OPTIMIZING THE ICU DURING EPIDEMICS: DEVELOPING FRAMEWORKS SUPPORTING HEALTHCARE SYSTEMS DURING PUBLIC HEALTH EMERGENCIES

by

Haleigh Noelle West-Page

A dissertation submitted to the faculty of
The University of North Carolina at Charlotte
in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in
Applied Mathematics

Charlotte

2026

Approved by:

Dr. Kevin McGoff

Dr. Xingjie Li

Dr. Duan Chen

Dr. Shi Chen

ABSTRACT

HALEIGH NOELLE WEST-PAGE. Optimizing the ICU during epidemics: developing frameworks supporting healthcare systems during public health emergencies. (Under the direction of DR. KEVIN MCGOFF)

Public health emergencies like the COVID-19 pandemic motivate the development of tools supporting the healthcare system. Mathematical modeling plays a role in creating these frameworks aiding the healthcare system at multiple levels. In this work, we explore how different models can optimize workflow and resources through the intensive care unit in the face of an epidemic.

Decision support tools can provide clinicians with augmented views of their patients, allowing them to deliver precise medical care even during an emergency. At the clinician level, we leverage machine learning to create decision support tools in aiding patient classification. Using patient-level data, we explore the predictability and interpretability of various machine learning models in their ability to classify COVID-19 severity types. These models provide useful predictions and insight on severely ill patients with COVID-19.

Public health emergencies place high levels of burden on the healthcare system, especially the intensive care unit. Aiming to support the systemic level, we develop an advanced framework simulating the intensive care unit under epidemic conditions. Investigating the tradeoff between ICU capacity and clinician staffing levels when constrained by an operating budget, we demonstrate a robust modeling foundation for optimizing critical care resources. This flexible framework provides the space for rich analysis of a diverse set of healthcare systems under various scenarios.

DEDICATION

To my husband Jacob- who was my guiding light when I was lost at sea.

To my family- to whom I promised I would someday be a doctor.

To all of the nerds at Baldwin, Hope Mills, and South View- we ran the table.

Power to the local dreamer |-/

ACKNOWLEDGEMENTS

I am deeply grateful to my dissertation chair, Dr. Kevin McGoff, for playing a decisive role in my path to success. I will forever be thankful for his mathematical insight, thoughtful modeling discussions, and unwavering support throughout this journey. I would like to express my sincerest gratitude to my committee member, Dr. Shi Chen, for his expertise in public health and epidemiology. I am so thankful for his contribution in supervising the applications of this work and for his continued support. I extend a special thanks to my committee members, Dr. Xingjie "Helen" Li and Dr. Duan Chen for their wise questions, comments, and suggestions. Their perspectives have proved to be invaluable in pushing this dissertation to its fullest potential.

I would like to acknowledge the funding sources that made this dissertation possible. I thank the support of the Graduate Assistant Support Plan (GASP) provided by the UNC Charlotte Graduate School. I also thank the support of U.S. Centers for Disease Control and Prevention: Building Mathematical Modeling Workforce Capacity to Support Infectious Disease and Healthcare Research. The project described was supported by cooperative agreement (U01CK000677) from U.S. Centers for Disease Control and Prevention (CDC). Its contents are solely the responsibility of the authors and do not necessarily represent the official views of the CDC.

TABLE OF CONTENTS

LIST OF TABLES x

LIST OF FIGURES xi

LIST OF ABBREVIATIONS xvii

CHAPTER 1: INTRODUCTION 1

1.1. Multi-feature clinical decision support 1

1.2. Modeling emergency preparedness in the healthcare system 2

1.3. Structure of dissertation 4

CHAPTER 2: MULTI-FEATURE CLINICAL DECISION SUPPORT 5

2.1. Introduction 5

2.2. Methods 7

2.3. Results 14

2.4. Discussion 19

2.5. Conclusion 26

CHAPTER 3: MODELING EMERGENCY PREPAREDNESS IN THE HEALTHCARE SYSTEM 27

3.1. Introduction 27

3.2. Methods 29

3.3. Results 50

3.4. Discussion 71

3.5. Conclusion 84

CHAPTER 4: CONCLUSIONS 85

REFERENCES 86

APPENDIX A: MULTI-FEATURE CLINICAL DECISION SUPPORT 96

APPENDIX B: MODELING EMERGENCY PREPAREDNESS IN THE HEALTHCARE SYSTEM 109

LIST OF TABLES

TABLE 2.1: Area under the Receiver Operating Characteristic Curve (AUC) scores 14

TABLE 3.1: SEIR compartment transition rate parameters and descriptions 32

TABLE 3.2: Table of model parameters 46

TABLE 3.3: Course sensitivity analysis for the low resource setting 51

TABLE 3.4: Course sensitivity analysis for the high resource setting 51

TABLE A.1: Standard Deviations of Area under the Receiver Operating Characteristic Curve (AUC) 96

TABLE A.2: Imported packages for ML pipeline 96

TABLE A.3: Frequent hyperparameters selected during ML model training by variant and data modality 97

LIST OF FIGURES

FIGURE 2.1: De-identified patient biomedical data were collected from Wuhan Union Hospital from January to March 2020 (original COVID-19 patients n=362) and December 2022 to January 2023 (Omicron variant patients n=1000). The data were grouped by their feature modality: biochemical, clinical or fusion. An 80:20 split was performed to generate training and testing sets, respectively. Four classifiers were developed and evaluated for their performance: LR=logistic regression, RF=random forest, kNN=k-nearest neighbors, and SVM=support vector machine. Upon hyperparameter tuning, each classifier was validated with the same-variant hold-out set, as well as the preprocessed cross-variant set. 11

