Optimization of Personnel Response in an eICU Environment

by

Phillip Ferguson

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Engineering in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2018

Author .... Signature redacted ....
Department of Electrical Engineering and Computer Science
May 23, 2018

Certified by .... Signature redacted ....
Amar Gupta
Research Scientist
Thesis Supervisor

Accepted by .... Signature redacted ....
Katrina LaCurts
Chair, Master of Engineering Thesis Comittee

Optimization of Personnel Response in an eICU Environment

by

Phillip Ferguson

Submitted to the Department of Electrical Engineering and Computer Science
on May 23, 2018, in partial fulfillment of the
requirements for the degree of
Master of Science in Computer Science and Engineering

Abstract

In this thesis, I analyzed multiple databases of eICU data and formulated methods of improving personnel response to the incidence of alerts in these environments. Multiple aspects of the eICU response workflow were considered such as the medications provided by the healthcare providers following the triggering of an alert, the settings selected by the personnel regarding the reactivation of the alerts, as well as specific responses and their correlations to specific alert types. A new algorithm for predicting the reactivation time of alerts has been developed, which takes into account the severity and type of alert, as well as the patient status.

Thesis Supervisor: Amar Gupta
Title: Research Scientist

Acknowledgments

This is the acknowledgements section. You should replace this with your own acknowledgements.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Many individuals in the United States and abroad lack access to healthcare [59, 60]. Both the cost of service and geographical accessibility are barriers to the distribution of health services. For example, most hospitals and trauma centers are located in urban areas, which then require those living in rural areas to travel significant distances in order to receive the medical attention that they need. In rural areas, there is a dearth of medical professionals, specialists, and modern medical technology. Rural victims were over seven times more likely to die before arrival (relative risk = 7.4, 95 percent confidence interval 2.4-22.8) if the emergency medical services' response time was more than 30 minutes [1]. Telemedicine is viewed as a potential solution that allows healthcare to be more readily available to underserved and isolated communities.

Telemedicine is defined as the use of telecommunications and information technology to provide healthcare to patients across significant distances. Though the exact legal definition has changed since its emergence, telemedicine, in all of its forms involves the use of remote communication or digital technologies to provide or enhance the provision of health care. Telemedicine was utilized in 1929, where the Royal Flying Doctors Service (RFDS) made use of two-way radios in order to provide healthcare services to remote communities in Australia. In the modern era, changes in technology, such as advancements in patient monitoring technology and the widespread availability of high-speed internet have allowed telemedicine initiatives to develop and proceed at a much faster pace than was possible before.

In a modern context, telemedicine is employed in a wide variety of services. The proliferation of two-way video communication technology has enabled online communication between doctor and patient as well as online consultations between care providers. The availability of long-distance medical monitoring technology has enabled patients to receive diagnoses, treatments, as well as examinations, allowing for healthcare providers to communicate with patients and handle their cases remotely.

1.1 eICU

One such implementation of telemedicine technology is the Tele-ICU, otherwise known as an eICU [40], in which a team of nurses and physicians remotely monitor patients in an Intensive Care Unit, 24/7, as a supplement to the bedside staff. The remote team is provided with technology that tracks the patients' vital signs and medical records in order to detect changes in the patients' conditions and use that information to correct or prevent potential complications. When a patient's condition changes, an alert is sent to the remote location, notifying the remote clinician (such as a nurse, PA, doctor) who can then examine the situation via a bedside video camera and alert the on-site care providers (doctors, nurses), if needed.

In this emerging field of telemedicine, there are many questions raised regarding the quality of service and the benefits and shortcomings. Telemedicine, in an effective implementation, offers the patient faster and longer distance care, which is particularly necessary for underserved populations such as lower-income or rural communities. At the same time, there valid concerns with the possible consequences of telemedicine, such as the over-reliance on the technology and it being used improperly as a replacement for face-to-face interactions, rather than as a supplement, as it is intended to be [41]. Difficulties in the adoption of telemedicine initiatives, especially given the legal landscape of the United States in the medical domain, where different states have different degrees of regulation, are also factors to consider.

Caregivers in an eICU environment are exposed to a constant stream of information and alerts from their patients. This situation can cause alert fatigue, which, in some cases, has resulted in the monitoring professionals turning off some of the alerts.

1.2 Purpose

This thesis will examine the following hypothesis: It is possible to apply Systems Engineering techniques to enhance the impact of telemedicine services in distributed hospital networks.

Telemedicine purports to offer medical services in a faster and more effective manner in environments with it than in environments without it. To test this hypothesis, I have compared the quantity of alerts and the effectiveness of response to them by the medical staff across a variety of hospitals, mainly those that have eICU. The intent of this work is to analyze the actions taken by the medical staff after an alert is generated, for a wide variety of alert situations, comparing moderate alerts to severe alerts, as well as considering the alerts that are dismissed as compared to those that are acted upon by the personnel. This includes the actions taken by the staff when a high volume of alerts occur in a brief period of time, how the alerts are prioritized or ignored in order to minimize alert fatigue, and how the alert systems can be made more effective in the future.

Currently, we have access to a significant set of data from the Philips eICU software system. Naturally, this data has been deidentified, that is, all Protected Health Information of the patient (name, DOB, etc.) has been removed or replaced with dummy identifiers. The Philips eICU software is used in many hospitals and hospital chains across the country. The Philips Healthcare's IntelliSpace eCareManager health care platform helps to enable triage assessments [3].

Data access agreements have been completed with multiple eICU locations, which will enable us to use the relevant data in our analyses. With this data, we will carry out various machine learning algorithms and methods in order to isolate the most relevant factors and variables in these environments. Machine learning is a subfield of computer science, in which new or existing algorithms are used in order to learn from a set of existing data and create generalizable models from the data. These models would ideally offer useful predictions or determine important patterns from the existing data. In this case, we would like to be able to predict such things as length of stay and mortality for patients with a given condition or with certain biographical information, such as age or sex.

Clinical decision support (CDS) systems have been installed in numerous practice settings since the development of these systems almost two decades ago. This advanced form of information technology is characterized by a series of algorithms derived from clinical best practices and applied to patient-specific situations using real-time physiological data [4]. Developing such systems is and objective of this study and other like it.

Chapter 2

Previous Research

Prior to this project, several publications and papers have explored topics of telemedicine and data usage, especially in an ICU environment and have expounded on many important aspects of data and algorithm usage in this field. The paper published in IEEE by Johnson et al. describes various limitations inherent in the current medical environment for the application of data science techniques and what would need to change in order to move past these obstacles, as well as some viable applications of data science in critical care analysis. Another paper, written by Awad et al. delineates some examples of extrapolative algorithms in eICU environments and details how measures such as LOS (length of stay) can be predicted

Prior studies have remarked that the eICU offers an unprecedented opportunity to make use of data science in the medical field [22], but that the acquisition, analysis, interpretation, and presentation of the relevant data in a clinically relevant and usable format is the premier challenge of data analysis in critical care [9]. One aspect of this challenge include the tendency of health organizations to compartmentalize patient data, storing different bits of information in different locations, even for the same patient, which means that integration between these different locations is tantamount. Another difficulty is the corruption of data that is likely to occur in these environment; sensor failures, changes due to interventions and incomplete measurements. In these scenarios, being able to identify and remove these irregularities has been shown to have a significant effect on the quality of the results from machine learning algorithms [7].

Another challenge is the complexity of the systems involved; all analysis of medical data is dependent on the specific characteristics of the patient and the patient's condition and thus, relies on many variables that are not likely to be completely represented in the collected data [35].

Others mention various applications of machine learning that have been able to reduce the impact of the aforementioned challenges, or surmount them entirely [49]. Regression models in particular have proven to be rather effective in analyzing trends in existing data. One such advancement is the utilization of multiple SVMs (Support Vector Machines) that were combined in a regression step in order to classify and predict patient mortality, yielding a minimum Se/PPV of 53.52 percent. Another method used to predict patient mortality was a tree-based classifier that achieved a score of 53.53 percent, meaning that it was capable of correctly predicting half of the patients that would die in-hospital. The success of these regression-based methods may be attributed to the greater time scales involved in the calculations, and more complex methods have been shown to not be quite as useful.

In other studies [10], the authors mention how LOS (length of stay in the ICU) can often be used as a proxy for other characteristics such as the severity of the patient condition, especially in cases for which that information is unavailable. As with any surrogate, it cannot be considered as a complete representation of the variable it is replacing. LOS reveals information about certain non-obvious aspects of the eICU experience, such as the utilization of resources and clinical quality, which is especially important considering the high cost of ICU care on both the providers and patients.

Several machine learning methods have been applied to LOS data (and more), with the intention of predicting patient mortality and patient length of stay above a certain threshold. The study reveals that the most useful method used was the Random Forest (RF), though Multiple Regression (MP), Decision Trees (DT), Artificial Neural Network (ANN) ensemble, and Support Vector Machines (SVM) were also utilized [58]. Through these methods, various correlations were determined, such as the positive correlation between age and LOS as well as a significant tendency for patients with lung or respiratory disorders and high blood pressure.

Additionally, they discuss the potential usefulness of having daily scores of patient wellness, referencing APACHE (Acute Physiology and Chronic Health Evaluation)[59] scores as one example. Techniques such as logistic regression have been used to predict the risk of mortality and adverse events for patients admitted to the ICU, and the predictions, though promising are not accurate enough to be used on patients [60]. Overall, it was determined that there was not a particular algorithm that surpassed others in this way, but rather that the best algorithm is dependent on multiple factors.

The papers examined here were a significant inspiration for this thesis. Potential applications of Machine Learning to the field of telemedicine will be examined. We have noted that the eICU data that we have received is very large in terms of the number of patients as well as the depth of information on each patient. With this, we also needed to contend with the previously mentioned challenges of incomplete data, or highly divided data. The opportunity to receive the data directly from the healthcare providers offers a unique perspective on how the operations of these organizations can be improved and areas in which the personnel can be partially relieved of their burden.

Many of the data science techniques mentioned in this work will serve as the basis for the experimental approach of this thesis.

Chapter 3

Vision

Some machine learning algorithms that will be used include Principal Component Analysis, logistic regressions, and clustering. Principal Component Analysis is defined as a multivariate technique that analyzes a data table in which observations are described by several inter-correlated quantitative dependent variables. Its goal is to extract the important information from the table and to represent it as a set of new orthogonal variables called principal components, and to display the pattern of similarity of the observations and of the variables as points in maps [4]. The goal of logistic regression is to find the best fitting (yet biologically reasonable) model to describe the relationship between the dichotomous characteristic of interest (dependent variable) and a set of independent variables [5]. Clustering will also be employed, which involves dividing a population or set of data points into a number of groups such that data points in the same group are more similar than data points in other groups. That is, to segregate the groups with similar traits and assign them into various clusters.

Another aspect of this project was the development of visualization tools for the purpose of better understanding the data, as well as potentially serving as a useful resource for the eICU professionals for which we are performing this study. The purpose of this is twofold, one reason being that the size of the dataset may be a limiting factor for gaining an intuitive grasp of the data, and so a visualization may offer insights that were previously unnoticed.

Furthermore, visualizations may offer evidence of correlations that the eICU professionals can observe and analyze, which may positively influence their behavior. In the following sections, I will expand on a few visualization tools that I developed.

As an overview of the topics of research, there are two main aspects of the eICU workflow that are of particular interest.

