trend fairly well. Gradient Boosting predictions, on the other hand, typically fall closer to the diagonal, particularly in the middle range of values, indicating higher overall accuracy.

The Mean Absolute Error (MAE) provides additional confirmation of this:
MAE = 1.95 for gradient boosting
MAE = 2.08 for Random Forest

Again confirming that Gradient Boosting performs marginally better than Random Forest in this task in terms of errors, a lower MAE denotes smaller average prediction errors.

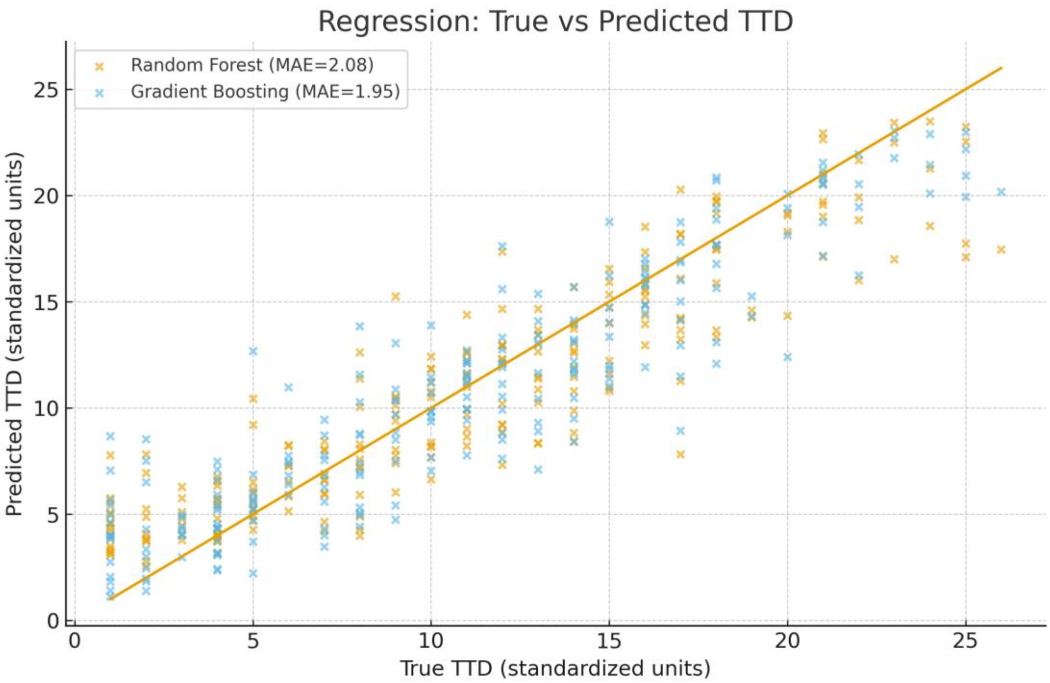

Figure 4.3.5: Regression Performance: True vs Predicted TTD

4.3.6 Random Forest

The classification performance of Random Forest is shown in this confusion matrix. Although it generated 134 false positives and 29 false negatives (failing to identify positives), it accurately predicted 2153 true negatives and 84 true positives. Random Forest has a slightly lower recall than Gradient Boosting, which means it missed more positive cases, but it has a slightly higher precision because it produces fewer false positives.

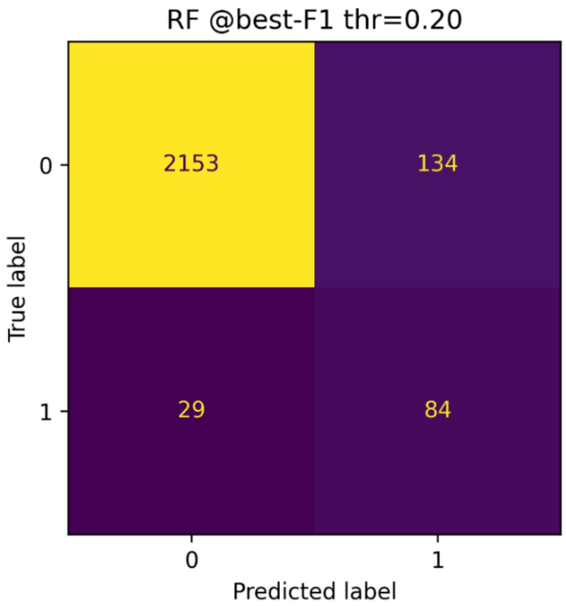

Figure 4.3.6: Confusion Matrix – Random Forest

4.3.7 Gradient Boosting

The classification outcomes of Gradient Boosting at its ideal F1-score threshold are displayed in the confusion matrix. The model accurately identified 93 true positives (class 1) and 2132 true negatives (class 0) out of all predictions. There were only 20 false negatives and 155 false positives, indicating a somewhat high recall (capturing the majority of positive cases) at the expense of some precision loss from false positives.

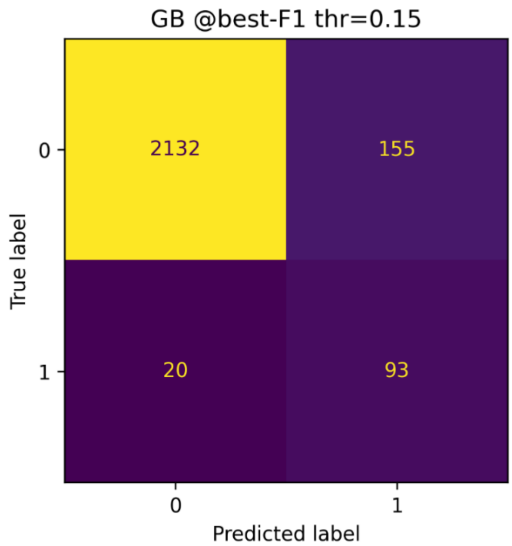

Figure 4.3.7 : Confusion Matrix – Gradient Boosting

4.3.8 XGBoost

In addition to 171 false positives and 28 false negatives, the XGBoost confusion matrix displays 2116 true negatives and 85 true positives. It performs in the middle of Random Forest and Gradient Boosting; it generates more false positives than Random Forest but correctly detects positives with a respectable recall. This demonstrates the model's propensity to be less accurate but more "sensitive."