FIGURE 2.2: This figure shows the mean ROC curves plotted for each of the four training-testing combinations. Each figure contains the ROC curve of each data modality (biochemical-red, clinical-green, and fusion-blue.) Panel (a) plots the curve of classifiers trained and tested on original data only (i.e., original-original combination). Panel (b) plots the curve of classifiers of the original-Omicron combination. Panel (c) plots the curve of classifiers of the Omicron-original combination. Panel (d) plots the curve of classifiers of the Omicron-Omicron combination. 16

FIGURE 2.3: This panel presents the feature importance rankings of the logistic regression (LR) and random forest (RF) models. Panel (a) plots feature weigh coefficients for LR when trained on Original COVID-19 patient data. Panel (b) plots feature weigh coefficients for LR when trained on Omicron COVID-19 patient data. Panel (c) plots root features for RF when trained on Original COVID-19 patient data. Panel (d) plots root features for RF when trained on Omicron COVID-19 patient data. 18

FIGURE 3.1: This visualization of the queueing system displays the flow of patients through our simulated ICU. Critically ill patients are requested for admission either due to an epidemic related illness ($\lambda_{\text{epidemic}}$) or non-epidemic related illness ($\lambda_{\text{baseline}}$), such as surgery or trauma. Patients are placed in an ICU bed whenever possible. If no beds are available, patients wait in an unlimited length queue. Patients may leave the queue without ICU treatment (ρ) due to recovery in a general ward, transferring to another facility, or suffering mortality. Patients in the ICU depart at a rate relative to their patient type (μ_{epidemic} or μ_{baseline}). Epidemic-related patients are considered infectious and may infect healthcare providers (η). If enough healthcare providers are ill and unable to work, the capacity of the ICU is reduced until they return (ν). 40

FIGURE 3.2: This figure shows the simulated proportion of the community population infected with a hypothetical novel pathogen over time. Scenario B represents a pathogen with a basic reproduction number $R_0 = 2.73$ (left). Scenario C represents a pathogen with a basic reproduction number $R_0 = 3.28$ (right). 43

FIGURE 3.3: Scenario A represents a pre-epidemic baseline where there is no wave of infectious patients. Scenario B represents an outbreak of a hypothetical novel pathogen, simulating a wave of critically ill infectious patients. Scenario C represents an outbreak of a hypothetical novel pathogen with increased transmissibility. For A, B, and C, these scenarios are applied to two simulated facility settings, one with a high amount of resources and one with a low amount of resources. 44

FIGURE 3.4: Fine-grain sensitivity analysis of baseline arrival rate, baseline departure rate, and epidemic departure rate 53

FIGURE 3.5: This figure displays the heatmaps of accumulated queue abandonments simulated over 365 days of a low resource facility in three case scenarios. Scenario A represents a pre-epidemic baseline in which there is no infectious disease among the community population. Scenario B represents an outbreak of a novel respiratory pathogen. Scenario C represents an outbreak of a novel respiratory pathogen with a 20% increase in transmissibility. Areas of dark green indicate the relative lowest accumulated queue abandonments, while areas of dark red indicate the highest. 55

FIGURE 3.6: This figure displays the simulated behavior of a Markovian queueing model of an ICU with 10 beds. This low resource setting is simulated under the conditions of Scenario A, the pre-epidemic baseline case. The first subplot graphs the mean (solid red) and individual sample paths tracking the number of patients in the queue. The second subplot graphs the mean (solid green) and individual sample paths tracking the number of total patients in the ICU. The third subplot graphs the mean (solid blue) and individual sample paths tracking the number of total accumulated queue abandonments in the ICU. The fourth subplot graphs the mean (solid purple) number of active clinicians assigned to the ICU. 57

FIGURE 3.7: This figure displays the simulated behavior of a Markovian queueing model of an ICU with 10 beds. This low resource setting is simulated under the conditions of Scenario B, the initial outbreak case. The first subplot graphs the mean number of total patients in the queue (solid red), as well as mean number of infectious patients (dashed coral). The second subplot graphs the mean number of total patients in the ICU (solid green), as well as mean number of infectious patients (dashed light green). The third subplot graphs the mean number of total accumulated queue abandonments (solid blue), as well as mean number of accumulated abandonments who were infectious patients (dashed light blue). The fourth subplot graphs the mean (solid purple) number of active clinicians assigned to the ICU. 59

FIGURE 3.8: This figure displays the simulated behavior of a Markovian queueing model of an ICU with 10 beds. This low resource setting is simulated under the conditions of Scenario C, the initial outbreak case with higher transmissibility. The first subplot graphs the mean number of total patients in the queue (solid red), as well as mean number of infectious patients (dashed coral). The second subplot graphs the mean number of total patients in the ICU (solid green), as well as mean number of infectious patients (dashed light green). The third subplot graphs the mean number of total accumulated queue abandonments (solid blue), as well as mean number of accumulated abandonments who were infectious patients (dashed light blue). The fourth subplot graphs the mean (solid purple) number of active clinicians assigned to the ICU. 61

FIGURE 3.9: This figure displays the heatmaps of accumulated queue abandonments simulated over 365 days of a high resource facility in three case scenarios. Scenario A represents a pre-epidemic baseline in which there is no infectious disease among the community population. Scenario B represents an outbreak of a novel respiratory pathogen. Scenario C represents an outbreak of a novel respiratory pathogen with a 20% increase in transmissibility. Areas of dark green indicate the relative lowest accumulated queue abandonments, while areas of dark red indicate the highest. 63

FIGURE 3.10: This figure displays the simulated behavior of a Markovian queueing model of an ICU with 29 beds. This high resource setting is simulated under the conditions of Scenario A, the pre-epidemic baseline case. The first subplot graphs the mean (solid red) and individual sample paths tracking the number of patients in the queue. The second subplot graphs the mean (solid green) and individual sample paths tracking the number of total patients in the ICU. The third subplot graphs the mean (solid blue) and individual sample paths tracking the number of total accumulated queue abandonments in the ICU. The fourth subplot graphs the mean (solid purple) number of active clinicians assigned to the ICU. 64