One is the activation of alerts, which are intended to signal to the employees observing them that some action ought to be taken, or at the very least, that the patient should be observed in order to determine the cause of the alert. The alerts are based primarily on certain biological measures, such as MAP (mean arterial pressure) or O2/Sat (oxygen saturation), these measurements are taken from various devices and monitoring equipment that are hooked up to the patient [37]. Typically alerts will activate once one or more of these vital signs is outside of an accepted range. This accepted range has a default value, but can be customized by the eICU intensivist, which is often necessary due to the differences in patients and their conditions. There are 6 types of alerts.

Another important aspect of the eICU process is the reactivation time. After an eICU intensivist handles an alert, he or she can choose how long they wish for that (particular type of) alert to be disabled until it can be triggered again. One can imagine that there may different reasons for choosing a specific length of time for specific situations. For example, in the case of an alert being triggered by a sensor device falling off, the intensivist may wish to have the alert reactivate immediately, knowing that the patient is otherwise fine. Whereas, if the alert is triggered by a worsening of the patient's condition, he or she may elect to instead choose a longer reactivation time, knowing that the patient will need more time to recover, and thus, their vital signs will be abnormal for a significant period of time, and they do not wish to receive extraneous alerts while this is the case [36].

For the purposes of this thesis, the relationship between these two elements of the eICU experience has been explored in order to better understand how alerts are responded to and how this process may be improved or standardized in some way so as to improve patient outcome and prevent instances of phenomena such as alert fatigue. Apparently, a significant portion of the alerts that are generated are false and thus, do not require any action [38].

3.1 Steps

The following are the steps that have been taken to appropriately test and confirm the aforementioned hypothesis:

Interviews with caregivers and staff in an eICU environment have been conducted, both the in-person team and the remote team, in order to have a better operational understanding of how the system works and what the existing procedures are. These interviews have been conducted with staff from a wide variety of hospitals, such as Emory, UMass Worchester, and several others. This corresponds mainly with the phase of capturing the needs and requirements of the system, as was covered in the Systems Engineering process.

Following this, more data has been collected from these different environments in order to have a more robust analysis. Administrative approvals were necessary in order to access this data. Our next goals were to carry out some statistical and machine learning analysis so that we can determine the main correlations and causations in the system. This was needed in order to to determine the involvement and role of the various medical personnel, such as the bedside physicians and nurses and the remote physicians and nurses. With this, we defined the ideal positions and functions of the medical professionals in terms of maximizing patient outcomes.

At the conclusion of this work, the expected contributions include a set of analyses regarding clinician response in the eICU environment, as well as suggestions for improvements on the existing system. This will provide a baseline for understanding what is important in alerting and remote advice, with as much specificity as possible.

Another contribution is the visualization tools that were developed for the statistical analysis. This would serve the professionals in the eICUs and at Philips. Some examples of visualizations are timelines of patient activity and caregiver response, which would demonstrate the evolution of a particular patient's case.

With the data and information acquired from our machine learning techniques as well as from interviews with eICU professionals, a systems engineering analysis on the eICU alert-response process was conducted in order to determine the best practices for an eICU workflow and how such systems can be utilized most effectively in terms of patient satisfaction, minimizing strain on personnel, and determining the optimal conditions for collaboration between the central remote site and the bedside, where the patients are located.

Systems Engineering is an interdisciplinary engineering discipline that is concerned with the creation and development of a process that satisfies the needs of the stakeholder and the consumer and is reliable, cost efficient, and high in quality. Stakeholders in this endeavor would include the eICU personnel, hospital staff and patients, etc. All systems contain needs and requirements, which define the goals and objectives of the systems engineering endeavor. These expectations can be captured through a variety of methods such as interviews and prototyping, which we will be conducting in tandem with the eICU personnel involved. The requirements are then translated into a functional description of the system or product in development. In this work, functional descriptions of the eICU response process will be created and various visualization tools will be developed in order to capture the differences between different implementations as well as between the responses to different types of alerts. Functional analysis and allocation allows for a better understanding of what the system has to do, in what ways it can do it, and the priorities and conflicts associated with lower-level functions [8].

3.2 Data Sources

The sources of data I used were from the MIMIC-III database, the eICU Collaborative Research Database, as well as some data that was provided by specific hospital chains.

MIMIC-III (Medical Information Mart for Intensive Care III) is a freely available database of deidentified (all protected health information such as the name, date of stay, etc. are removed) health-related data associated with over forty thousand patients who stayed in critical care units of the Beth Israel Deaconess Medical Center between 2001 and 2012 [13]. The database contains information pertaining to patient demographic, vital signs throughout the patient's stay, test results, imaging procedures, notes, and mortality. This data has been made publicly available in order to allow for studies such as this one to be performed, ideally so that the results of such studies may be utilized for the benefit of the critical care unit providers. Researchers that wish to utilize this data must complete a CITI (Collaborative Institutional Training Initiative) course to verify their knowledge of HIPAA protocols as well as appropriate behavior with sensitive data of this nature. MIMIC-III is the third iteration of the MIMIC initiative, and currently contains data from 53,423 distinct hospital admissions for adult patients (aged 16 years or above).

For the purposes of my research, the data tables that pertained to patient demographics was heavily used, as well as those pertaining to chart events and diagnoses.

The eICU Collaborative Research Database is a database comprised of deidentified patient information from the Philips eICU Research Institute database. The database includes information such as demographics, vital sign measurements made at the bedside, laboratory test results, procedures, and medications. As with the MIMIC-III database, the data is freely available to researchers, so long as they complete the CITI course. Currently the database holds about 200,000 patient entries [13].

For the purposes of my research, a significant subset of the eICU Collaborative Research Database data was used, including the tables pertaining to drug infusion, patient demographics, the hospitals that at which the patients were located, and alert activation.

The hospital chain data was comprised of deidentified patient information from the hospital chain's database. The database was comparable in size to the others used, possessing nearly 500,000 patient entries. Though it was from a different source, the data present in this database contained similar fields concerning patient demographics, and the triggering of alerts, which I made use of for the purposes of this research.

Chapter 4

Initial Analysis

4.1 Introduction

Reactivation time is an important consideration for the eICU specialists, as they must decide how long they wish for an alarm to be inactive after addressing it. For example, if a MAP (Mean Arterial Pressure) alert activates, and the intensivist answers the alert, he or she may choose to deactivate the alarm for some amount of time, such as 5 minutes, 10 minutes, an hour, or more, or no time at all, depending on their selection. With this selection, no MAP alerts will activate for the chosen duration, even if the biological requirements for the alerts activating are fulfilled [24]. This is all the more important considering the nature of alert fatigue, a tendency of providers to dismiss potentially useful alerts due to a high number of unhelpful false alarms [38]. Naturally, healthcare providers in these situations must balance the need to set the reactivation time to be short enough so that they do not miss any important alerts that may arise, and to be long enough to prevent repeated alerts from occupying the attention of the eICU specialist, especially since the specialist would already be aware of the issue that the alert was expressing [25]. In short, one potential area of study is to determine methods to minimize the load on the eICU professionals and to free up their time in order to focus on more relevant signals.

In order to properly understand the current state of eICU alert and response, I compiled some initial statistics regarding the choice in reactivation time as correlated

Table 4.1: Incidence of reactivation with respect to alert type for a dataset received from a hospital collaborator.

	Combined O2/RR	Heart Rate Limit	Mean Arterial Pressure	Heart Rate Trend	Mean Arterial Pressure Trend
Reactivate Now	759	689	1611	302	748
Reactivate in 5 minutes	156	70	159	28	69
Reactivate in 15 minutes	40	23	46	4	4
Reactivate in 30 minutes	5	3	2	2	0
Reactivate in 1 hour	4	2	2	0	0
Reactivate in 6 hours	0	0	1	0	0

Table 4.2: Incidence of reactivation with respect to alert type for a dataset received from another hospital collaborator.

	Combined O2/RR	Heart Rate Limit	Mean Arterial Pressure	Heart Rate Trend	Mean Arterial Pressure Trend
Reactivate Now	1068	449	978	300	782
Reactivate in 5 minutes	303	58	131	25	79
Reactivate in 15 minutes	139	18	66	4	26
Reactivate in 30 minutes	30	3	13	0	1
Reactivate in 1 hour	7	2	3	0	0
Reactivate in 6 hours	1	1	0	0	0

to alert type. The alert types examined were the Combined O2/Sat, Heart Rate Limit and Mean Arterial Pressure.

4.2 Results

The program used to perform this task is based in Python, and makes use of various packages such as numpy and pandas in order to store the data, as well as sklearn and statmodels in order to perform the necessary statistical work and analysis. The data the algorithm was tested on was acquired from collaborating hospitals.

4.3 Analysis

As can be observed from Table 4.1 and Table 4.2, the eICU professionals overwhelmingly select to reactivate immediately across all alert types, which will often cause the same alerts to trigger for a given patient. According to prior analyses involving eICU professionals, this may be due to a desire to have a constant reminder of the alert that is occurring, especially in a quickly evolving situation. Others claim that this may be a result of the Reactivate Now setting being the default value, so the users select it out of convenience, or out of a desire to not make adjustments to the settings of the program [37].

Chapter 5

Post-Alert Follow-Up

5.1 Motivation

Another important aspect of eICU response is the follow-up action performed by the eICU personnel as they respond to the triggered alert. The purpose of this research is to optimize methods of personnel response to the alerts, which makes the immediate actions performed after the alerts have been sounded all the more important [22]. The responses in question are the administration of medicine to the patient. The following is a compilation of the most common alert responses and the amount of time necessary to enact the response.

5.2 Method

The program used to perform this task is based in Python, and makes use of various packages such as numpy and pandas in order to store the data. The data was combed to find the timestamps of alert and medicine administration occurrences for a given patient. I searched for alerts that were directly followed by instances of medicine being administered and isolated those cases, seeing those as the likeliest scenarios in which a medication was administered due to the occurrence of an alert, and thus the closest representation of the behavior of eICU personnel after an alert has been sounded.

The occurrences of these follow-up actions were then sorted and classified based on the type of the alert that was triggered and the specific medical response. It is important to note that the following statistics describe the most commonly ordered medications immediately after an alert is activated and that they do not suggest a causative relationship between the alert and the medication, merely that these were the next medications ordered, and thus may only represent a coincidental relationship.

Table 5.1: Most common follow-up actions after activation of Heart Rate Trend alert. Compiled for the purposes of this research.

Medication	Prevalence
Fentanyl	5.5539216000000002E-2
Vasopressin	2.8968190000000001E-2
Norepinephrine	2.706093E-2
Midazolam	1.9599938000000001E-2
Insulin	1.33694771523828E-2

Table 5.2: Most common follow-up actions after activation of Mean Arterial Pressure Trend alert.

Medication	Prevalence
Metoprolol	5.9336800000000002E-2
Vancomycin	3.9997199999999997E-2
Atorvastatin	8.1374500000000002E-2
Midazolam	8.0528000000000006E-3
Lisinopril	7.2963999999999998E-3

Table 5.3: Most common follow-up actions after activation of Combined O2/RR alert.

Medication	Prevalence
Oxygen Therapy (> 60%)	9.7873000000000002E-2
Amiodarone	5.3872999999999997E-2
Dexmedetomidine	3.8398000000000002E-2
Lisinopril	2.0133000000000002E-2
Norepinephrine	1.1109000000000001E-2

Table 5.4: Most common follow-up actions after activation of Heart Rate Limit alert.

Medication	Prevalence
Atorvastatin	6.1875242422241197E-2
Carvedilo	4.9517674397843499E-2
Dexmedetomidine	4.8537408581969199E-2
Insulin	2.8272052291139001E-2
Lisinopril	2.74199409874392E-2

Table 5.5: Most common follow-up actions after activation of Mean Arterial Pressure Limit alert.