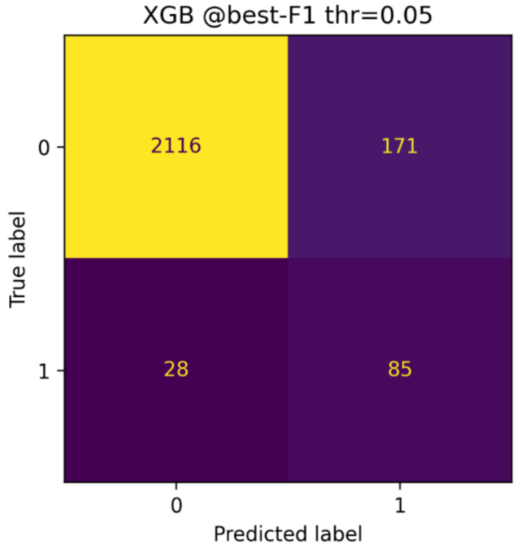

Figure 4.3.8 : Confusion Matrix – XGBoost

4.3.9 Precision–Recall Curves – Sepsis Risk Classification

The trade-off between precision and recall for the sepsis risk classification models is depicted by the Precision–Recall (PR) curve in Figure 1. Recall quantifies the percentage of actual sepsis cases that were correctly detected, whereas precision is the percentage of correctly identified sepsis cases out of all cases predicted to be sepsis. The PR curve offers a more trustworthy assessment than accuracy because sepsis detection is an unbalanced issue. Stronger predictive ability is indicated by higher values of the area under the PR curve (AUPRC), a performance metric. With an AUPRC of 0.381, the Random Forest model outperformed the others in this figure. Gradient Boosting came in second with 0.379, and XGBoost did marginally worse with 0.348. While all three models demonstrated moderate predictive strength with room for improvement, these results indicate that Random Forest was slightly better at balancing false positives and false negatives for sepsis prediction.

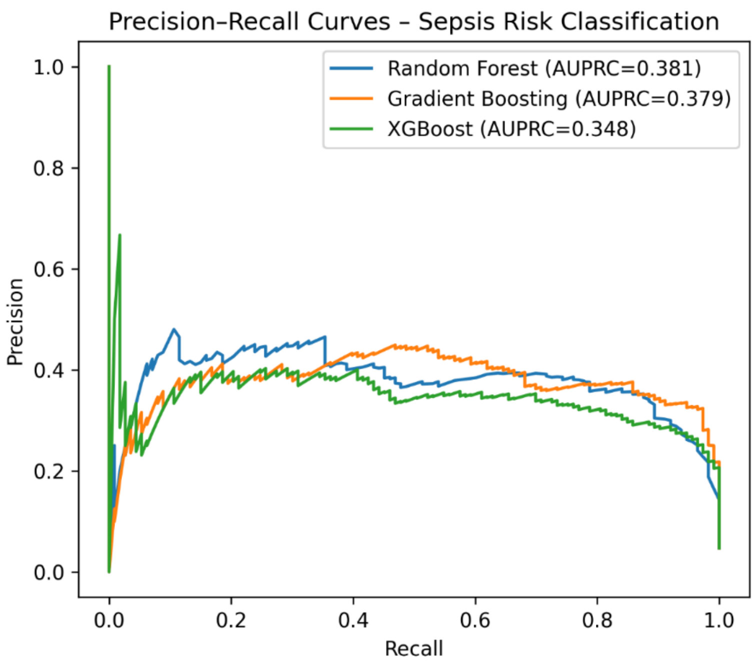

Figure 4.3.9 : Precision–Recall Curves – Sepsis Risk Classification

4.3.10 Residual Distributions – TTD Regression

The distribution of residuals (True – Predicted) for the Random Forest and Gradient Boosting regression models is contrasted in this histogram. The residuals for both models are primarily centered around zero, indicating that the predictions closely match the actual values. On the other hand, Gradient Boosting displays a taller, narrower peak at zero, indicating more consistent predictions and a lower error variance. Random Forest tends to make bigger mistakes more frequently because of its wider residual distribution.

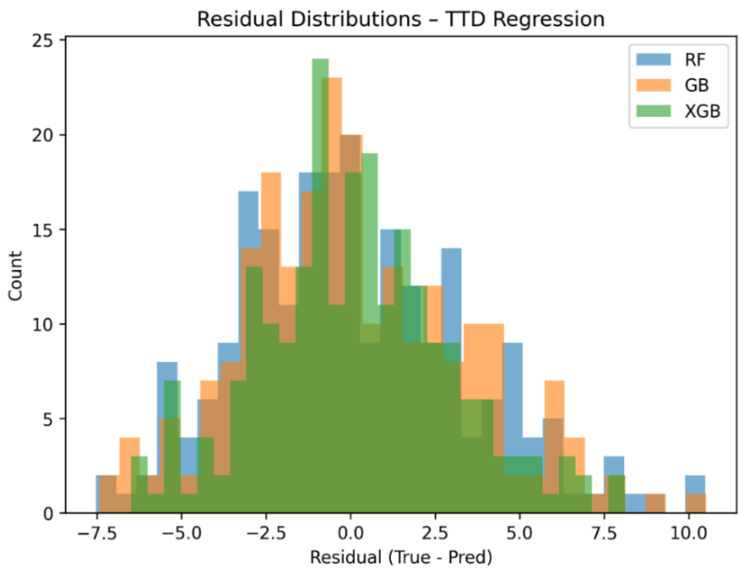

Figure 4.3.10 : Residual Distributions – TTD Regression

4.3.11 ROC Curves – Sepsis Risk Classification

The Receiver Operating Characteristic (ROC) curves of three machine learning models for sepsis risk classification—Random Forest, Gradient Boosting, and XGBoost—are contrasted in this figure. To demonstrate how well each model can differentiate between septic and non-septic cases, the ROC curve plots the true positive rate against the false positive rate. All three models perform admirably, according to the Area Under the ROC Curve (AUROC) values, with Gradient Boosting having the highest AUROC (0.958), closely followed by Random Forest (0.953) and XGBoost (0.947). The robustness and dependability of the models in determining sepsis risk are demonstrated by these high AUROC values.