FIGURE 3.11: This figure displays the simulated behavior of a Markovian queueing model of an ICU with 29 beds. This high resource setting is simulated under the conditions of Scenario B, the initial outbreak case. The first subplot graphs the mean number of total patients in the queue (solid red), as well as mean number of infectious patients (dashed coral). The second subplot graphs the mean number of total patients in the ICU (solid green), as well as mean number of infectious patients (dashed light green). The third subplot graphs the mean number of total accumulated queue abandonments (solid blue), as well as mean number of accumulated abandonments who were infectious patients (dashed light blue). The fourth subplot graphs the mean (solid purple) number of active clinicians assigned to the ICU. 66

FIGURE 3.12: This figure displays the simulated behavior of a Markovian queueing model of an ICU with 29 beds. This high resource setting is simulated under the conditions of Scenario C, the initial outbreak case with higher transmissibility. The first subplot graphs the mean number of total patients in the queue (solid red), as well as mean number of infectious patients (dashed coral). The second subplot graphs the mean number of total patients in the ICU (solid green), as well as mean number of infectious patients (dashed light green). The third subplot graphs the mean number of total accumulated queue abandonments (solid blue), as well as mean number of accumulated abandonments who were infectious patients (dashed light blue). The fourth subplot graphs the mean (solid purple) number of active clinicians assigned to the ICU. 68

FIGURE 3.13: Summary of heatmaps displaying the mean accumulated queue abandonments for a diverse array of resource allocations at both the low resource setting (top) and high resource setting (bottom). 69

FIGURE 3.14: Summary of single queueing models with the top row representing the low resource setting and the bottom row representing the high resource setting. For either facility, scenario A (baseline) is left, scenario B (initial outbreak) is middle, and scenario C (increased transmissibility) is right. 70

FIGURE A.1: ROC plots from Logistic Regression (LR) models 98

FIGURE A.2: ROC plots from k-Nearest Neighbors (kNN) models 99

FIGURE A.3: ROC plots from Support Vector Machine (SVM) models 100

FIGURE A.4: Feature importances identified by LR model trained on original biochemical data only 101

FIGURE A.5: Feature importances identified by LR model trained on original clinical data only 102

FIGURE A.6: Feature importances identified by LR model trained on Omicron biochemical data only 103

FIGURE A.7: Feature importances identified by LR model trained on Omicron clinical data only 104

FIGURE A.8: Feature importances identified by RF model trained on original biochemical data only 105

FIGURE A.9: Feature importances identified by RF model trained on original clinical data only 106

FIGURE A.10: Feature importances identified by RF model trained on Omicron biochemical data only 107

FIGURE A.11: Feature importances identified by RF model trained on Omicron clinical data only 108

LIST OF ABBREVIATIONS

kNN An acronym for k-nearest neighbors.

AUC An acronym for area under the curve.

CDSS An acronym for clinical decision support system.

CTMC An acronym for continuous time Markov chain.

FD An acronym for frequency-dependent transmission.

ICU An acronym for intensive care unit.

LOS An acronym for length of stay.

LR An acronym for logistic regression.

ML An acronym for machine learning.

QT An acronym for queueing theory.

RF An acronym for decision tree based random forest.

ROC An acronym for receiver operating characteristic curve.

SD An acronym for standard deviation.

SEIR An acronym for the Susceptible-Exposed-Infected-Recovered/Removed disease transmission model.

SVM An acronym for support vector machines.

USD An acronym for United States dollar.

CHAPTER 1: INTRODUCTION

Public health emergencies like the COVID-19 pandemic motivate the development of tools supporting the healthcare system. Mathematical modeling plays a role in creating these frameworks aiding the healthcare system at multiple levels. In this work, we explore how different models can optimize workflow and resources through the intensive care unit in the face of an epidemic.

1.1 Multi-feature clinical decision support

Accurately differentiating severe from non-severe COVID-19 clinical types is critical for the healthcare system to optimize workflow. Current techniques lack the ability to accurately predict COVID-19 patients' clinical type, especially as SARS-CoV-2 continues to mutate. We explore predictability and interpretability of multiple state-of-the-art machine learning (ML) techniques trained and tested under different biomedical data types and COVID-19 variants.

Comprehensive patient-level data were collected from 362 patients (148 severe, 214 non-severe) with the original SARS-CoV-2 variant in 2020 and 1000 patients (500 severe, 500 non-severe) with the Omicron variant in 2022-2023. The data included 26 biochemical features from blood testing and 26 clinical features from patients' clinical characteristics and medical history. Different ML techniques, including penalized logistic regression (LR), random forest (RF), k-nearest neighbors (kNN), and support vector machines (SVM), were applied to build predictive models based on each data

modality separately and together for each variant. Fifty randomized train-test-splits were conducted per scenario and performance results were recorded.

The fusion (hybrid) characteristic modality yielded the highest mean area under the curve (AUC), achieving 0.915, while the biochemical modality alone and the clinical modality alone had AUCs of 0.862 and 0.818, respectively. All ML models performed similarly under different testing scenarios and were robust when cross-tested with original and Omicron variant patient data. Our models ranked elevated d-dimer (biochemical), elevated high sensitivity troponin I (biochemical), and age greater than 55 years (clinical) as the most predictive features of severe COVID-19.

ML is a useful tool for predicting severe COVID-19 based on comprehensive individual patient-level data. Further, ML models trained on the biochemical and clinical modalities together witness enhanced predictive power. The improved performance of these ML models when trained and cross-tested with Omicron variant data suggests that ML may be used as a clinical decision support tool even as the virus continues to evolve.

1.2 Modeling emergency preparedness in the healthcare system

Public health emergencies place high levels of burden on the healthcare system, especially the intensive care unit. Aiming to support the systemic level, we develop an advanced framework simulating the intensive care unit under epidemic conditions.