Medication	Prevalence
Vasopressin	9.8340628556184806E-2
Dexmedetomidine	3.9015637282951703E-2
Insulin	2.0749854287568901E-2
Midazolam	1.8772854077133399E-2
Norepinephrine	1.6533592386900001E-2

5.3 Analysis

Table 5.1 contains information regarding the soonest ordering of medicine after the time of the alert. In this case, the five most common responses are shown for each alert type, as well as the average amount of time between the activation of the alert and the time of the medication being administered. It is important to note that more than 95 percent of alerts did not have any appreciable action that followed their occurrence, and this should be taken into consideration in future studies of this issue.

Chapter 6

Reactivation Prediction

6.1 Introduction

In order to minimize the load on the eICU professionals, a decision support algorithm has been developed that predicts the reactivation time that will be chosen based on the alert that has been activated. To do so, regression techniques will be used.

Linear regression is a technique in statistics and machine learning designed to determine whether there is a linear relationship between a dependent variable and one or more independent variables. As is often the case in the evaluation of treatment modalities, there exists multiple outcomes of interest. Utilizing a multivariate statistical tool facilitates the examination of the effect of a specific intervention on multiple variables and tests hypotheses of correlation between the variables.

Regression analysis of the data is performed to test the validity of the hypothesized model and specifically determine the effect of the independent (predictor) variables on the dependent (response) variables [4]. To examine the stated theoretical associations, path analysis is implemented to examine the causal relationships among the variables of interest [5]. Additionally, error terms suggesting lack of model fit are identified and revisions to the model are made accordingly. For the purpose of this work, I intended to determine significant correlations between the different types of alerts and the responses on the part of the attending professionals. As mentioned before, there are six different alert types, and I sought to determine the most correlated events for each of them.

According to past studies , linear regression has been used in eICU contexts in the past in order to determine the correlations between specific aspects of the eICU workflow and patient outcomes, including mortality and readmission. Although ICU telemedicine adoption resulted in a small relative overall mortality reduction, there was heterogeneity in effect across adopting hospitals, with large-volume urban hospitals experiencing the greatest mortality reductions [40, 49]. This, however, would be the first instance of using linear regression in this particular fashion.

6.2 Method

Linear regression is a basic and commonly used type of predictive analysis. The overall idea of regression is to examine two things: (1) does a set of predictor variab les do a good job in predicting an outcome (dependent) variable? (2) Which variables in particular are significant predictors of the outcome variable, and in what way do they, as indicated by the magnitude and sign of the beta estimates, impact the outcome variable? [17]

For such methods, the data is typically split into test and training sets. The training set is used in order to form the model of the data and then the model is tested on the test set, which determines the degree of error and accuracy of the model; if the data is split evenly, then a good performance on both sets is then indicative of a model that is robust and is not a victim of overfitting.

This is often due to an overly complex model (i.e. too many features/variables compared to the number of observations). As a result, this model will be ill-equipped to handle data that is different from the training set. It is because this model is not generalized, meaning that inferences cannot be made on other data, an undesirable outcome.

The tool used to perform the linear regression was the sklearn linear model package in Python, which has various linear regression functions. I also made use of sklearn cross validation package, which allowed me to split the data into test and training sets in order to verify my conclusions.

Logistic regression is most useful for predictions in which the dependent variable is binary in nature. Like all regression analyses, the logistic regression is a predictive analysis. Logistic regression is used to describe data and to explain the relationship between one dependent binary variable and one or more nominal, ordinal, interval or ratio-level independent variables [6].

The program used to do so, is based in Python, and makes use of various packages such as numpy and pandas in order to store the data, as well as sklearn and statmodels in order to perform the necessary statistical work and analysis. The program makes use of logistic regression in order to form its predictions. Logistic regression makes use of past records in order to predict classifications for a set of data, usually in a binary fashion, though in this case, I made use of a multi-class version, in which the past records were the alert events, and the classifications being predicted were the reactivation times that were selected by the professional. The data the algorithm was tested on was acquired from a collaborating hospital. The data was split into training and testing sets, with the regression trained on the training set.

Another highly relevant aspect of the algorithm is how it is intended to combat the issue of repeated alerts in a short period of time. One issue that often arises in the case of eICU monitoring is the issue of alert fatigue in the context of receiving multiple unactionable alerts. Naturally, receiving many instances of the same alert in a short period of time is not likely to be particularly helpful for the professional answering these alerts, as they would already be aware of situation that resulted in the earlier alerts. In response, the algorithm has included an analysis of the degree of change in the status of the patient. In order to measure this, the prediction algorithm accounts for several possible changes that can occur, namely:

- A change in the severity of the alert type, that is, between the 'Moderate' and 'Severe' settings.

- A change in the alert type. As mentioned before, there are 6 different alert types (MAP, HR, etc.).

Table 6.1: General Linear Regression Results

Response	R squared
Reactivate Now	.84346472456
Reactivate in 5 minutes	0.55427067017700005
Reactivate in 15 minutes	0.119918237565
Reactivate in 6 hours	4.7662893354299997E-3

- A change in the DRS score. As previously stated, the DRS score is a composite of a variety of physiological signals [33]. Early studies of the DRS score show a high NPV (negative predictive value), suggesting that they may be well suited to determining low risk patients that can be safely transferred out of the eICU [16].

Considering all of these changes, the algorithm will make a prediction of the reactivation time so that the next alert to be sounded will demonstrate a change in one of these three factors.

6.3 Results

Table 6.1 contains information regarding the correlation of response types with the occurrence of an alert. In it, we can see that the shorter response times are more highly correlated with the incidence of an alert, meaning that the shorter times are more often chosen by the professionals, indeed, in a predicable fashion.

Figure 6-1 depicts the correlation between the O2/Sat alert type and the 6 hour reactivation time. We can see that the R-squared value is .001, very low, which is an indicator of the infrequency of the 6 hour alert time being chosen for the reactivation. However, the F-statistic in this case is .993, which indicates that the reactivation time variable is highly significant to any predictions made about the alerts, which demonstrates that the other reactivation values are more highly correlated. Compare this with figure 6-2, which depicts the correlation between the O2/Sat alert type and the 15 minute reactivation time, in which the R-squared value is 4 times as high,

```
                            OLS Regression Results
==============================================================================
Dep. Variable:     alertName_Combined O2 Sat/RR   R-squared:                       0.001
Model:                                      OLS   Adj. R-squared:                 -0.000
Method:                           Least Squares   F-statistic:                    0.9334
Date:                          Sat, 31 Mar 2018   Prob (F-statistic):              0.334
Time:                                  18:29:04   Log-Likelihood:                -654.26
No. Observations:                          1000   AIC:                             1313.
Df Residuals:                               998   BIC:                             1322.
Df Model:                                     1
Covariance Type:                      nonrobust
=========================================================================================
                                coef    std err          t      P>|t|      [95.0% Conf. Int.]
-----------------------------------------------------------------------------------------
const                         0.3186      0.015     21.604      0.000         0.290     0.348
response_Reactivate in 6 hours  -0.3186   0.330     -0.966      0.334        -0.966     0.329
==============================================================================
Omnibus:                       81.561   Durbin-Watson:                   1.639
Prob(Omnibus):                  0.000   Jarque-Bera (JB):              181.798
Skew:                           0.779   Prob(JB):                     3.34e-40
Kurtosis:                       1.609   Cond. No.                         22.4
==============================================================================
```

Figure 6-1: Linear regression results for O2/Sat alert correlated with 6 hour reactivation time

showing the greater degree of correlation that exists in that case. The rest of the correlations are similar in nature, and so I have opted to not include them for the sake of readability. A similar trend is notable among the other alert types and the other possible selections for the reactivation time.

The overall results of the classification are as can be seen in Tables 6.1 and 6.2, where the measures of the specificity, sensitivity, f-score, overall accuracy, positive predictive value and negative predictive value are encapsulated. In brief, the columns represent the classifications being made by the program and the numerical values are different measures of accuracy in relation to that particular prediction made by the logistic classifier. To summarize the terms:

- TP is the True Positive Value, which represents the proportion of positives that are correctly identified as such (e.g. the percentage of sick people who are correctly identified as having the condition) [9]. Naturally, in this case, this relates to the ability of the classifier to predict the reactivation time correctly with the actual value.

- TN is the True Negative Value, which represents the proportion of negatives that are correctly identified as such (e.g. the percentage of healthy people who are correctly identified as not having the condition)

Table 6.2: Logistic Regression Results

	Reactivate Now	Reactivate in 5 min
TP	0.6	0.1
TN	0.77778	1
PPV	0.230769	1
NPV	0.9459	0.81632
FPR	0.222222	0
FNR	0.4	0.9
FDR	0.769231	0
ACC	0.76	0.82
F1	0.333332	0.18181800

Table 6.3: Logistic Regression Results p.2

	Reactivate in 60 min	Reactivate in 6 hours
TP	1	1
TN	0.89189200000000002	1
PPV	0.764706	1
NPV	1	1
FPR	0.108108	0
FNR	0	0
FDR	0.235294	0
ACC	0.92	1
F1	0.86666699999999997	1

```
                            OLS Regression Results
==============================================================================
Dep. Variable:     alertName_Combined O2 Sat/RR   R-squared:                       0.006
Model:                              OLS   Adj. R-squared:                  0.005
Method:                   Least Squares   F-statistic:                     6.063
Date:                  Sat, 31 Mar 2018   Prob (F-statistic):             0.0140
Time:                          18:44:10   Log-Likelihood:                -651.70
No. Observations:                  1000   AIC:                             1307.
Df Residuals:                       998   BIC:                             1317.
Df Model:                             1
Covariance Type:              nonrobust
======================================================================================================
                                      coef    std err          t      P>|t|      [95.0% Conf. Int.]
------------------------------------------------------------------------------------------------------
const                               0.3099      0.015     20.572      0.000         0.280     0.339
response_Reactivate in 15 minutes   0.1693      0.069      2.462      0.014         0.034     0.304
==============================================================================
Omnibus:                    63678.561   Durbin-Watson:                   1.632
Prob(Omnibus):                  0.000   Jarque-Bera (JB):              178.616
Skew:                           0.778   Prob(JB):                     1.64e-39
Kurtosis:                       1.635   Cond. No.                         4.69
==============================================================================
```

Figure 6-2: Linear regression results for O2/Sat alert correlated with 15 min reactivation time

- PPV is the Positive Predictive Value, which represents the proportion of positively classified results that are true positives, that is, correctly identified.
- NPV is the Negative Predictive Value, which represents the proportion of negatively classified results that are true negatives, that is, correctly identified.
- FPR is the False Positive Rate, which represents the proportion of incorrectly classified results among the positives.
- FNR is the False Negative Rate, which represents the proportion of incorrectly classified results among the negatives.
- FDR is the False Discovery Rate, which represents the is the expected proportion of type I errors, that is, false positives.
- ACC is Accuracy which is the proportion of true positives and negatives to the whole population, that is, the proportion of correct classifications.
- F1 is the F1-score, a measure of the harmonic mean of the precision and recall of the experimental classifier.