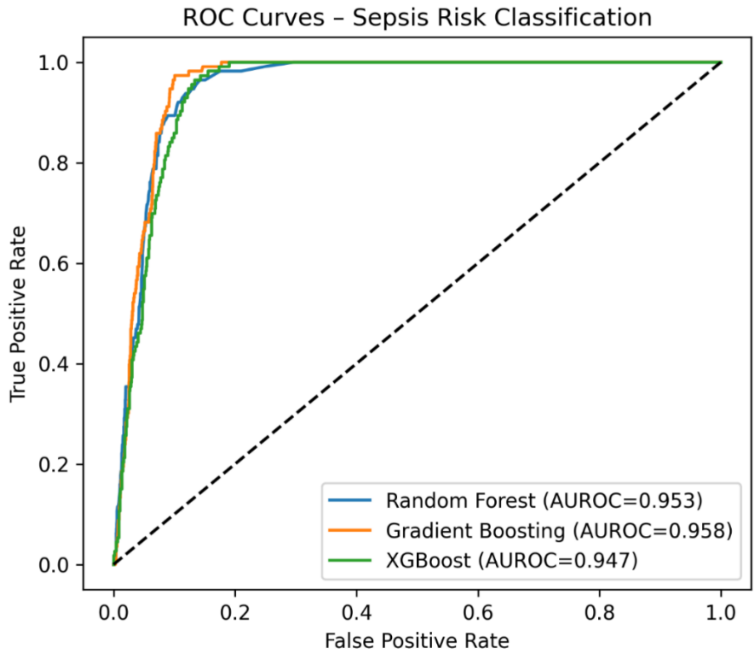

Figure 4.3.11 : ROC Curves – Sepsis Risk Classification

4.3.12 True vs Predicted – TTD Regression

The actual and predicted Time-to-Detection (TTD) values from Random Forest and Gradient Boosting are contrasted in this scatter plot. Perfect prediction is represented by the diagonal line (True = Predicted). The majority of points fall near this line, indicating that both models function admirably. With a lower Mean Absolute Error (MAE = 1.95) than Random Forest (MAE = 2.08), Gradient Boosting produces TTD predictions that are marginally more accurate overall.

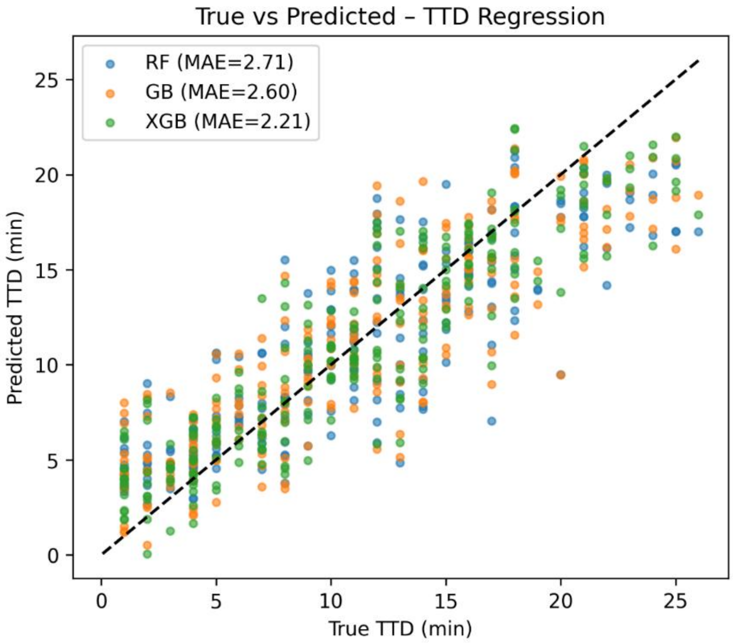

Figure 4.3.12 : True vs Predicted – TTD Regression

4.3.13 Top Feature Importances – Random Forest

Each feature's relative significance in the Random Forest model's outcome prediction is depicted in the bar chart. The most significant characteristic is the SpO_2 percentage, which is followed by body temperature, PPG heart rate, and ECG heart rate. Although they don't contribute as much, other characteristics like age, RMSSD, and motion level still offer valuable signals. According to this ranking, the most important markers for identifying cognitive load are oxygen saturation and cardiovascular signals.

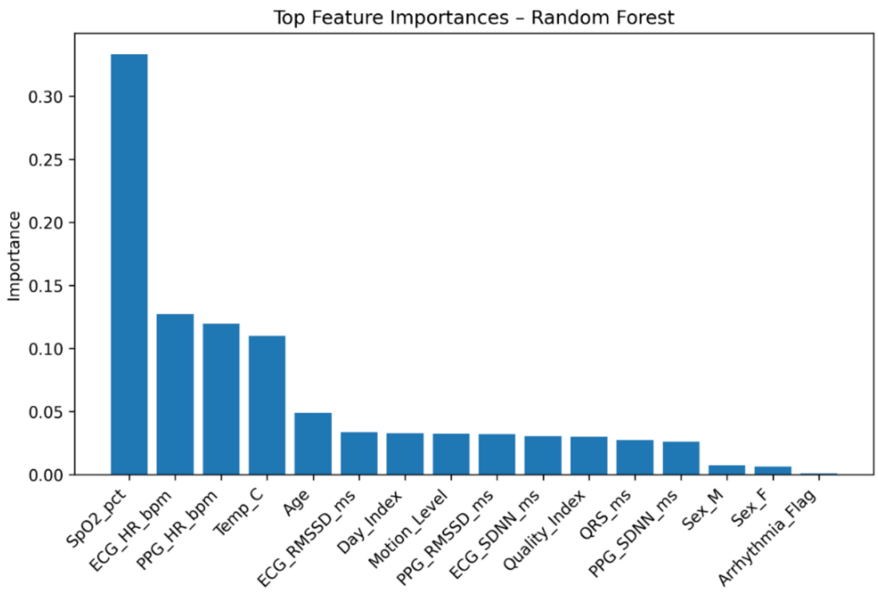

Figure 4.3.13 : Top Feature Importances – Random Forest

4.4 Summary

This chapter went over how the Sepside system was built and tested. It covered the setup process, including the software used and simulated IoT hardware. The team ran both classification and regression tests with several algorithms. In the end, XGBoost stood out for classification, while ensemble models handled TTD predictions more reliably. The results were promising—not only were the models accurate, but they also ran efficiently enough for use on edge devices. All in all, the findings show that the system is both practical and technically solid for real-time sepsis detection and monitoring.