Healthcare systems are significantly impacted by health emergencies. During the COVID-19 pandemic intensive care units (ICUs) experienced sudden high utilization rates, resulting in the need to optimize workflow and resources. We explore the queueing theoretic framework for modeling the ICU under different scenarios. We

construct an expanded queueing theoretic model of the ICU with operating budget constraints to find optimal resource allocation in various scenarios.

Assuming a Markovian arrival and departure process with reneging, we establish a queueing theoretic model of an ICU. We expand the model by endowing it with dynamic arrival and departure rates relevant to an epidemic. To this aim, we incorporate a deterministic disease transmission model of a novel pathogen through the simulated community population surrounding the ICU. We add healthcare providers as a secondary resource to the ICU who are susceptible to the novel pathogen. After parameterizing the model with values found in published literature, we simulated two facilities with different resource levels under various epidemic scenarios.

This framework may be able to identify optimal allocations of ICU clinicians and beds when constrained by an operational budget under various conditions. This framework can estimate the burden on an ICU assuming non-emergency conditions, as well as two distinct community outbreak scenarios. Our model captures trends for both low and high resource settings, including projected queue length, ICU utilization, ICU capacity fluctuations, queue abandonment, and clinician illness. We compare these trends between scenarios aiming to quantify the effect of the community epidemic on the ICU.

This expanded queueing model provides a rich basis for a realistic model of the ICU. Upon calibration and validation with real world data, this tool can provide insight to hospital administrators and public health policymakers evaluating healthcare systems for emergency preparedness.

1.3 Structure of dissertation

This dissertation closely details the relevant work, methods, results, and discussion of the two projects introduced above. The work done for creating data-driven clinical decision support is explored in Chapter 2. The work done for establishing a queueing-theoretic framework of the ICU to conduct scenario modeling is found in Chapter 3. Final discussions and conclusions are drawn in Chapter 4.

CHAPTER 2: MULTI-FEATURE CLINICAL DECISION SUPPORT

2.1 Introduction

The COVID-19 pandemic caused by the SARS-CoV-2 virus has impacted healthcare systems everywhere. Since 2019, several major SARS-CoV-2 variants and subvariants have manifested, with the Omicron variant being the most persistent since November 2021 [1]. A critical effect of the pandemic has been the sudden increased burden on healthcare facilities, mostly hospitals. The influx of severe COVID-19 patients overwhelms intensive care units which results in increased mortality [2], especially in regions with fewer health resources [3, 4].

In current clinical practice, patients with COVID-19 are typically classified as having severe illness by features such as shortness of breath, low oxygen saturation, and low PaO2/fraction of inspired oxygen. However these few features cannot sufficiently distinguish between severe and non-severe types of patients with COVID-19, as some severe types may lack these or any symptoms upon admission [5]. Without suitable medical intervention, these severe types may progress quickly to a critical condition, resulting in a high risk of mortality [6]. This uncertainty motivates a predictive method of patient types that is reliable and efficient, while also making use of alternative features. Early determination of patient types may enable healthcare professionals to improve their treatment plans and optimize facility resources.

The interest in integrating machine learning (ML) into general clinical practice

has grown rapidly in recent years [7]. Particularly, studies of implementing ML as a method for clinical decision support systems (CDSS) are ongoing. While these studies have shown great potential, their greatest limitation is a lack of interpretability [7]. Given the weight of their decisions, clinicians are hesitant to rely on "black-box" systems.

In response, the subfield of explainable artificial intelligence has emerged to provide clinicians with more transparent ML models [8]. This direction of work seeks to incorporate mechanisms within ML pipelines that output both reliable predictions and understandable decision processes. Early detection of COVID-19 patient severity using ML is often studied using a singular ML technique or data modality [9, 10]. Among the studies using multiple data types or ML techniques [11], many only utilized data from the early waves of infection. Some attempts to provide interpretable models were found [11], but many lacked this feature.

In this study, we investigated the performance and feature importance of various ML techniques for COVID-19 severity prediction, and then we evaluated feature modalities that provide the most predictive and reliable results. We trained ML models employing different techniques with patient-level biochemical and clinical feature modalities, both separately and together as a fusion modality. We applied logistic regression (LR), decision tree-based random forest (RF), k-nearest neighbors (kNN), and support vector machines (SVM), and we evaluated their abilities to predict severe COVID-19. We developed these ML models from data collected from original and Omicron COVID-19 patients to investigate model robustness across different variants.

2.2 Methods

Data collection

Our study uses two distinct datasets covering two time periods with distinct dominant viral variants. All patients were confirmed to be positive for COVID-19 by two independent quantitative reverse transcriptase-polymerase chain reaction (qRT-PCR) tests before inclusion in this study. The first set includes 362 patients infected with the original SARS-CoV-2 variant upon admission to Wuhan Union Hospital in China from January to March 2020. This dataset was previously described and analyzed by Chen et al. [5] and serves as a comparative baseline in this study.

Among these 362 patients, 148 were in severe condition according to guidelines established by the National Health Commission of China and the American Thoracic Society [12, 13], while the remaining 214 were designated non-severe types. Severe types were categorized by meeting at least one of the following criteria: (i) respiratory rate >30 breaths per minute; (ii) oxygen saturation $<93\%$ at rest; or (iii) PaO/fraction of inspired oxygen <300 mm Hg (40 kPa).

As this dataset contains data from patients infected by the original SARS-CoV-2 variant, this set is referred to as "original" in the rest of this study. The second dataset consists of 1000 patients admitted to Wuhan Union Hospital in China from December 2022 to January 2023, during which time patients were diagnosed with the SARS-CoV-2 Omicron variant. Using the same guidelines outlined earlier, 500 of these patients were severe COVID-19 type, while the other 500 were non-severe type.