In Figure 6-4, one can see the general results for the logistic regression analysis, containing information relating to the coefficients, standard errors, and confidence

Logit Regression Results

Dep. Variable:	response	No. Observations:	200
Model:	Logit	Df Residuals:	196
Method:	MLE	Df Model:	3
Date:	Wed, 23 May 2018	Pseudo R-squ.:	0.1037
Time:	10:11:20	Log-Likelihood:	-124.25
converged:	True	LL-Null:	-138.63
		LLR p-value:	2.523e-06

	coef	std err	z	P>|z|	[95.0% Conf.	Int.]
alertName_Combined 02 Sat/RR	-2.0680	0.569	-3.636	0.000	-3.183	-0.953
alertName_Heart Rate Limit	-1.0877	0.580	-1.876	0.061	-2.224	0.049
alertName_Mean Arterial Pressure Limit	-0.2543	0.286	-0.889	0.374	-0.815	0.306
severity	2.3442	0.541	4.334	0.000	1.284	3.404

Figure 6-3: Logistic regression results for reactivation time and alert type.

intervals of the experimental classifier. The p-values for the severity and the Combined 02 Sat variable are smaller than 0.05, therefore, they are significant to the model. We can consider these to significantly explain the variance in the dependent variable, that is, the response to the alert in the form of reactivation time. Additionally, each estimated coefficient is the expected change in the independent variable as compared to the dependent variable, and so they describe the relationship between the two aspects of the analysis.

In Figure 6-5, a heatmap describing the relationship between different elements is shown. A heatmap is a two-dimensional representation of a dataset in which different elements of the data are colored differently. It provides an immediate summary of a set of data, especially one in which it is important to note repeating values. A heatmap of the smartAlert data can be found in Figure 6-3. This heatmap was intended to map the different alert types to the different reactivation responses, namely to reactivate immediately, or after some period of time, such as 5 or 10 minutes. This was intended to offer further understanding into the behaviors of the eICU clinicians as they were presented with the different situations that arose in the eICU environment. They are very useful plots for visualizing the measurements for a subset of rows over all the samples.

In the future, plots such as this may be utilized in the effort to convert the existing data regarding eICU responses to more visually understandable forms.

From this, we can form some conclusions as to how these alerts are generally responded to. For example, we can see that the Heart Rate Limit Alert is most

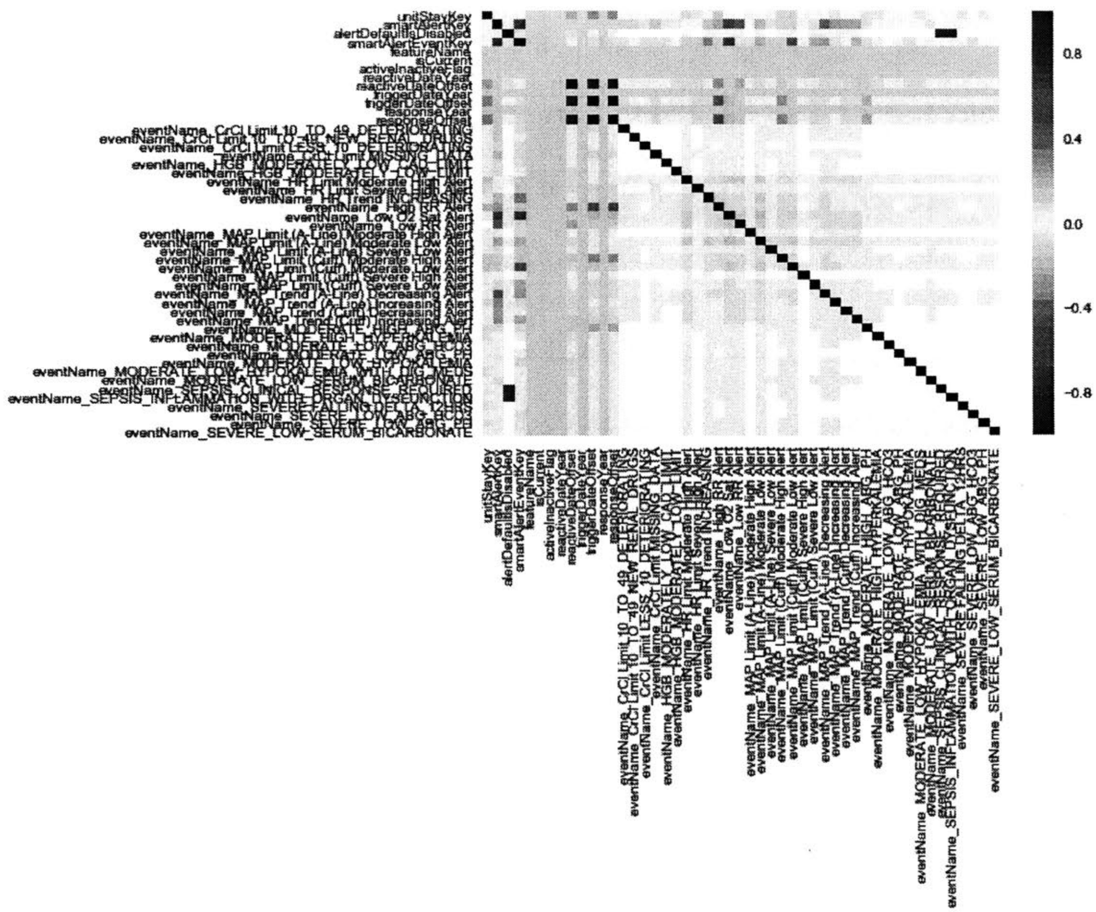

Figure 6-4: Graph showing a heatmap corresponding to the incidence of alert events.

correlated with the Reactivate Now option, the Heart Rate Trend is most correlated with the Reactivate in 5 Minutes option. Additionally, one may note that in the case of the Mean Arterial Pressure Limit and Mean Arterial Pressure Trend, that the Reactivate Now and the Reactivate in 5 Minutes are the most heavily correlated with the occurrences of these particular alert types. It may also be seen that in the case of the O2 Sat that the most commonly correlated alert types are those related to the Reactivate in 5 minutes and the Reactivate in 15 minutes options. It is also noteworthy that the alert with the highest reactivation time correlation would be the alert for sepsis, which is the Reactivate in 72 hours option. Naturally, this may be due to unique clinical challenges pertinent to dealing with sepsis [62].

The program needed to create the graph was written in Python, and made use of

Table 6.4: A table depicting the most highly correlated reactivation type for each alert as well as the coefficient of correlation for it.

Alert Type	Most Correlated Reactivation	Coefficient of Correlation
Heart Rate Limit	Reactivate Now	.13
Heart Rate Trend	Reactivate in 5 Minutes	.04
Mean Arterial Pressure Limit	Reactivate Now	.07
Mean Arterial Pressure Trend	Reactivate Now	.057
O2/Sat Limit	Reactivate in 15 Minutes	.078
Sepsis	Reactivate in 72 Hours	.71

the seaborn and numpy libraries in order to collect the data, and the matplotlib.pyplot package in order to plot the resulting graph.

The confusion matrix gives the number/proportion of instances between the predicted and actual class. A confusion matrix consists of an n by n matrix, where n is the number of possible classes in the data. Each row or column of the matrix represents one of the classes that the classifier can select. In the case of this particular classifier, the values represent 0, 5, 10, 30, and 60 minutes respectively. Thus entry (i,j) in the matrix represents the number of items of class i predicted as a member of class j by the classifier. In this case, one can see that the diagonal of the matrix is of particular importance, because it represents the number of items that are correctly identified as a member of their own class. The selection of the elements in the matrix feeds the corresponding instances into the output signal. This way, one can observe which specific instances were misclassified and how. Recall can be defined as the ratio of the total number of correctly classified positive examples divide to the total number of positive examples. Recall indicates the class is correctly recognized (small number of false negatives). To get the value of precision we divide the total number of correctly classified positive examples by the total number of predicted positive examples. High precision indicates an example labeled as positive is indeed positive (small number of false positives) [63].

One can see the summary for the confusion matrix as applied to a subset of 50 items in the following table, which contains information pertaining to the precision and recall of the data. In this case, it is of particular note that the classes representing longer reactivation times namely thirty and sixty minutes are able to avoid any

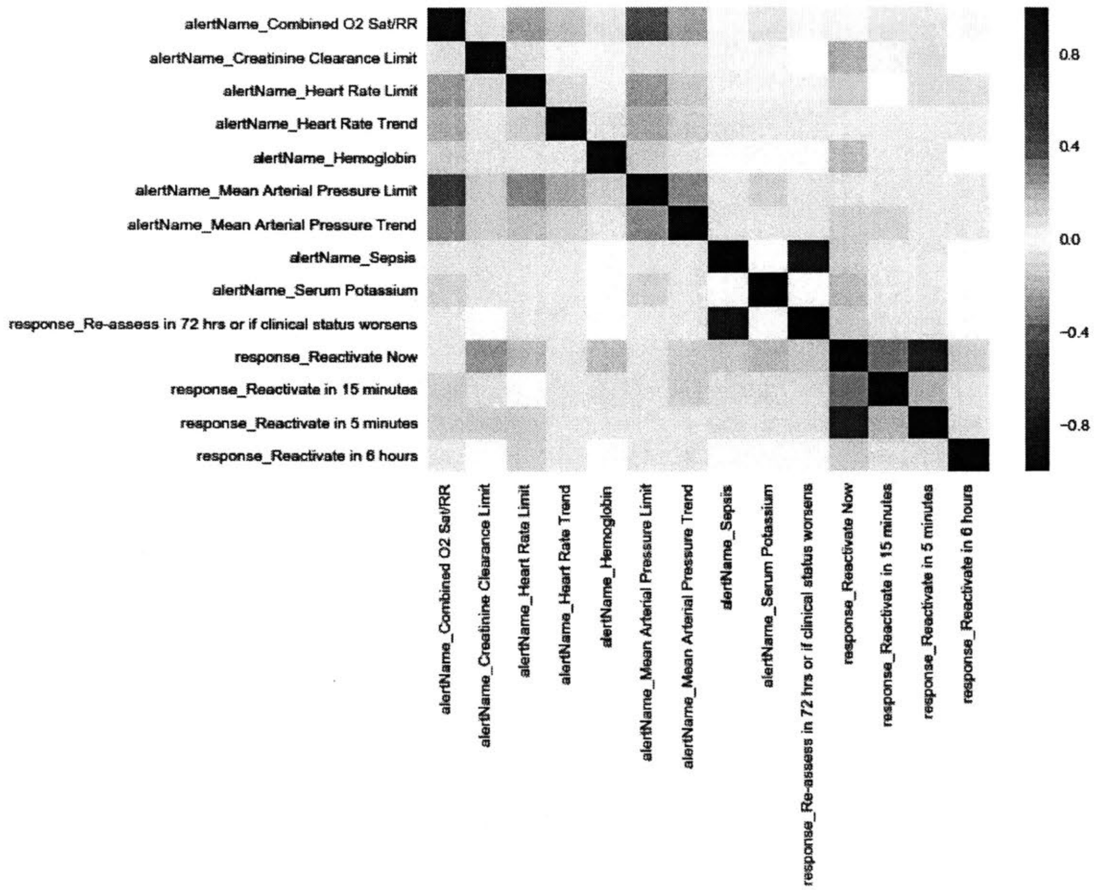

Figure 6-5: Graph showing a heatmap corresponding to the incidence of a small number of alert types.

instances of false positives and false negatives, which implies that there these longer time durations are only selected in very specific situations. As a result, the classifier can predict these situations with a high accuracy. Naturally, the lower durations of time are not as easy to accurately predict, which suggests that they are often employed in a wide variety of situations, which corroborates the previous statements made with respect to their usage. A concern of setting longer reactivation times may be that if they silence the alarms for longer periods of time, the eICU clinicians lose the ability to identify noteworthy trends earlier.