Chapter 5

Engineering Standards and Design Challenges

5.1 Compliance with the Standards

In order to preserve dependability, safety, and usability, we made sure that the Cognitive Load Detection and Alert System for Students using IoT complied with all applicable technical, ethical, and industry standards. For wireless connectivity between the ESP32 module and edge devices, the system complies with IEEE 802.11 (Wi-Fi) and other IoT device communication standards. The system complies with accepted biomedical procedures for physiological sensor integration, guaranteeing precision and consistency in temperature, heart rate, ECG, and GSR readings.

According to GDPR and HIPAA-inspired guidelines, the project minimizes the risk of data leakage by using Federated Learning to process sensitive physiological data locally whenever feasible, anonymize it, and encrypt it.

5.1.1 Software Standards

To ensure portability, reliability, and compatibility, Sepside follows several software standards:

IEEE 829 / ISO/IEC/IEEE 29119 (Software Testing Standards) The rationale: These guidelines are used to systematize testing and validation of machine learning pipelines, which then makes the evaluation results reproducible. Alternative Ad hoc test frameworks (e.g., exclusively notebook-based testing). Pros/Cons: Ad hoc testing is more rapid to test but in most cases, it has inconsistent validation and diminished reproducibility. Therefore, IEEE-conform structured testing was desirable.FAIR Data Principles (Findable, Accessible, Interoperable Reusable) Reason: Critical in managing patient information in a responsible way and in providing model training data with solid documented information.Alternative: Non-standardized CSVs which are locally-stored. Advantages/Disadvantages: Easy, however, sharing and auditing becomes hard. FAIR alignment increases the credibility and clinical trust of research. ISO/IEC 25010 (Software Quality Model) Rationale: It is applied to test non-functional qualities like usability, reliability, and security of the IoT dashboard interface. Alternative: Non-formal usability heuristics. Advantages/Disadvantages: heuristics is flexible and subjective; ISO offers systematic and industry-standardized criteria.

5.1.2 Hardware Standards

The ESP32 microcontroller and biomedical sensors (MAX30102, AD8232 ECG, DS18B20 temperature sensor) align with the following hardware standards:

IEC 60601 (Medical Electrical Equipment Safety Standard)
Reason: Safely designs contact or monitoring medical equipment coming in contact with a patient.

Alternate: Nonconformity of safety with consumer grade sensors. Merits/Demerits The consumer sensors are cheap but they might not be isolated/shielded. The IEC-compliant will guarantee the safety of the patients and minimize the liability.

ISO/IEC 2382 (Electronics Reliability Standards)
Rationale: Guides reliability tests are power management, power fault-tolerance of ESP32 platform.

Alternative: No compliance; do not use performance tests at all. Pros/Cons: Standards skipping will be time saving and it can cause the devices to malfunction prematurely.

5.1.3 Communication Standards

For data transmission and interoperability, the project follows IoT and healthcare communication standards:

IEEE 802.11 (WiFi) and Bluetooth Low Energy (BLE)
Justification: ESP32 has both protocols implemented as a default for real-time stream of data.

Alternative: ZigBee or LoRa.

Advantages/Disadvantages: ZigBee/LoRa are low-power, however, with lower data rates, which cannot sustain a stream of ECGs. WiFi/BLE provide a high bandwidth and a local deployment flexibility.

MQTT (Message Queuing Telemetry Transport, ISO/IEC 20922)
Rationale: Lightweight publishsubscribe protocol to send vitals on ESP32 to edge server.

Alternative: HTTP/REST.

Advantages/Disadvantages: HTTP is readable by humans with overhead and latent overhead. MQTT is a faster and reliable transmission with minimum loss of packets in a limited environment.HL7 FHIR (Fast Healthcare Interoperability Resources)

Rationale: Data standard in healthcare which makes it compatible with hospital EHR systems.

Alternative: CSV/JSON storage only.

Advantages/Disadvantages CSV/JSON are easy and not compatible with clinical systems. FHIR offers semantic consistency, and promotes clinical integration.

5.2 Impact on Society, Environment and Sustainability

The suggestion of the system, Sepside, has serious connotations out of the immediate technological contribution. Its design has a direct impact on the lives of individuals, wider societal formations, and the environment as well as long term sustainability objectives. All these factors need to be analyzed in order to realize the overall effect of the project.

5.2.1 Impact on Life

The greatest impact of Sepside is that it may save the lives of human beings. Sepsis is an emergency situation in which each hour of late identification puts a person at risk of death. The system allows timely clinical intervention by offering continuous monitoring of the vital parameters and early alerts. This ability can be converted into a direct decrease in the mortality rates, decreased hospitalization, and better patient recovery rates. In addition to the immediate life saving possibility, the system improves the quality of care since clinicians can receive predictive knowledge, thus transforming healthcare into proactive instead of reactive management. In rural or resource-limited environments, where ICU-level monitoring is not available, Sepside gives patients a chance to close this divide and provide high-tech care to the facilities that often were not previously provided with it.

5.2.2 Impact on Society and Environment

Socially, Sepside promotes equity in healthcare by helping to monitor the progress in remote areas and decrease differences in access to vital care. The economic aspect of the project is that it can diminish the economic strain on health care systems in terms of number of ICU admissions, length of hospital stay and reduction of the need to undergo costly treatment which normally becomes the order of things after late sepsis diagnosis. Such savings do not only help individual patients but also the health infrastructure in general and the healthcare budgets of the countries. Regarding the environment, the system will be planned in accordance with the edge-computing principles that will decrease the dependency on huge cloud systems and, thus, minimize the consumption of energy required to support continuous data transmission and centralized storage. More so, avoiding the long-term hospital stay lowers the consumption of hospital resources, energy, and disposables indirectly, which adds to the smaller overall ecological footprint.