All patients were comprehensively evaluated before being admitted to the hospi-

tals. Their fully de-identified, anonymous biomedical data were extracted from the electronic health record system. All participants were informed about the study, agreed to participate, and signed written consent. An institutional review board (IRB) application was submitted and approved by the Wuhan Union Hospital, Tongji College of Medicine, Huazhong University of Science and Technology (IRB approval #IEC-J-345), where the data were collected.

In our study, the input data were complete and without any missing feature data. The de-identified patient information comprised two main modalities of biomedical features. The first feature modality had 26 distinct laboratory testing features from blood tests, most of which were continuous values of the readings. The specifics of these tests are reported in detail in our prior study [5]. We refer to this feature modality as "biochemical" hereinafter. The second is a total of 26 features of one-hot encoded binary values indicating the presence of pre-existing conditions, comorbidities, symptoms, and other key risk factors such as demographic information. This modality was referred to as the "clinical" features. A complete description of these features across the two modalities is present in the supplementary materials of our prior study [5].

Together, features from both modalities were appended into a single corpus of de-identified patient data with 52 multi-modal features. This was referred to as the "fusion" set, as it fused across the continuous, real-valued biochemical feature modality and the binary clinical feature modality. We note that the specific features of respiratory rate, oxygen saturation, and fraction of inspired oxygen were excluded from our predictive feature list, as they were the original clinical standard

to determine COVID-19 severity.

ML pipeline development, validation and interpretation

We developed, evaluated, and compared the performance of several state-of-the-art ML classification techniques, including RF, kNN, and SVM. All ML techniques were implemented as supervised binary classification problems. To acknowledge LR's popularity in the field [14, 15, 16, 17, 18, 19, 20], we included it as a benchmark method. LR is generally sensitive to highly correlated variables, also known as multicollinearity, making predictions less precise [21]. This technique is also rather susceptible to overfitting, and thus it may be less capable of generalizing to unseen prediction sets. A powerful method of reducing overfitting is through regularization or penalization of regression. When LR is penalized (usually using l1 or l2 norms), the multicollinearity issue can be reduced [16].

To contrast with LR, we evaluate RF, kNN, and SVM. RF is an ensemble learning algorithm that adapts to nonlinearities within the data [22]. Within the "forest," each individual binary decision tree is built by recursively partitioning subsets of the training data according to the features that yield the most information gain. Like LR, RF is interpretable after training and is a frequent choice for CDSS tasks. The kNN classifier assigns samples to their predicted classes with which they share the most similarities, as determined by a chosen distance function [23]. This technique's performance depends on the choice of number of "neighbors" to the sample that the algorithm consults when predicting a class. We included kNN since it is a popular classifier due to its ease of use and effectiveness when dealing with larger datasets [24]. Lastly, SVMs are classifiers trained to create boundaries between classes in

the high dimensional feature space. SVM aims to maximize the distance between samples of different classes. This decision boundary is referred to as a hyperplane, and its geometry allows for application to both linear and non-linear problems. SVM is also a frequent choice of classification algorithm for CDSS tasks [25, 26].

Utilizing Python 3.10 and a variety of ML-related libraries from Sci-Kit Learn [27], we developed an end-to-end ML framework for each of the four classifiers to predict COVID-19 severity types from patients' clinical, biochemical, and fusion feature modalities. A full list of packages is included in the Appendix. We provide a graphical display of this pipeline in Figure 2.1.

The non-severe and severe COVID-19 types were labeled as 0 and 1, respectively. Each classifier was constructed to predict the outcome (0 or 1) based on the input features provided. For a given dataset, the corpus was first randomly partitioned into training and hold-out testing sets by an 80% to 20% split, respectively. Each feature type- biochemical, clinical and fusion were then preprocessed by a standard scaler separately prior to training to ensure consistency across different features.

Following the scaling step, a grid search method of hyperparameter tuning was utilized during training to maximize the ML model's performance. The resulting optimal hyperparameters for each classifier were detailed in the Appendix. Upon completion, the model with trained hyperparameters was then applied to the 20% hold-out data for testing. This process was repeated 50 times to generate different random training-testing splits and hyperparameter searches, each resulting in an independent model. To summarize, each model only witnessed 80% of a given dataset for training, and was evaluated for its performance on the remaining 20% of unseen

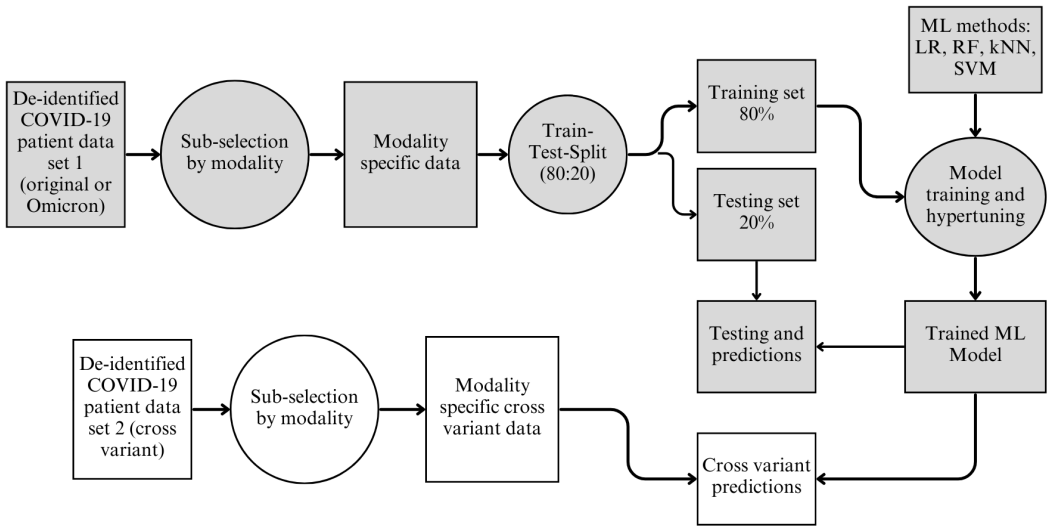

Figure 2.1: De-identified patient biomedical data were collected from Wuhan Union Hospital from January to March 2020 (original COVID-19 patients n=362) and December 2022 to January 2023 (Omicron variant patients n=1000). The data were grouped by their feature modality: biochemical, clinical or fusion. An 80:20 split was performed to generate training and testing sets, respectively. Four classifiers were developed and evaluated for their performance: LR=logistic regression, RF=random forest, kNN=k-nearest neighbors, and SVM=support vector machine. Upon hyperparameter tuning, each classifier was validated with the same-variant hold-out set, as well as the preprocessed cross-variant set.

data. This process was conducted to avoid overfitting [28] and establish an average performance for each setting.