In this work, a method for predicting the choices of the eICU responders has also been determined, which suggests that one can develop decision support tools that can have a default reactivation time value based on the available parameters of the alert

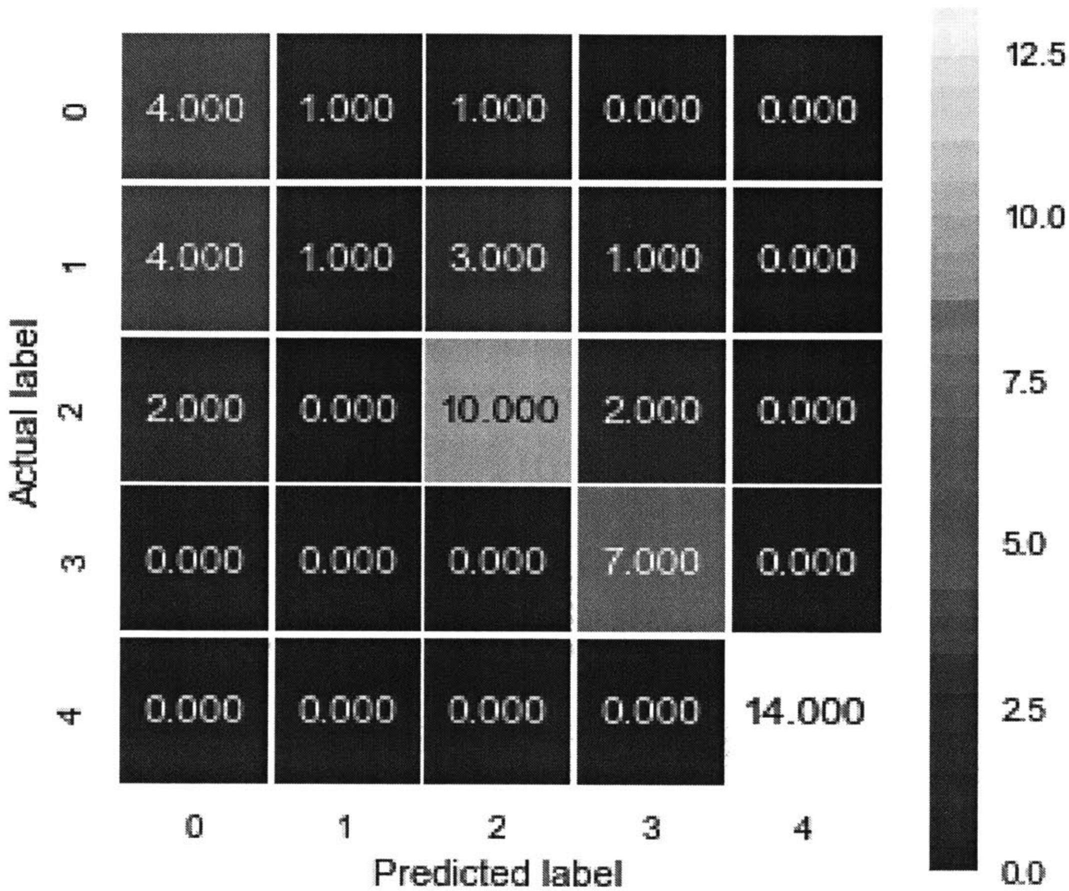

Figure 6-6: Graphic depicting confusion matrix pertaining to results of reactivation analysis.

incidence and will select the appropriate response based on the actions previously performed by the doctors in similar situations as well as based on the prescribed medicine and amount of time typically required to see relevant changes in the patient given their condition.

Figure 6-7 depicts the receiver operation characteristic (ROC) for the analysis as well as the AUC (area under the curve) for it, which shows that the classifier performs at a better than average rate. ROC curves are frequently used to show in a graphical way the connection/trade-off between clinical sensitivity and specificity for every possible cut-off for a test or a combination of tests. In addition the area under

Table 6.5: A summary of the confusion matrix statistics, namely the precision and the recall.

Reactivation Type	Precision	Recall
Reactivate Now	40%	40%
Reactivate in 5 minutes	11.111%	50%
Reactivate in 10 minutes	71.429%	71.429%
Reactivate in 30 minutes	100%	70%
Reactivate in 60 minutes	100%	100%

Table 6.6: Area under the curve for the receiver operating characteristic for each alert type.

Reactivation Type	AUC
Reactivate Now	.58
Reactivate in 5 Minutes	.59
Reactivate in 15 Minutes	.64
Reactivate in 30 Minutes	.72
Reactivate in 60 Minutes	.54

the ROC curve gives an idea about the benefit of using the classifier in question. The area under the curve measures discrimination, that is, the ability of the test to correctly classify the test data. The receiver operation characteristic randomly picks two data points; one from the positive group and one from the negative group and test the classifier on both.The area under the curve is the percentage of randomly drawn pairs for which this is true (that is, the test correctly classifies the two data points in the random pair). As a result, every point on the ROC curve represents a chosen cut-off or a trade-off between sensitivity and specificity. It is a plot of the True Positive Rate (on the y-axis) versus the False Positive Rate (on the x-axis) for every possible classification threshold.

In this case of the ROC curve, they are most often used to describe the behavior of a binary classifier. As mentioned earlier, the classifier being analyzed here is intended to classify across multiple classes. As a result, the ROC and AUC have been calculated for each of the possible classifications that can be made, and the data has been placed in a binarized format such that the correct value for a particular test is a 1 and all of the other possible classifications is a 0. The AUC is depicted for each possible classification in the following table.

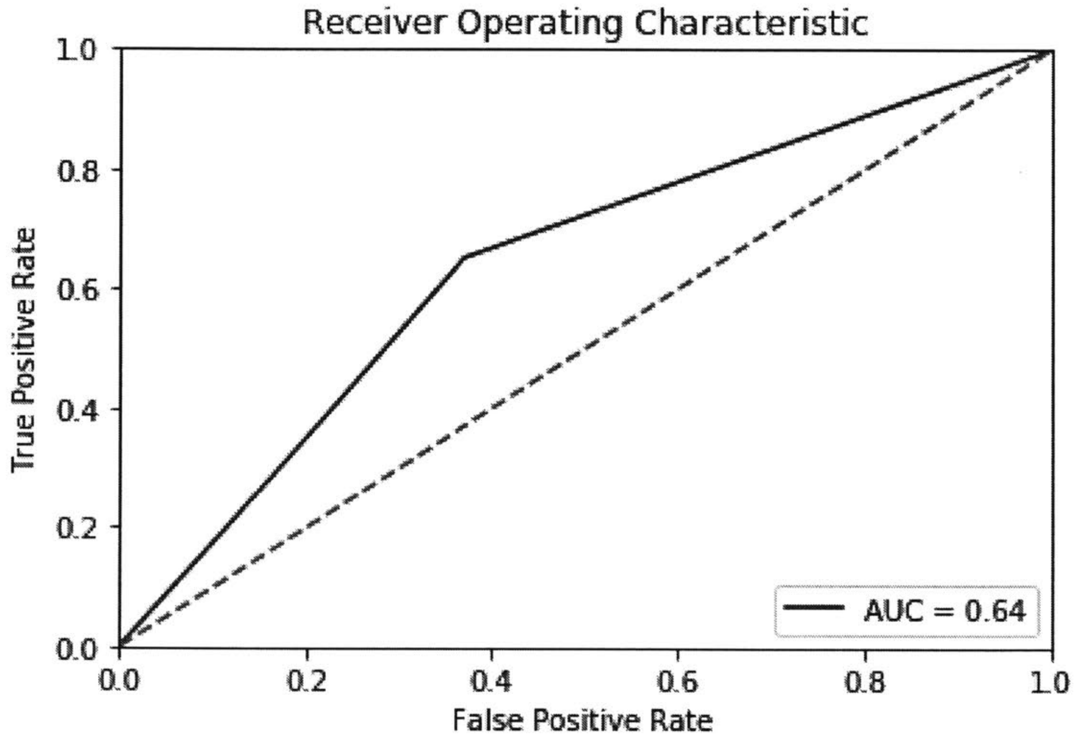

Figure 6-7: Receiver operating characteristic for logistic response analysis. This particular curve encompasses all of the classification options

From the table available, one can clearly see that all of the prediction values that correspond to the AUC are over .5, which suggests that the classifier is viable across all alert types.

6.4 Conclusions

The results from the linear regression continue to support our hypothesis that the selections for the reactivation time are shorter than necessary, especially given the results of Chapter 5, which imply that the most significant actions do not occur until a significant amount of time (more than 30 minutes) after the alert has been triggered, and so the choices of the reactivation time are excessively short, which may well act to the detriment of the healthcare provider, who may be overwhelmed with the alerts and suffer from alert fatigue, as has been mentioned before.

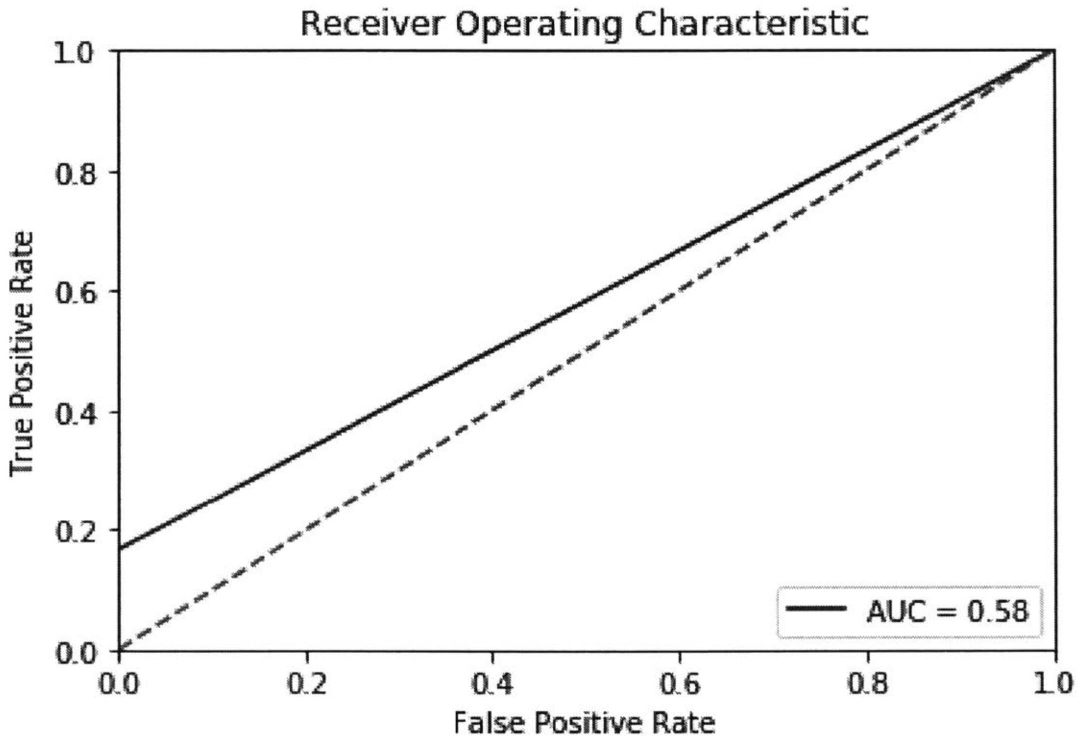

Figure 6-8: Receiver operating characteristic for logistic response analysis. This graph measures the effectiveness of predicting a 0 minute reactivation time.

While useful, linear regression may not be ideal for some aspects of this analysis, as it presumes linearity in the data, and is not well suited to variables that have a limited number of possible values, that is, categorical variables. Since categorical variables were widely used in the data we were provided, another technique, logistic regression will be used.