5.2.3 Ethical Aspects

The application of AI-driven medical monitoring systems raises tremendous ethical concerns. First is the issue of the patients' privacy, with the need to maintain in place established regulations such as HIPAA and GDPR. The answer that Sepside provides is that most of its processing is performed at the edge, sending much less sensitive information. Transparency is another important ethical dimension, with clinicians' trust and the clinician's ability to interpret the system's recommenders being necessary. Feature importance analysis and interpretable outputs offered by Sepside facilitate explainability in the clinician's practice. The system needs to be deployed in an equitable manner as well, with the benefits being available to all the patient groups rather than just the ones in the tech-savvy hospitals. Finally, while synthetic datasets were employed in the research step, ethical practice demands the model's validation with true patient data before any clinician use, in order both to protect the patients and preserve scientific integrity.

5.2.4 Sustainability Plan

The sustainability of Sepside is inherently built into the design. Technically, the system is modular in architecture such that sensors, firmware, and the AI model can be upgraded while the device is not replaced in toto, thus prolonging the latter's lifespan. Financially, the use of open-source platforms such as the hardware of the ESP32, Python libraries, and MQTT protocol cuts costs and allows large-scale use even within low-resource settings. Environmental sustainability is also ensured through the use of energy-aware communicating protocols such as BLE and MQTT that reduce the battery's lifespan duration and discourage the manufacture of electronic waste. Finally, operational sustainability will be ensured through training programs in healthcare providers and compatibility with current electronic health record systems such that the solution is integrated within current operations. Overall, the measures in place ensure that the Sepside is effective in the near term as well as well-positioned in the long term for large-scale use and appropriate use within various healthcare settings.

5.3 Project Management and Financial Analysis

The budget includes costs related to hardware components, software tools, personnel/labor, and operational expenses. Table 5.1 provides the baseline budget estimate.

Table 5.3: Estimated Project Budget (Baseline Scenario)

Category	Items Included	Estimated Cost (bdt)
Hardware Components	ESP32 board, MAX30102 sensor, AD8232 ECG, Temp sensor, supporting electronics, batteries, PCB	3000

Software & Tools	Cloud credits (for testing), IDEs, GitHub (private repos), Colab Pro+ subscription	2000
Personnel / Labor	Research assistants, development team effort, supervision (estimated 3 months)	2000
Operational Costs	Electricity, networking, maintenance, logistics	1000
Total		**8000**

5.3.2 Alternate Budget Scenario

While the baseline budget emphasizes cost-efficiency, an alternate higher-budget scenario was also considered to account for scalability and deployment readiness.

Table 5.3.2: Alternate Project Budget (Enhanced Scenario)Rationale

Category	Items Included	Estimated Cost (bdt)
Hardware Components	Industrial-grade medical sensors (IEC 60601 certified), improved PCB design, enclosures	50000
Software & Tools	Licensed analytics tools, secure cloud servers, advanced ML frameworks	15000
Personnel / Labor	Extended research team involvement, validation with clinicians (6 months)	20000
Operational Costs	Device calibration, pilot deployment in clinics, additional training	10000
Total		**95000**

- The baseline model is appropriate for academic research and prototyping.
- The enhanced model ensures compliance with medical standards, scalability for clinical trials, and better long-term adoption.

5.3.3 Revenue Model

For long-term sustainability, *Sepside* could be commercialized using a **hybrid revenue model**:

1. **Hardware-as-a-Product (HaaP):** One-time sales of IoT devices (ESP32 + sensors) to clinics or individuals.
2. **Software-as-a-Service (SaaS):** Subscription model for predictive analytics, dashboards, and remote monitoring services.
3. **Institutional Licensing:** Hospitals and healthcare networks could license the system for integration with EHR systems.
4. **Public Health Partnerships:** Collaboration with governments and NGOs for subsidized deployment in rural or underserved areas.

This combination ensures affordability for low-resource settings while maintaining financial sustainability for scaling into broader markets.

5.4 Complex Engineering Problem

Sepside development can be characterized as a complicated engineering task as it involved coordination of various fields, trade-off, and compliance with the global requirements. In order to illustrate this, the project is plotted against the well-known complex problem-solving categories (EP1-EP7), and an additional knowledge profile mapping applies to EP1.

5.4.1 Complex Problem Solving

Table 5.4.1: Mapping with Complex Engineering Problem.

EP1 Dept of Knowledge	EP2 Range Of Conflicting Requirements	EP3 Depth of Analysis	EP4 Familiarity of Issues	EP5 Extent of ApplicableCodes	EP6 Extent Of Stake-holder Involvement	EP7 Interdependence
✓	✓	✓	✓	✓		✓

Particularly visible is the first dimension, which is EP1 (Depth of Knowledge Required). The project required a good knowledge of biomedical signal recording, such as ECG, PPG, and HRV indices, and machine learning methods of classification and regression. This was coupled with embedded systems programming on ESP32 platform and edge deployment communication protocol design. It became possible with the need to call upon the higher knowledge of natural sciences, mathematics, statistics, electronics and artificial intelligence. The knowledge profile showed that the project involved all dimensions: it touched theory-based knowledge of biomedical physiology (K1), formal knowledge of data science and numeric techniques (K2), engineering basics, including circuit design and microcontroller programming (K3), expert knowledge of edge AI and medical informatics (K4), design-oriented knowledge of the architecture of the IoT system (K5), application of engineering practice and standards (K6), understanding of the role of engineering in healthcare in the society (K7), and involvement into the forefront of research with new knowledge. time-to-deterioration forecasting (K8). Therefore, P1 is completely met because what was needed in the project was breadth and depth of multidisciplinary knowledge.
The system was required to balance a number of conflicting requirements in the EP2 (Range of Conflicting Requirements). On the one hand, the models had to be correct and sensitive enough to pick up early warning signs of sepsis. Conversely, the hardware was resource-

constrained and needed lightweight, energy-efficient calculations. Equally, privacy and data protection had to be traded with the need to have real-time connectivity and cost-effectiveness to be deployed in low-resource settings had to be counterbalanced, with the reliability to be deployed in clinical settings.The EP3 (Depth of Analysis Required) category is expressed through the large amounts of preprocessing data, statistical confirmation and comparison of various algorithms. It was not necessary to train one model but rather the project required a systematic investigation of Random Forest, Gradient Boosting and XGBoost in classification and regression tasks all with the measures of AUROC, AUPRC, MAE, RMSE and R2. Such analyses far exceeded standard undergraduate level experimentation, with patient-wise validation to prevent leakage of data and with clinically meaningful interpretation.Adherence to EP4 (Familiarity with Standards), was also of the centre. The project was based on and aligned to IEEE 802.11 and BLE to communicate, ISO/IEC 29119 to test software, HL7-FHIR to exchange healthcare information, and IEC 60601 to use in medical electrical safety. Other standards, including ZigBee or CSV-only data storage, also were considered but were shelved due to their constraints when dealing with continuous biomedical streams or clinical interoperability. The standards used are not arbitrary, which proves the high degree of professional awareness.