ML classifiers' performances were evaluated by the receiver operating characteristic (ROC) curve and computing the area under the curve (AUC) of ROC. In this study, AUC was selected as the main performance metric (as opposed to F-measure or accuracy) because AUC has been shown to be more reliable than the other metrics [29]. For both the original and Omicron variant data sets, we evaluated ML classifier performance of training and testing on each modality separately and fused.

Each data set (binary clinical feature modality alone, continuous biochemical feature modality alone, and fusion modality that incorporates both feature modalities) underwent the pipeline defined earlier. To evaluate the robustness of the developed ML classifiers, we swapped and cross tested with testing data from the other variant. For example, models trained on the original SARS-CoV-2 variant data were also tested with Omicron variant data, and vice versa: this process was mirrored for classifiers trained on the Omicron variant data. For cross testing, the testing data was standardized according to the scaling scheme from the classifier's training data. During cross testing, the entire corpus was used as a hold-out testing set, since the classifiers were only trained with one variant and never trained with the other variant's data. The cross-set testing was evaluated for its performance exactly as the same-set testing.

One of the advantages of certain ML classification techniques is their interpretability in addition to performance. Of the classifiers that were developed in this study, we gathered insights from LR and RF classifiers. The resulting feature coefficient

vector obtained from training LR indicated what the classifier "learned" from the data. This appears as the largest absolute coefficient corresponding to the most influencing feature in predicting the severe COVID-19 type [30]. Pertaining to RF, feature importance was quantified and ranked by the feature's mean decrease in Gini impurity [31], which is commonly used in feature selection tasks [32]. During RF's training, these feature importances were computed using Sci-Kit Learn's feature_importances_ package [31]. By averaging the feature rankings (i.e., importance in predicting severe COVID-19 type) over 50 runs, we compared feature importance identified by LR versus RF, as well as different feature importances between the original and Omicron variant data. These comprehensive investigations enable us to validate findings from the ML classifiers by cross-checking results with other studies performing traditional statistical studies aimed at identifying predictive features of the severe COVID-19 type.

2.3 Results

ML classifier performance

Upon running each ML classifier pipeline for 50 independent repetitions, the average AUC value was calculated (Table 2.1). We validated that the computed average AUC was equal to the AUC of the composite ROC curve, hence there is no need for their distinction. These values were tabulated according to which SARS-CoV-2 variant dataset and modality were used for training each ML classifier. The standard deviations (SD) of the AUC did not exceed 0.06 across all testing scenarios. A summary of the standard deviations is found in the Appendix.

Our tuned ML classifiers demonstrated overall high AUCs in predicting severe COVID-19. The AUC scores can be found in Table 2.1. In general, ML models trained from Omicron variant data performed the best across all scenarios. The highest AUC among classifiers trained on Omicron data was 0.915, compared to the original's highest at 0.827. Performance of different ML classifiers trained on the same dataset showed minor differences. All ML classifiers developed from either original or Omicron data were almost indistinguishably accurate when tested on the Omicron data.

Table 2.1: Area under the Receiver Operating Characteristic Curve (AUC) scores

	ML method	LR		RF		kNN		SVM	
	Training	Original	Omicron	Original	Omicron	Original	Omicron	Original	Omicron
Modality	Testing								
Biochemical	Original	0.667	0.681	0.678	0.739	0.608	0.664	0.671	0.676
	Omicron	0.746	0.849	0.777	0.862	0.690	0.800	0.740	0.853
Clinical	Original	0.720	0.768	0.708	0.757	0.668	0.735	0.728	0.808
	Omicron	0.754	0.818	0.746	0.792	0.724	0.789	0.782	0.769
Fusion	Original	0.749	0.754	0.697	0.763	0.692	0.733	0.739	0.749
	Omicron	0.798	0.915	0.809	0.893	0.791	0.858	0.827	0.908

For each ML classifier, ROC plots were generated for the same-variant and cross-variant testings, yielding a total of 4 training-testing combinations (original-original, Omicron-Omicron, original-Omicron, and Omicron-original, where the latter two were cross-variant testings). Each ROC plot visualizes comparisons between the three feature modalities (clinical alone, biochemical alone, and fusion). All 16 plots are presented in the Appendix. For brevity, we discuss four graphs from RF in Figure 2.2. Each graph shows the composite ROC plot over 50 independent repetitions for each modality. The shaded regions denote one standard deviation from the mean of the true positive rate (TPR). The red, green, and blue lines represent the performance of models with biochemical, clinical, and fusion feature modalities, respectively.

Classifiers trained with the fusion feature modality generally demonstrated higher predictive power than either biochemical or clinical feature modality alone. The clinical feature modality performed slightly worse than the fusion modality, while biochemical feature modality alone had the least predictive power. Regardless of ML classification technique, models trained and tested with original variant data experienced the largest variation in their performance. We note that this combination also had the lowest performance of the four training-testing combinations. Classifiers trained with original data had improved performance when cross-testing with Omicron data.