Logistic regression was also used in the development of the reactivation prediction algorithm. Rather than choosing parameters that minimize the sum of squared errors (like in ordinary regression), estimation in logistic regression chooses parameters that maximize the likelihood of observing the sample values.

The results of this examination point to the conclusion that the current actions being taken by the eICU professionals can be predicted rather accurately with an algorithm that takes in prior data regarding their responses. More important is the fact that these results, especially the low PPV and FPR, imply that the actions of the

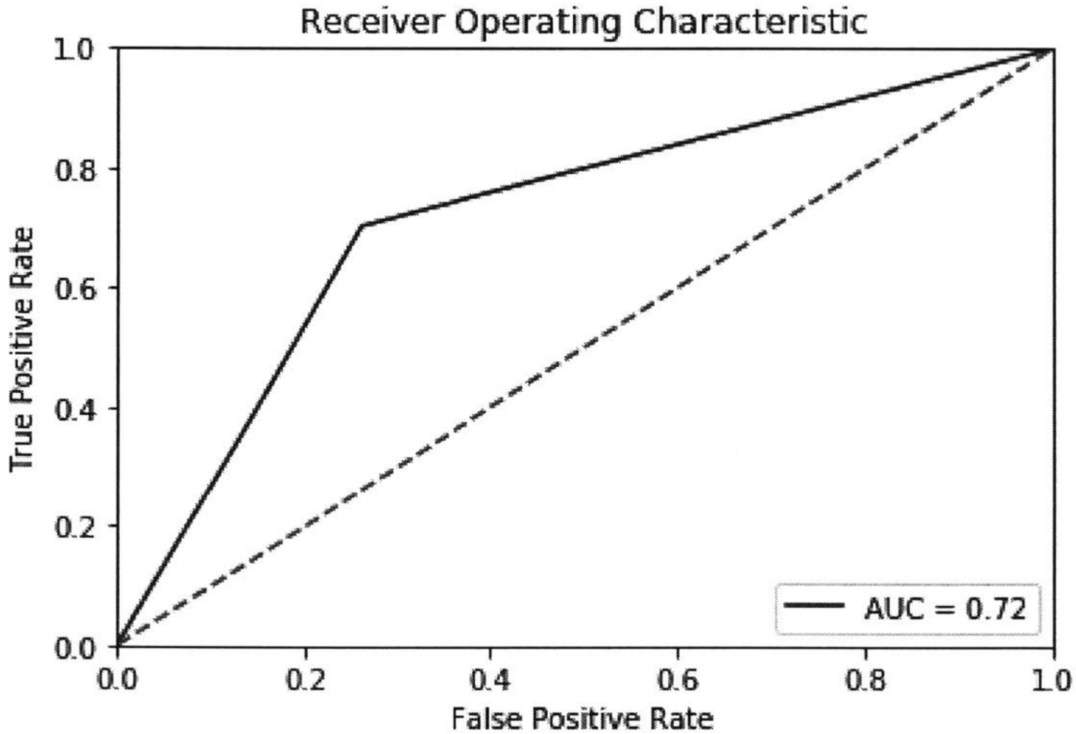

Figure 6-9: Receiver operating characteristic for logistic response analysis. This graph measures the effectiveness of predicting a 30 minute reactivation time.

personnel are incorrect in that they should delay the reactivation of their alerts for longer periods of time. In this vein, there now exist further statistical evidence that the reactivation times ought to be increased, especially in the wake of information regarding the effects of alert fatigue [38], and the repeated detriments that it has on performance in medical environments.

As mentioned earlier, the results of the logistic regression offer more detail and insight then those of the linear regression on account of the multimodal environment in which the alerts are deployed. Naturally, healthcare providers in these situations must balance the need to set the reactivation time to be short enough so that they do not miss any important alerts that may arise, and to be long enough to prevent repeated alerts from occupying the attention of the eICU specialist, especially since the specialist would already be aware of the issue that the alert was expressing. In order to minimize the load on the eICU professionals, a decision support algorithm has been developed that predicts the reactivation time that will be chosen based on the alert that has been activated.

Chapter 7

Contributions and Future Work

7.1 Contributions

7.1.1 Care Provider Response Analysis

An important result from this thesis is the determination of the most common follow-up medications and the time necessary to administer them after an alert is sounded as can be seen in Figure 5-1. This result has been determined for the 5 most common types of alerts and the relative incidence of the different medications and the average amount of time for the medication to be administered afterwards. With this, one can imagine a guide to alert response and treat this guide as a potential back-up to assist the professionals in the eICU environments in unusual cases, or as a beginning guide to assist those who may be starting out in their careers in the eICU.

In this work, a method for predicting the choices of the eICU responders has also been determined, which suggests that one can develop decision support tools that can have a default reactivation time value based on the available parameters of the alert. A wide variety of correlation analyses was also performed on the data regarding the follow-up of eICU care providers. For example, the correlations between the reactivation time and alert type has been quantified in a variety of manners. From the heatmap, it could be seen that the Reactivate Now and the Reactivate in 5 Minutes were the most highly correlated with most of the alert types, with the exception of the O2 Sat.

7.1.2 Reactivation Prediction

As can be seen in the Conclusions sections of the prior chapters, a primary cause of the difficulties encountered by the eICU personnel in terms of the volume of alerts can be found in their choice of reactivation time, often setting it to the lowest option, even when the more immediate actions that would need to be taken often occur much later. In this period of dead time, in which the responder is aware of the alert and before the administration of the medication to the patient, a high volume of repeat alerts is often received by the eICU personnel.

Some possible reasons for this reticence to choose a longer reactivation time may include a fear of the possibility of missing an important alert, even if it is repeated [35]. Indeed, many healthcare providers may feel that even if they receive thousands of false or unactionable alerts [54], that the one or two actionable alerts that they do receive may be worth the extra effort. Additionally, there is a concern for the potential legal ramifications in the case of a patient that has a poor outcome under these circumstances, the responders feeling that there may be a potential liability in the case of a patient that may have had a relevant alert occur in the reactivation period, but did not due to the choice in reactivation time. Furthermore, there is the issue of inertia, and the tendency of humans to select the most readily available option, which is the Reactivate Now option [42]. As a result, the convenience of choosing the first option that is seen may be culprit for the less than optimal choice in alert reactivation time.

In this work, a method for predicting the choices of the eICU responders has also been determined, which suggests that one can develop decision support tools that can have a default reactivation time value based on the available parameters of the alert incidence and will select the appropriate response based on the actions previously performed by the doctors in similar situations as well as based on the prescribed medicine and amount of time typically required to see relevant changes in the patient given their condition.

Decision support tools of this nature may serve to hasten the process of responding to alerts and ensure that the actions of the care providers are all the more efficient, and reduce alarm fatigue. This is especially prescient considering that alarm fatigue is a widely reported problem in the literature [2, 38, 42, 54], and it is reasonable to expect that the telemedicine systems will generate the highest volume alarms for the most deadly illnesses, and thus, will be likelier to cause mistakes and errors on the part of the providers due to the effects of alarm fatigue. As mentioned earlier, some studies state that more than 60 percent of clinical alarms are false, which makes the work of the responding personnel rather difficult, in terms of sorting through the alerts [53].

The completed work took DRS information into account, predicting the incidence of alerts that coincided with a significant change in patient status. This change is an improvement upon the previous development system that had alerts activate blindly, that is, with no consideration of new developments on the part of the patient, resulting in repetitive alerts, which are well known to result in fatigue. Additionally, the algorithm took into account potential change in the severity and type of the alert, predicting when a different type of alert would be activated.

7.2 Future Work

7.2.1 Care Provider Response Analysis

In prior chapters, a compilation was made of the types of responses the care providers made to the different types of alerts that occur in a Tele-ICU environment. In this case, the response was measured by the medications requested by the care providers at the point in time following an alert. Naturally, medication is not the only way that medical personnel can respond to an alert, and future examinations of this field can consider various softer approaches such as communication from eICU to bedside or verbal communication to patients as other forms of response that can be performed as the result of an alert activation.

Naturally, such an analysis may require data that is not necessarily logged in the default software.

Another contribution could be the development of visualization tools relating to eICU workflow. This would serve the professionals in the eICUs and at Philips. Some examples of visualizations may include timelines of patient activity and caregiver response, which would demonstrate the evolution of a particular patient?s case. Another visualization may be one that displays the workflow for the various personnel in the eICU, showing the contribution and collaboration present between the members of an eICU team in their response to an alert and how they work toward their common goal.

In the work of this thesis, actions taken by the bedside professionals was not accounted for. The bedside receives alerts that are more immediate in nature, often referencing a particular physiological signal falling below a set value. For example, a patient may suffer from an oxygen saturation below the set alarm. in contrast, the eICU alarms intended to reference the evolution of the biological parameter the last 10 minutes. Prior research has suggested that practicing bedside nurses that work in telemedicine environments are optimistic about the benefits of the technology, however, they are concerned about the lack of familiarity and privacy that may be present in such a configuration [64,73, 89]. As a result, possible future innovations can take the bedside into account, especially where the determination of the best practices of the combination of bedside and remote monitoring in the eICU is concerned. Different hospitals have different configurations for their eICUs and it may be fruitful to examine these on a case-by-case basis in order to find the best methods available for having the two forces work in tandem.

Innovations in these fields may have particular relevance to the unique situations encountered by rural hospitals, in which the resources made available by telemedicine may serve well to allow certain rarer medical specialties to be accessed by those who do not have such personnel on-site, especially in remote areas [76, 86]. As a result, improvements in the workflow between the personnel located in person and at eiCUs may serve to shore up any gaps in the availability in healthcare, especially in critical care situations.

7.2.2 Reactivation Prediction

In the work previously discussed, an alert system was described that takes into account the patient status in the form of the DRS score, using it to predict the next alert occurrence that corresponds with a change in patient status. In future iterations of this algorithm, other measures can be included in the determination of patient status, such as change in measured vital signs, and consideration of past patterns of alerts in order to refine this calculation.

One significant possibility for the future of analytics in this domain is the usage of patterns of alert to act as a diagnostic tool. In this case, an algorithm can take into account a particular sequence of alerts occurring and use that information to predict a diagnosis based on prior knowledge as well as an understanding of previous instances of this type of occurrence. For example, if a patient's heart rate instantly falls to zero, and there is no notable change in the other measures such as blood pressure, one may assume that the lead that measures the blood pressure has fallen off. Similar technology has been utilized in the prediction of sepsis in patients [55], in which the measured biological signals from the patients are used to synthesize a real-time assessment of a patient's vulnerability to sepsis. It is possible that a similar method can be carried out for other specific diseases or issues in the critical care environment. Sepsis is the most expensive disease to treat in hospital [56], and other costly medical issues can be the basis for future advancements, such as osteoarthritis (listed as 2nd most expensive), or complication of device, implant or graft (listed as 4th most expensive).

Sepsis in particular is responsible for costs over 20 billion dollars each year in treatment, making it the most expensive diagnosis for hospitals and it carries with it an average mortality rate of 45 percent [92]. Other eICU collaboratives have made efforts to combat the occurrences of sepsis in hospitals. One project in particular was intended to decrease mortality rates for septic patients from 41.2 to 18.8 percent [93]. The method relied on technological as well as organization innovation, focusing on the methods in which an interdisciplinary team can best work in tandem in order to realize their goals. Furthermore, the Clinical Nurse Leader at this particular hospital was expected to act as a liaison between members of the health care team, and to ensure effective communication with a bedside response team consisting of nurses, physicians, pharmacists and laboratory personnel in order to manage necessary early treatment in this vulnerable patient population.