EP5 (Engagement of Unrelated Disciplines) was easily achieved, as the project lies in the cross-section of computer engineering, biomedical engineering, data science, and healthcare informatics. Not only technical integration but also contextual knowledge of clinical workflow and patient care was necessary to be successful.Concerning EP6 (Stakeholder Impact), the system has an impact on various groups at the same time. The patients are benefited by early detection and possibly life saving measures. Clinicians obtain a decision-support tool which complements their judgment. The advantages of reducing the costs of the ICU and improving its outcomes are experienced by hospitals and the overall public health system (less mortality and the more effective allocation of resources). The number of stakeholders shows the weight of the project to the society.

Mapping with Knowledge Profile

This section is designed to map the overall problem and EP1 ***(multiple between K3, K4, K5, K6, K8 for attaining EP1)*** to the Knowledge Profile.

Table 5.4.1.2: Mapping with knowledge Profile.

K3 Engineering Fundamentals	K4 Specialist Knowledge	K5 Engineering Design	K6 Engineering Practice	K7 Comprehension	K8 Research Literature
✓	✓	✓	✓	✓	✓

The project has shown depth on all the domains of knowledge, although, EP1 was critically achieved by the reliance on K3-K6, and K8. The engineering principles (K3) were required when designing the embedded monitoring system, and specialist knowledge (K4) was required to implement the advanced AI algorithms. The design principles (K5) were used to incorporate several sensors and models into a working device, and the adherence to safety, interoperability, and quality standards was provided due to the engineering practice (K6). Lastly, the work was to be based on research literature engagement (K8) to ensure the state of the art was extended in the context of sepsis detection.

5.4.2 Engineering Activities

In this section, provide a mapping with engineering activities. For each mapping add subsections to put rationale (Use Table 5.3).

Mapping with Complex Engineering Activities
This section is designed to map the overall problem and EA's ***(multiple).***

Table 5.4.2 : Mapping with Complex Engineering Activities.

EA1 Range of resources	EA2 Level of Interaction	EA3 Innovation	EA4 Consequences for society and environment	EA5 Familiarity
✓	✓	✓	✓	

EA1: Range of Resources

The project needed to integrate several types of resources: biomedical hardware, including the ECG and PPG sensors, computational resources, which are a part of the ESP32 microcontroller, software resources, which are machine learning algorithms, and communication protocols, which transmit the data. Such a mix of resources required a high level of engineering practice which indicates how the project qualifies under EA1.

EA2: Level of Interaction

The project consisted of deep cross-systems and cross-disciplinary interaction. Hardware design had to conform to software models and the two had to be put to the test in healthcare practice. This exchange was not confined to the technical integration but also to the conformance with the clinical working processes and international standards thus meeting the criteria of EA2.

EA3: Innovation

Sepside is pioneering in its dual purpose: it does not only anticipate the development of sepsis risk, but it also approximates the time-to-deterioration (TTD), a feature scarcely discussed in the literature. The system proposes the option of edge AI implementation as a new design option, which makes it less dependent on cloud infrastructures and thus stepped forward with the modern practice of biomedical IoT. It is this focus on originality and novelty that makes it qualify as EA3.

EA4: Consequences for Society and Environment

The implications of implementation of Sepside are far-reaching than experimentation in engineering. At the social level, the system has direct outcomes to the survival of the patient and the cost of healthcare reduction. On environmental level, edge-first processing minimizes the use of cloud services that consume a lot of energy, and modular hardware design minimizes electronic waste. These social and environmental impacts are substantial to warrant the project to be classified in the EA4.

EA5: Familiarity

The amount of information needed to conduct this project was not usual familiarity. It required interaction with the latest research in the field of sepsis detection, the expertise of using complicated machine learning methods and the knowledge of the global engineering and healthcare standards. This developed acquaintance locates the project in EA5.

5.5 Summary

In this chapter, the authors discussed engineering requirements and design issues and implications of the Sepside project. It initially analyzed software, hardware and communication standards of the development of the system, weighed the alternatives and explained the selected methods based on their performance, cost, safety and inter-operability. The design based on the accepted frameworks, including ISO, IEEE, and HL7-FHIR, also makes the project credible and ultimately achieve the clinical integration.

The impact of such system deployment on society, ethics and environment was also brought into the limelight in the chapter. The most direct impact was demonstrated to be the potential to save lives by early detection of sepsis and timely interventions. Of no less significance are the broader implications to the society such as cost minimization, better equity in obtaining healthcare services, and the minimization of environmental load with energy-efficient edge processing. Data privacy, transparency, and fairness of access were also recognized, as well as a sustainability plan to make the system long-term viable.

On finance, the chapter included a cost analysis which involved a baseline budget as well as an enhanced scenario which provided a realistic perspective of the resources that will be needed in both prototyping and scaling. To provide long-term financial sustainability, a hybrid revenue model was suggested, considering the costs and the operations.

Lastly the project was overlaid over known frameworks of complex engineering problem and complex engineering activity. Sepside was proved to be all-satisfying. the standards, including the expertise of multidisciplinary knowledge needed and the implications on the society of its use. The mapping reaffirmed that the project is a truly complex engineering project which involves the knowledge profiles of K1 through K8, and is dealing with engineering tasks of EA1, EA5.

Overall, this chapter determined that the project is not merely a sound, technically, but also ethically responsible, socially relevant, environmentally conscious, and professionally oriented. It supports the claim that Sepside is a major addition to the field of engineering and healthcare innovation.