Feature importance ranking

During each replication of the ML classifier, the feature coefficient vectors from the tuned LR classifier and the mean decrease of Gini impurity score vectors from the RF classifier were recorded and averaged after 50 replications. The averaged

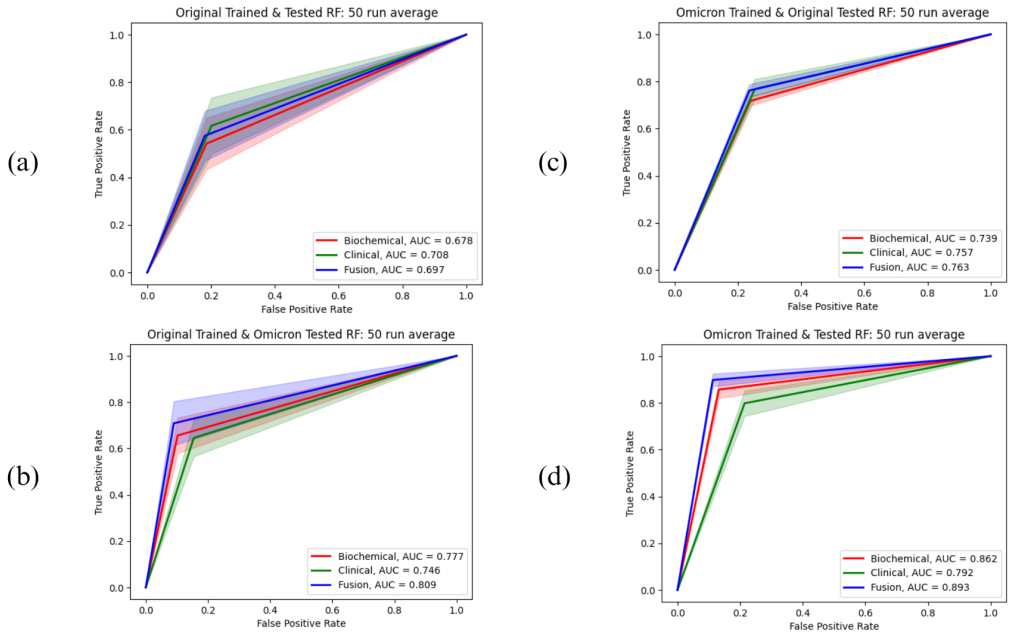

Figure 2.2: This figure shows the mean ROC curves plotted for each of the four training-testing combinations. Each figure contains the ROC curve of each data modality (biochemical-red, clinical-green, and fusion-blue.) Panel (a) plots the curve of classifiers trained and tested on original data only (i.e., original-original combination). Panel (b) plots the curve of classifiers of the original-Omicron combination. Panel (c) plots the curve of classifiers of the Omicron-original combination. Panel (d) plots the curve of classifiers of the Omicron-Omicron combination.

coefficients were then used to identify important features that differentiate severe from non-severe COVID-19 and to evaluate how different training data (original or Omicron variant) and feature modalities influence the feature rankings. Note that LR's associated coefficient vector has real values, while RF's reduction of Gini importance is interpreted as probabilities in [0, 1].

In this study, feature importances were determined by the fusion modality, as feature fusion demonstrated the highest predictive power for severe COVID-19. These feature importance rankings are presented in Figure 2.3. Regardless of ML classifiers

(LR and RF) or SARS-CoV-2 variant, features such as d-dimer (DD; biochemical modality), high sensitivity Troponin I (hsTNI; biochemical), and age $>$ 55 years (OLD; clinical modality) were consistently ranked in the top five most predictive features for COVID-19 severity. Features that often appeared among the top ten also include high sensitivity C-reactive protein (hsCRP; biochemical) and hypertension (HYP; clinical).

There were also some disagreements in feature rankings between the two techniques. For instance, LR suggested a history of chronic obstructive pulmonary disease (CPD; clinical) as an important predictive feature of the severe COVID-19 when trained on both the original and Omicron variant data, but CPD was not identified as a top feature in RF. Conversely, only RF identified elevated lymphocytes (LY; biochemical), ferritin (FERR; biochemical), and interleukin-6 (IL-6; biochemical) as important features regardless of SARS-CoV-2 variants.

By comparing the feature rankings across variants, LR trained on the original variant data identified low-, mid- and high-grade fever (LOD, MDF, HIF; clinical) all among the top 10 most predictive features, while its counterpart trained on Omicron variant data identified elevated procalcitonin (PCT; biochemical), percent of neutrophil (NE.1; biochemical) and white blood cell (WBC; biochemical) as the most predictive features. Such discrepancies in feature rankings were not observed in results from RF classifiers trained on different variants' datasets.

Lastly, there were some differences in the range of feature importance (quantified by coefficients in LR and Gini impurity scores in RF) across the two variants. LR's feature coefficients on average fell in the range [-0.95, 2.30] for the original variant,

whereas the range was [-0.70, 2.85] for the Omicron variant. Mean decreases of Gini impurity were in the range of [0, 0.10] and [0, 0.12] for the original and Omicron variants, respectively.