One may imagine that similar efforts may be undertaken in the future to combat the most lethal and costly diseases and complications that may occur in an eICU environment. Developing a set of best practices for preventing the occurrence of these types of maladies may be a fruitful endeavor.

Additionally, further examination can be made in the field of eICU as it related to increasing access to healthcare, especially in rural and underserved communities around the world. Most ICU Units are staffed by only one consultant, who cannot be expected to be responsible for the unit at all hours of the day. Night cover duty doctors are also present but many lack adequate expertise in the necessary specialities [79]. This lack of personnel is all the more prescient in emergency situations. In another effort to improve eICU outcomes, a hospital conducted an analysis of the difference in patient outcomes in a CCU (coronary care unit), before and after a remote monitoring implementation was added [94]. In this case, in the six months before the installation of the eICU, there were 18 out of 90 fatalities of patients that had entered due to myocardial infarction, as opposed to 4 of 92 in the six months following the addition of the eICU. In this case, the severity of the patient cases and the mean patient ages were mainly the same. As a result, when availability to healthcare is limited by a lack of expert manpower the EICU can potentially shore up this gap and provide support in order to minimize casualties.

Another possible future application in the field of telemedicine may be the inclusion of decision support algorithms in the remote monitoring process. Naturally, one can imagine the technology and data present in the eICU used in a capacity to assist the making of critical decisions regarding patient care, both to improve the quality of care as received by the patient and to ensure that the resources available in the hospital are being utilized in as efficient a manner as possible. APACHE based patient acuity predictions have already been employed to great effect, seeing a reduction in the usage of invasive procedures performed on patient as well as a decrease in length of stay, critical care complications and mortality [95]. Additional progress can be made in the utilization of the various alerts and physiological signals in order to determine the best course of actions for the care of the patient. Other publications have made mention of big data and how it can be used in the process of responding to patient needs [96].

An innovation that would make use of the data available would be the implementation of a natural language processor for the purpose of developing notes that the eICU clinicians (as well as perhaps the bedside professionals) could make use of in order to efficiently change shifts. Improving handoff between shifts has been addressed before. In one such endeavor, a patient admitted to the hospital experienced severe hypoglycemia having received an insulin dose that was greater than the prescribed amount [99]. Had the insulin syringe and the dosage been checked, this critical error in care would have likely been averted. Research has been conducted with regarding to standardizing handoff and correcting errors in the system [100], though a universally standardized protocol has not yet been determined. Structured tools, however, have been designed for the purposes of transferring patients in a postoperative fashion, and one can imagine that similar tools can be designed for the purposes of shift change in the eICU [101]. Furthermore, such a development would be useful in the case of transfers between different departments within a hospital. In order further elaborate on the development of a standardized handoff tool, interviews with relevant professionals must be conducted so that the ideal characteristics of handoff can be distilled.

Chapter 8

References

1. Grossman DC et al., Urban-rural differences in prehospital care of major trauma, J Trauma. 1997 Apr;42(4):723-9.

2. Kizzier-Carnahan et al., Frequency of Passive EHR Alerts in the ICU: Another Form of Alert Fatigue? J Patient Saf. 2016 Jun 22.

3. Thompson H, Philip's Orb Offers Hospitals A New Vision for ICU Assessments, Medical Software, March 7, 2013.

4. A.A. Miranda, et al., New Routes from Minimal Approximation Error to Principal Components, Volume 27, Number 3 / June, 2008, Neural Processing Letters, Springer.

5. Pampel FC, Logistic regression: A primer. Sage University Papers Series on Quantitative Applications in the Social Sciences, 07-132. Thousand Oaks, CA: Sage Publications, 2000.

6. Auffarth, B. Clustering by a Genetic Algorithm with Biased Mutation Operator. Wcci Cec. IEEE. . July 18 23, 2010.

7. Faisandier, A. Systems Architecture and Design. Belberaud, France: Sinergy'Com. 2012.

8. Department of Defense Systems Management College. System Engineering Fundamentals Defense Acquisition University Press. 2001.

9. A. E. W. Johnson, M. M. Ghassemi, S. Nemati, K. E. Niehaus, D. A. Clifton and G. D. Clifford, Machine Learning and Decision Support in Critical Care, in Proceedings of the IEEE, vol. 104, no. 2, pp. 444-466, Feb. 2016.

10. Awad et al. Patient Length of Stay and Mortality Prediction: a survey, University of Portmouth. 2016.

11. Hawkins, Helen A. et al., ICU Telemedicine Comanagement Methods and Length of Stay. CHEST, Volume 150, Issue 2, 314 - 319.

12. Lilly CM, McLaughlin JM, Zhao H, et al. A multicenter study of ICU telemedicine reengineering of adult critical care. Chest 2014;145:5007.

13. MIMIC-III, a freely accessible critical care database. Johnson AEW, Pollard TJ, Shen L, Lehman L, Feng M, Ghassemi M, Moody B, Szolovits P, Celi LA, and Mark RG. Scientific Data (2016). DOI: 10.1038/sdata.2016.35.

14. Nassar BS, Vaughan-Sarrazin MS, Jiang L, et al. Impact of an intensive care unit telemedicine program on patient outcomes in an integrated health care system. JAMA Intern Med 2014;174(7):11607.

15. Schoch, O.D. et al. 114 Telemetrically triggered interventions in the first month of CPAP: 6-months results of a prospective, randomized controlled trial, CHEST, Volume 151, Issue 5 , A10 - A11

16. Palagiri, Ashok Gudmestad, Donna Lakshmanan, Rekha Veremakis, Christopher Hassan, Erkan Badawi, Omar. (2013). Telemedicine utilization of a Discharge Readiness Score may facilitate safer ICU discharges.

17. Zou et al. Correlation and Simple Linear Regression, 10.1148/radiol.2273011499 Radiology 2003; 227:617 628.

18. Martin et al. A systematic review of the factors that influence the quality and effectiveness of telesupervision for health professionals. Journal of Telemedicine and Telecare.

19. Machado, Shana et al., Utilization of eICU Telemedicine Carts in the Emergency Department. CHEST, Volume 150, Issue 4, 363A.

20. Lilly, Craig M. et al., Five-Year Trends of Critical Care Practice and Outcomes, CHEST, Volume 152, Issue 4, 723 - 735.

21. Awad, Omar et al., Patient's Acceptance for the Advanced Monitoring Systems in Electronic Intensive Care Units (eICU), CHEST, Volume 132, Issue 4, 444B.

22. Mukherjee, Sumit et al., Closing the Door to Tele-ICU: Analysis of Outcomes Before and After Closing a Single ICU to Tele-Medicine CHEST, Volume 148, Issue 4, 222A.

23. Zawada, Edward et al., Workload Update on the Most Rural Tele-Intensive Care Program, CHEST, Volume 142, Issue 4, 548A.

24. Halpern, Neil A., From Telemedicine to a Critical Care Database, CHEST, Volume 140, Issue 5, 1111 - 1113.

25. Hawkins, Helen A. et al., ICU Telemedicine Comanagement Methods and Length of Stay, CHEST, Volume 150, Issue 2, 314 - 319.

26. Garland, Allan et al., Staffing in ICUs CHEST, Volume 143, Issue 1, 214 - 221.

27. Kerlin, Meeta Prasad et al., Nighttime Intensivist Staffing, Mortality, and Limits on Life Support, CHEST, Volume 147, Issue 4, 951 - 958.

28. Venkataraman, Ramesh et al., Outcomes Related to Telemedicine in the Intensive Care Unit, Critical Care Clinics, Volume 31, Issue 2, 225 - 237.

29. Hoffmann, Stephen, Counterpoint: Should Tele-ICU Services Be Eligible for Professional Fee Billing? No, CHEST, Volume 140, Issue 4, 849 - 851.

30. Critical Care: ICU Monitoring/Diagnostic Techniques. CHEST, Volume 114, Issue 4, 323S - 326S.

31. Schoch, O.D. et al., 114 Telemetrically triggered interventions in the first month of CPAP: 6-months results of a prospective, randomized controlled trial. CHEST , Volume 151 , Issue 5 , A10 - A11.

32. Santos, Carolina et al. Telenursing Protocol to the Individual With Pulmonary Dysfunction and Heart Failure. CHEST, Volume 145, Issue 3, 202A.

33. Badawi O, Breslow MJ (2012) Readmissions and Death after ICU Discharge: Development and Validation of Two Predictive Models. PLoS ONE 7(11): e48758. https://doi.org/10.1371/journal.pone.0048758.

34. Aaronson, Michael L. et al., Role of a Telemedicine Intensive Care Unit Program On Glycemic Control in Seriously Ill Patients in a Rural Health System, CHEST, Volume 130 , Issue 4 , 226S.

35. Kelley, Mark A. et al. The Critical Care Crisis in the United States. CHEST , Volume 125 , Issue 4 , 1514 - 1517.

36. Ahn, Young, et al. Perspectives on the Electronic ICU. ICU Director, vol. 3, no. 2, 2012, pp. 64 69., doi:10.1177/1944451612440721. Cummings, Joseph, et al. Intensive Care Unit Telemedicine: Review and Consensus Recommendations. American Journal of Medical Quality, vol. 22, no. 4, 2007, pp. 239 250., doi:10.1177/1062860607302777.

37. Henderson-Mclennan, Leon. A Day in the Life of a Tele-Intensivist. ICU Director, vol. 1, no. 1, 2010, pp. 35 37., doi:10.1177/1944451609354182.

38. Kapoor, Sumit, et al. Sacred Pause in the ICU: Evaluation of a Ritual and Intervention to Lower Distress and Burnout. American Journal of Hospice and Palliative Medicineő, Apr. 2018, p. 104990911876824., doi:10.1177/1049909118768247.

39. Liu, Fy, et al. A Prospective Validation of National Early Warning Score in Emergency Intensive Care Unit Patients at Beijing. Hong Kong Journal of Emergency Medicine, vol. 22, no. 3, 2015, pp. 137-144., doi:10.1177/102490791502200302.

40. Marcin, James P., et al. The Role of Telemedicine in Treating the Critically Ill. ICU Director, vol. 3, no. 2, 2012, pp. 70-74., doi:10.1177/1944451612439207.

41. Martin, Priya, et al. A Systematic Review of the Factors That Influence the Quality and Effectiveness of Telesupervision for Health Professionals. Journal of Telemedicine and Telecare, vol. 24, no. 4, July 2017, pp. 271 281., doi:10.1177/1357633x17698868.

42. Shimabukuro, David. Acute Care Nurse Practitioners in an Academic Multidisciplinary ICU. ICU Director, vol. 2, no. 1-2, 2011, pp. 28-30., doi:10.1177/1944451611405049.

43. Shrope-Mok, Shaffer R., et al. APACHE IV Versus PPI for Predicting Community Hospital ICU Mortality. American Journal of Hospice and Palliative Medicineő, vol. 27, no. 4, Mar. 2009, pp. 243?247., doi:10.1177/1049909109350177.

44. Angus, DC, Kelley, MA, Schmitz, RJ et al, Current and projected workforce requirements for care of the critically ill and patients with pulmonary disease: can we meet the requirements of an aging population?. JAMA. 2000;284:2762 2800.

45. Halpern, Neil A. Innovative Designs for the Smart ICU. CHEST, Volume 145, Issue 4, 903 - 912.

46. Devereaux, Asha V. et al. Engagement and Education. CHEST, Volume 146, Issue 4, e118S - e133S.

47. Bhattacharjee, Poushali et al. Identifying Patients With Sepsis on the Hospital Wards. CHEST, Volume 151, Issue 4, 898 - 907.