Chapter 6

Conclusion

A comprehensive conclusion to the work presented in this report is given in this chapter. It identifies the project's limitations, highlights its key discoveries and contributions, and suggests future lines of inquiry and advancement.

6.1 Summary

Sepside: An Edge-AI Internet of Things System for Time-to-Deterioration and Early Sepsis Warning The goal of the prediction project was to find a solution to the pressing medical issue of detecting sepsis in its early stages. The system will provide a continuous flow of patient vital signs data and predictive information on the onset of sepsis and the remaining time to clinical deterioration by utilizing IoT devices, biomedical sensors, and embedded AI models.

Biomedical signals, data preprocessing, model training, and comparative analysis of the three were all used in this study.

In order to compare the three algorithms—Random Forest, Gradient Boosting, and XGBoost—in classification and regression, this study combined biomedical signals, data preprocessing, and model training. The results also showed that XGBoost performed consistently better in the time-to-deterioration regression (MAE = 1.87, RMSE = 2.55, R2 = 0.845) and sepsis risk prediction (AUROC = 0.965, AUPRC = 0.402). These findings highlight the clinical significance and technological resilience of edge AI models by confirming their potential for use in life-critical monitoring.

In addition to its technical success, the project was evaluated based on its financial viability, ethical considerations, overall social and environmental impact, and compliance with engineering standards. According to the analysis, the system qualifies as a complex engineering.

In addition to its technical success, the project was evaluated based on its financial viability, ethical considerations, overall social and environmental impact, and compliance with engineering standards. According to the analysis, the system can be characterized as a complex engineering problem with multiple stakeholders, a wide range of knowledge domains, and important societal ramifications.

6.2 Limitation

Despite the encouraging results, a number of limitations must be noted. First, the study's dataset was synthetic, which does not accurately represent the noise and variability of real-world patient data, even though it is useful for testing and prototyping. Therefore, prior to clinical deployment, validation using actual patient datasets is crucial. Second, although functional, the small-scale hardware implementation prototype did not go through Despite the positive outcomes, there are a few things to keep in mind. First of all, the dataset used in this study was artificial and, while useful for testing and prototyping, it lacked the noise and variability found in actual patient data. Therefore, they must be validated using actual patient datasets before being used in clinical settings. Second, a prototype of the hardware was created. Second, although the hardware was successfully prototyped on a small scale, it was not thoroughly tested in hospital clinical settings. Third, the models were adjusted for accuracy and generalization within the constraints of the data that was available; however, issues such as data imbalance, patient groups that were hidden, and hardware noise remain to be resolved. Last but not least, the project raised sustainability issues; however, long-term field testing is necessary to ascertain the devices' longevity, upkeep, and function in medical procedures.

6.3 Future Work

Addressing these constraints and broadening the project's scope will be the main goals of future work. The gathering and verification of actual patient data in partnership with medical facilities is a top priority since it allows models to be improved for clinical dependability. In order to ensure interoperability with current infrastructure, integration with hospital electronic health record (EHR) systems via HL7-FHIR standards will also be sought.

Technically, future studies will look into how to increase predictive power without going over hardware constraints by utilizing deep learning architectures tailored for edge computing, such as transformer-based models or lightweight CNNs. More biosensors could be added to increase accuracy and provide richer datasets, such as blood pressure and oxygenation trend monitoring.

Addressing these constraints and broadening the project's scope will be the main goals of future work. The gathering and verification of actual patient data in partnership with medical facilities is a top priority since it allows models to be improved for clinical dependability. In order to ensure interoperability with current infrastructure, integration with hospital electronic health record (EHR) systems via HL7-FHIR standards will also be sought.

Technically, future studies will look into how to increase predictive power without going over hardware constraints by utilizing deep learning architectures tailored for edge computing, such as transformer-based models or lightweight CNNs. More biosensors could be added to increase accuracy and provide richer datasets, such as blood pressure and oxygenation trend monitoring.

References

[1] M. Singer, et al., "The Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3)," *JAMA*, vol. 315, no. 8, pp. 801–810, 2016. doi: 10.1001/jama.2016.0287.

[2] C. W. Seymour, et al., "Assessment of Clinical Criteria for Sepsis: For the Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3)," *JAMA*, vol. 315, no. 8, pp. 762–774, 2016. doi: 10.1001/jama.2016.0288.

[3] K. E. Henry, et al., "A targeted real-time early warning score (TREWScore) for septic shock," *Science Translational Medicine*, vol. 7, no. 299, pp. 299ra122, 2015. doi: 10.1126/scitranslmed.aab3719.

[4] M. M. Churpek, et al., "Multicenter development and validation of a risk stratification tool for ward patients," *Am. J. Respir. Crit. Care Med.*, vol. 190, no. 6, pp. 649–655, 2014. doi: 10.1164/rccm.201406-1022OC.

[5] Royal College of Physicians, *National Early Warning Score (NEWS) 2: Standardising the assessment of acute-illness severity in the NHS*, London, UK: RCP, 2017.

[6] National Institute for Health and Care Excellence (NICE), "National Early Warning Score systems that alert to deteriorating adult patients in hospital," Medtech Innovation Briefing, 2020.

[7] M. A. Reyna, et al., "Early Prediction of Sepsis from Clinical Data: the PhysioNet/Computing in Cardiology Challenge 2019," *Comput. Cardiol.*, vol. 46, pp. 1–4, 2019. doi: 10.22489/CinC.2019.400.

[8] PhysioNet, "Challenge 2019: Early prediction of sepsis," Available: https://physionet.org/challenge/2019/.

[9] J. Futoma, et al., "Learning to detect sepsis with a multitask Gaussian process RNN classifier," in *Proc. ICML*, 2017.

[10] J. Futoma, et al., "Improved sepsis detection in the ICU using multimodal deep learning," in *Proc. MLHC*, 2017.

[11] T. Desautels, et al., "Prediction of sepsis in the ICU using minimal electronic health record data: InSight," *JMIR Med. Inform.*, vol. 4, no. 3, e28, 2016. doi: 10.2196/medinform.5909.

[12] S. Nemati, et al., "An interpretable machine learning model for accurate prediction of sepsis in the ICU," *Crit. Care Med.*, vol. 46, no. 4, pp. 547–553, 2018. doi: 10.1097/CCM.0000000000002936.