48. Parthasarathy, Sairam et al. Sleep and Well-Being of ICU Housestaff. CHEST, Volume 131, Issue 6, 1685 - 1693.

49. Sanchez-Pinto, L. Nelson et al. Big Data and Data Science in Critical Care. CHEST, Volume 0 , Issue 0.

50. Kerlin, Meeta Prasad et al. Twenty-four-Hour Intensivist Staffing in Teaching Hospitals. CHEST, Volume 141, Issue 5, 1315 - 1320.

51. Simmons, Robin L. et al. I See Heart Failure?: Tele-ICU eNurse Alerts Making a Difference. CHEST, Volume 138, Issue 4, 183A.

52. Yau, Simon et al. Bag Score: An ICU Survival Scoring System in a Veteran Population. CHEST, Volume 150, Issue 4, 214A.

53. Cho, Ok Min et al. Clinical Alarms in Intensive Care Units: Perceived Obstacles of Alarm Management and Alarm Fatigue in Nurses. Healthcare Informatics Research 22.1 (2016): 46-53. PMC. Web. 14 May 2018.

54. Bridi, Adriana Carla, Thiago Quinellato Louro, and Roberto Carlos Lyra da Silva. Clinical Alarms in Intensive Care: Implications of Alarm Fatigue for the Safety of Patients . Revista Latino-Americana de Enfermagem 22.6 (2014): 1034-1040. PMC. Web. 14 May 2018.

55. Makam, Anil N., Oanh K. Nguyen, and Andrew D. Auerbach. Diagnostic Accuracy and Effectiveness of Automated Electronic Sepsis Alert Systems: A Systematic Review. Journal of hospital medicine 10.6 (2015): 396-402. PMC. Web. 14 May 2018.

56. National Inpatient Hospital Costs: The Most Expensive Conditions by Payer, 2013 204. Emergency Department Visits for Adults with Diabetes, 2010 - Statistical Brief 167, May 2016.

57. https://en.wikipedia.org/wiki/Evaluationofbinaryclassifiers

58. D. Bera and M. M. Nayak, "Mortality risk assessment for ICU patients using logistic regression," 2012 Computing in Cardiology, Krakow, 2012, pp. 493-496.

59. Knaus WA, Draper EA, Wagner DP Zimmerman JE. Apache II: a severity of disease classification system. Critical Care Medicine 1985; 13: 818 - 29.

60. Lemeshow S, Teres D, Pastides H, Avrunin JS Steingrub JS. A method for predicting survival and mortality of ICU patients using objectively derived weights. Critical Care Medicine 1985; 13: 519 - 25.

61. Institute of Medicine, Committee on Monitoring Access to Personal Health Care Services. Access to Health Care in America. Millman M, editor. Washington, DC: National Academies Press; 1993.

62. National Healthcare Quality Report, 2013 [Internet]. Chapter 10: Access to Healthcare. Rockville (MD): Agency for Healthcare Research and Quality; May 2014. Available from: http://www.ahrq.gov/research/findings/nhqrdr/nhqdr15/access.html.

63. Chung KK, Grathwohl KW, Poropatich RK, et al. Bedside nurse perceptions of intensive care unit telemedicine. Crit Care Med. 2009;37(12):A441.

64. Wilson L, Hansen SK, Skinner MW. Design and evaluation of a broadband telemedicine system for supporting ICUs in regional hospitals. Paper presented at: American Telemedicine Association Thirteenth Annual Meeting and Exposition. April 2008 Seattle, WA.

65. Rogove H, Atkins C, Kramer J, et al. Enhanced access to neurointensivists through a telemedicine program. Crit Care Med. 2009;37(12 suppl):A1.

66. Marttos A, Wilson K, Krauthamer S, et al. Telerounds in a trauma ICU dept. Crit Care Med. 2008;36(12):A53.

67. Creswell JW. Research Design: Qualitative, Quantitative, and Mixed Methods Approaches. Thousand Oaks, CA: Sage; 2009.

68. Tashakkori A, Teddlie C. Mixed Methodology: Combining Qualitative and Quantitative Approaches. Thousand Oaks, CA: Sage; 1998.

69. Society of Critical Care Medicine. Tele-ICU Monitoring and Care. Mount Prospect, IL: Society of Critical Care Medicine; 2009.

70. Thomas EJ, Lucke JF, Wueste L, Weavind L, Patel B. Association of telemedicine for remote monitoring of intensive care patients with mortality, complications, and length of stay. JAMA. 2009;302(24):2671?2678.

71. Duke EM. Report to Congress: The Critical Care Workforce: A Study of the Supply and Demand for Critical Care Physicians. Washington, DC: US Department of Health and Human Services HRSA; 2006.

72. Gutsche JT, Kohl BA. Who should care for intensive care unit patients? Crit Care Med. 2007;35:S18-S23.

73. Berenson RA, Grossman JM, November EA. Does telemonitoring of patients in the eICU improve intensive care? Health Aff (Millwood) 2009;28(5):w937-w947.

74. Breslow MJ. Remote ICU care programs: current status. J Crit Care. 2007;22(1):66 76.

75. Ali J, Sorvari A, Sandrine C, et al. Telemedicine as a potential medium for teaching the Advanced Trauma Life Support (ATLS) course. J Surg Educ. 2013;70(2):258 64.

76. Odetola FO, Davis MM, Cohn LM, et al. Interhospital transfer of critically ill and injured children: an evaluation of transfer patterns, resource utilization, and clinical outcomes. J Hosp Med. 2009;4:164 70.

77. Dayal P, Hojman NM, Kissee JL, et al. Impact of telemedicine on severity of illness and outcomes among children transferred from referring emergency departments to a children?s hospital PICU. Pediatric Crit Care Med. 2016.

78. Coletti C, Elliot DJ, Zubrow MT. Resident perceptions of a tele-intensive care unit implementation. Telemed J E Health. 2010;16:894 7.

79. James JT. A new, evidence-based estimate of patient harms associated with hospital care. J Patient Saf. 2013;9:122 8.

80. Breslow MJ , Rosenfeld BA, Doerfler M, Burke G, Yates G, Stone DJ, et al. Effect of a multiple-site intensive care unit telemedicine program on clinical and economic outcomes: an alternative paradigm for intensivist staffing. Crit Care Med . 2004 ; 32 (1):

81. Angus DC . Caring for the critically ill patient: challenges and opportunities. JAMA . 2007 ; 298 (4): 456 8.

82. Rubenfeld GD, Angus DC. Are intensivists safe? Editorial. Ann Intern Med . 2008 ; 148 (11): 877 9.

83. Rosenfeld BA , Dorman T, Breslow MJ, Pronovost P, Jenckes M, Zhang N, et al. Intensive care unit telemedicine: alternate paradigm for providing continuous intensivist care. Crit Care Med . 2000 ; 28 (12): 3925 31.

84. Angus DC , Kelley MA, Schmitz RJ, White A, Popovich J Jr. Current and projected workforce requirements for care of the critically ill and patients with pulmonary disease: can we meet the requirements of an aging population? JAMA . 2000 ; 284 (21): 2762 70.

85. Marcin JP, Nesbitt TS, Kallas HJ, Struve SN, Traugott CA, Dimand RJ: Use of telemedicine to provide pediatric critical care inpatient consultations to underserved rural Northern California. J Pediatr. 2004, 144: 375-380. 10.1016/j.jpeds.2003.12.017.

86. Kohl BA, Gutsche JT, Kim P, Sites FD, Ochroch EA: Effect of telemedicine on mortality and length of stay in a University hospital [abstract]. Crit Care Med. 2007, 35: A111.

87. Norman V, French R, Hassan E, Kabani N, Mizell P, Stone DJ, Griebel J, Tragico E: Effect of a telemedicine facilitated program on intensive care unit

length of stay (LOS) and financial performance [abstract]. Crit Care Med. 2009, 37: A3.

88. Rincon F, Vibbert M, Childs V, Fry R, Caliguri D, Urtecho J, Rosenwasser R, Jallo J: Implementation of a model of robotic tele-presence (RTP) in the Neuro-ICU: effect on critical care nursing team satisfaction. Neurocrit Care. 2012.

89. Young LB, Chan PS, Cram P: Staff acceptance of tele-ICU coverage: a systematic review. Chest. 2011, 139: 279-288. 10.1378/chest.10-1795.

90. Wilcox ME, Adhikari NK: Effect of telemedicine on mortality in critically ill patients: systematic review and meta-analysis. Am J Respir Crit Care Med. 2012, 185: A2867.

91. Afrefian, H., Heublein, S., Scherag, A., Brunkhorst, F. M., Younis, M. Z., Moerer, O., Hartmann, M. Hospital related cost of sepsis: A systematic review. Journal of Infection, 2017, 74(2), 107 117.

92. Harrison, A. M., Thongprayoon, C., Kashyap, R., Chute, C. G., Gajic, O., Pickering, B. W., Herasevich, V. Developing the surveillance algorithm for detection of failure to recognize and treat severe sepsis. Mayo Foundation for Medical Education and research. 2015. 90(2): 166 175.

93. Chronaki C, Shahin A, Mark R. Designing Reliable Cohorts of Cardiac Patients across MIMIC and eICU. Computing in cardiology. 2015;42:189 192. doi:10.1109/CIC.2015.7408618.

94. Yoo, E., Dudley, R. Evaluating telemedicine in the ICU. 2009. JAMA, 302, 2705-2706.

95. Yau, Simon et al. Bag Score: An ICU Survival Scoring System in a Veteran Population. CHEST, Volume 150, Issue 4, 214A.

96. Gupta Shamit; Varma, Amit; Dewan, Sandeep; Kaushal, Anjali. eICU Can Save Lives in Myocardial Infarction in Remove Areas in the Developing World. Critical Care Medicine: December 2013 doi: 10.1097/01.ccm.0000439420.28386.60.

97. Lilly, Craig M. et al. Five-Year Trends of Critical Care Practice and Outcomes. CHEST, Volume 152, Issue 4, pp. 723 735.

98. Sanchez-Pinto, L. Nelson et al. Big Data and Data Science in Critical Care. CHEST, Volume 0, Issue 0.

99. Berkenstadt, Haim et al. Improving Handoff Communications in Critical Care*: Utilizing Simulation-Based Training Toward Process Improvement in Managing Patient Risk. CHEST, Volume 134, Issue 1, 158 162.

100. Carroll, Christopher et al. Promoting a Culture of Safety in the Pediatric ICU by Improving the Quality of Postoperative Handoffs. CHEST, Volume 144, Issue 4, 406A.

101. Berkenstadt, Haim et al. Improving Handoff Communications in Critical Care: Utilizing Simulation-Based Training Toward Process Improvement in Managing Patient Risk. CHEST, Volume 134, Issue 1, 158 162.

102. Dunn, William et al. The Patient Handoff: Medicine's Formula One Moment. CHEST, Volume 134, Issue 1, 9 12.

103. Hess, Dean R. et al. The Value of Adding a Verbal Report to Written Handoffs on Early Readmission Following Prolonged Respiratory Failure. CHEST, Volume 138, Issue 6, 1475 1479.

104. Frattini, John et al. Patient Handover Protocol From Intensive Care Unit, Effect on ICU Readmission Rate, and Predictability of Clinical Variables in ICU Readmitted Patients. CHEST, Volume 144, Issue 4, 563A.

105. Shah, Koonj et al. Studying Current Handoff Practices Amongst Critical Care Fellows. CHEST, Volume 140, Issue 4, 337A.

106. Brown, Sydney E.S. et al. Assessing the Utility of ICU Readmissions as a Quality Metric. CHEST, Volume 147, Issue 3, 626 636.