[13] M. Moor, et al., "Predicting sepsis using deep learning across international sites: external validation," *EClinicalMedicine*, vol. 55, 2023. doi: 10.1016/j.eclinm.2023.101783.

[14] A. Boussina, et al., "Impact of a deep learning sepsis prediction model on quality of care and survival," *NPJ Digital Medicine*, vol. 7, no. 1, p. 39, 2024. doi: 10.1038/s41746-024-00901-0.

[15] S. Nemati, et al., “AI algorithm COMPOSER: reliable early sepsis prediction with calibrated uncertainty,” *NPJ Digital Medicine*, vol. 4, no. 1, 2021. doi: 10.1038/s41746-021-00423-8.

[16] S. P. Shashikumar, et al., “Machine learning to improve patient care in sepsis: validation across multiple hospitals,” *JAMIA Open*, vol. 3, no. 3, pp. 326–334, 2020. doi: 10.1093/jamiaopen/ooaa030.

[17] Z. Yang, et al., “Machine learning for sepsis prediction: a systematic review and meta-analysis,” *BMC Infectious Diseases*, vol. 23, p. 275, 2023. doi: 10.1186/s12879-023-08172-y.

[18] H. Mohamadlou, et al., “DeepAISE: an interpretable recurrent neural survival model for early prediction of sepsis,” *Artificial Intelligence in Medicine*, vol. 117, 102083, 2021. doi: 10.1016/j.artmed.2021.102083.

[19] M. A. F. Pimentel, et al., “A hazard-based early warning system for detecting deteriorating patients in hospital,” *Am. J. Respir. Crit. Care Med.*, vol. 204, no. 1, pp. 44–52, 2021. doi: 10.1164/rccm.202003-0436OC.

[20] Task Force of the ESC/NASPE, “Heart rate variability: Standards of measurement, physiological interpretation, and clinical use,” *Eur. Heart J.*, vol. 17, pp. 354–381, 1996. doi: 10.1093/oxfordjournals.eurheartj.a014868.

[21] Maxim Integrated, “MAX30102: High-Sensitivity Pulse Oximeter and Heart-Rate Sensor,” Datasheet, 2019. Available: https://datasheets.maximintegrated.com.

[22] Melexis, “MLX90614 Infrared Thermometer: Datasheet Rev. 020,” 2025. Available: https://www.melexis.com.

[23] Analog Devices, “AD8232: Low Power, Single-Lead, Heart Rate Monitor Front End,” Datasheet, 2020. Available: https://www.analog.com.

[24] F. Garcés-Jurado, et al., “Integration of MAX30102 with ESP32 for IoT SpO_2 and HR monitoring,” *MDPI Proceedings*, vol. 81, no. 1, 2022. doi: 10.3390/proceedings2022081001.

[25] M. Giordano, K. Dheman, and M. Magno, “SepAl: Sepsis alerts on low-power wearables with TinyML,” *arXiv preprint arXiv:2403.01234*, 2024.

15% Overall Similarity

The combined total of all matches, including overlapping sources, for each database.

Match Groups

- **108** Not Cited or Quoted 13%
  Matches with neither in-text citation nor quotation marks
- **1** Missing Quotations 0%
  Matches that are still very similar to source material
- **26** Missing Citation 2%
  Matches that have quotation marks, but no in-text citation
- **3** Cited and Quoted 0%
  Matches with in-text citation present, but no quotation marks

Top Sources

10% Internet sources

5% Publications

14% Submitted works (Student Papers)

Integrity Flags

0 Integrity Flags for Review

No suspicious text manipulations found.

Our system's algorithms look deeply at a document for any inconsistencies that would set it apart from a normal submission. If we notice something strange, we flag it for you to review.

A Flag is not necessarily an indicator of a problem. However, we'd recommend you focus your attention there for further review.

23% detected as AI

The percentage indicates the combined amount of likely AI-generated text as well as likely AI-generated text that was also likely AI-paraphrased.

Caution: Review required.

It is essential to understand the limitations of AI detection before making decisions about a student's work. We encourage you to learn more about Turnitin's AI detection capabilities before using the tool.

Detection Groups

20 AI-generated only **15%**
Likely AI-generated text from a large-language model.

10 AI-generated text that was AI-paraphrased **7%**
Likely AI-generated text that was likely revised using an AI-paraphrase tool or word spinner.

Disclaimer
Our AI writing assessment is designed to help educators identify text that might be prepared by a generative AI tool. Our AI writing assessment may not always be accurate (it may misidentify writing that is likely AI generated as AI generated and AI paraphrased or likely AI generated and AI paraphrased writing as only AI generated) so it should not be used as the sole basis for adverse actions against a student. It takes further scrutiny and human judgment in conjunction with an organization's application of its specific academic policies to determine whether any academic misconduct has occurred.

Frequently Asked Questions

How should I interpret Turnitin's AI writing percentage and false positives?
The percentage shown in the AI writing report is the amount of qualifying text within the submission that Turnitin's AI writing detection model determines was either likely AI-generated text from a large-language model or likely AI-generated text that was likely revised using an AI paraphrase tool or word spinner.

False positives (incorrectly flagging human-written text as AI-generated) are a possibility in AI models.

AI detection scores under 20%, which we do not surface in new reports, have a higher likelihood of false positives. To reduce the likelihood of misinterpretation, no score or highlights are attributed and are indicated with an asterisk in the report (*%).

The AI writing percentage should not be the sole basis to determine whether misconduct has occurred. The reviewer/instructor should use the percentage as a means to start a formative conversation with their student and/or use it to examine the submitted assignment in accordance with their school's policies.

What does 'qualifying text' mean?
Our model only processes qualifying text in the form of long-form writing. Long-form writing means individual sentences contained in paragraphs that make up a longer piece of written work, such as an essay, a dissertation, or an article, etc. Qualifying text that has been determined to be likely AI-generated will be highlighted in cyan in the submission, and likely AI-generated and then likely AI-paraphrased will be highlighted purple.

Non-qualifying text, such as bullet points, annotated bibliographies, etc., will not be processed and can create disparity between the submission highlights and the percentage shown.