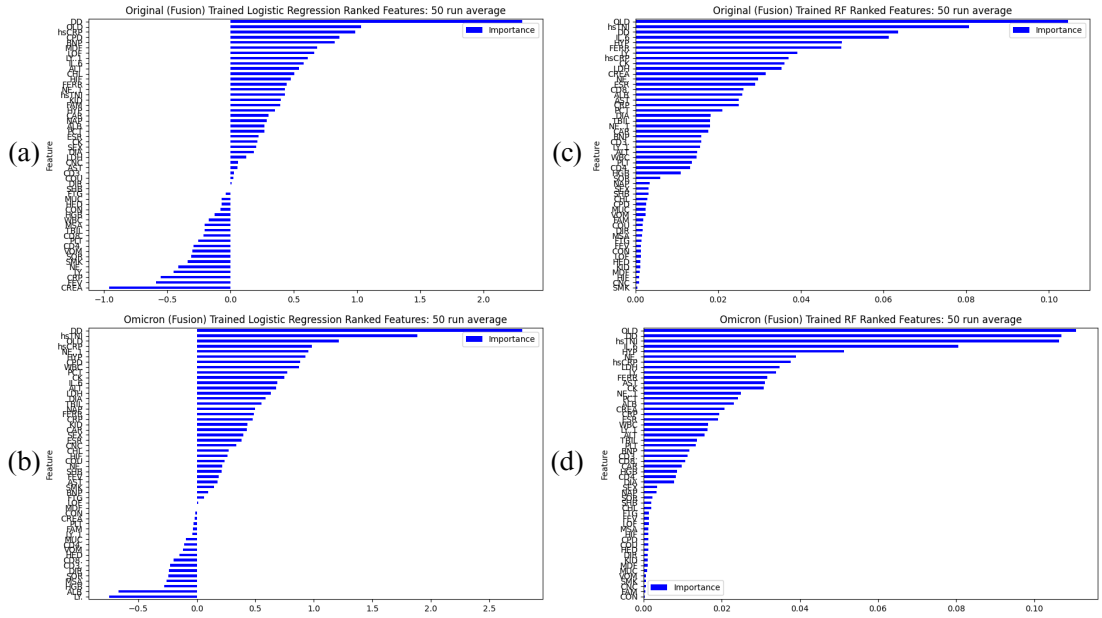

Figure 2.3: This panel presents the feature importance rankings of the logistic regression (LR) and random forest (RF) models. Panel (a) plots feature weigh coefficients for LR when trained on Original COVID-19 patient data. Panel (b) plots feature weigh coefficients for LR when trained on Omicron COVID-19 patient data. Panel (c) plots root features for RF when trained on Original COVID-19 patient data. Panel (d) plots root features for RF when trained on Omicron COVID-19 patient data.

2.4 Discussion

ML performance and interpretability

In this study, we evaluated the predictive power of multiple ML techniques when utilizing different feature modalities. We also found differences in model performance and interpretations across different SARS-CoV-2 variants. Overall, we found ML to be a useful tool for predicting COVID-19 severity based on comprehensive individual patient-level data. More importantly, we found that fusion of the biochemical and clinical modalities enhanced the predictive power of all types of ML models evaluated in this study. Models trained on multiple feature modalities have yielded the best performance in nearly every performance metric across all testing sets. These multimodal features are accessible by healthcare systems, especially with wide adoption of electronic health record systems. Results can be obtained efficiently from these systems, allowing the predictive ML model to be a fast and reliable CDSS tool to identify patients with high risk of severe COVID-19 [33].

The similarity of performance between the four ML techniques evaluated in this study suggests that the specific choice of modeling technique is not important for the task of predicting severe COVID-19 types. In general, LR, RF, and SVMs all showed relatively strong performance with their highest AUC scores being 0.915, 0.893 and 0.908, respectively. The kNN model exhibits the weakest performance of the methods considered here, with its highest AUC being 0.858. This may be due to the relatively small size of the datasets and requires further investigation with more samples [24].

Since model interpretability is important to the integrability of these ML models

into CDSS [7], LR and RF should be considered. The LR model offers the analyst information on which features are positively and negatively associated with the risk of severe COVID-19. However, LR is susceptible to multicollinearity between different features [21]. The RF model, on the other hand, is more resilient to the multicollinearity issue in the input data [22]. The RF model's reliability and robustness are consistent with the results produced by other similar studies [5, 17]. Upon validation, this study was among the first to use ML to identify critical biomedical features with the most predictive power of differentiating severe COVID-19 patients across different dominant variant phases.

The feature rankings provided by LR and RF are important for clinical decision making and provide clues to COVID-19 pathology. Our study indicates that elevated biomarkers such as D-dimer for coagulation, as well as hsTNI and hsCRP as indicators for heart damage, are strongly associated with severe COVID-19. This finding is compatible with other works which have shown that cardiovascular injury due to COVID-19 is highly associated with severe disease and adverse patient outcomes [34]. Other studies suggest that higher D-dimer is associated with higher risk of progressing to severe stage [35]. Our findings also suggest that patients' clinical information such as being 55 years or older, or having pre-existing conditions such as hypertension and COPD, could significantly increase the risk of progressing to severe COVID-19. Other studies also agree that age and hypertension are major risk factors for severe COVID-19 [36, 37, 38]. The identifications fit well within the work towards constructing explainable ML pipelines as they provide clinicians with the machine's decision-making process [7].

By identifying key risk factors associated with severe illness before its onset, ML can give clinicians augmented views of patient information and the possibility of personalized treatment [39]. Throughout the COVID-19 pandemic, COVID-19 patients in China were triaged based on their severity, in which severe types were treated at separate facilities from the non-severe types [40]. The designated hospitals for severe patients were part of a coordinated emergency response to the surge of infections. It was reported that these response measures and designated facilities resulted in improved recovery rates for severe patients [41]. These improved patient outcomes were made possible by accurate differentiation of patient types throughout the epidemic.

When comparing important features between patients infected by original and Omicron variants, we have identified an increase in the feature weight vector and Gini impurity values, which have not been reported before. This finding suggests that COVID-19 severity may have become more predictable in more recent variants. This might explain the higher variability and lower performance of models trained and tested on the original SARS-CoV-2 patient data. However, this claim requires significant further research with larger datasets over more detailed timelines. We also speculate that patient-level data may have higher quality in the Omicron wave than the original variant. We pose a potential explanation for this difference in data quality.

As previously mentioned, our data originates from Wuhan Union Hospital in China in January to March 2020 and December 2022 to January 2023. Throughout the epidemic, government responses to disease prevention and control, medical care protocols, and national guidelines changed regularly [40]. In particular, the pathology,

clinical manifestation, and diagnosis of SARS-CoV-2 evolved over this time frame [42]. According to one study detailing the timeline of such changes, clinical treatment protocols changed five times between January and March 2020 [42]. This same study reports only one change during the window of December 2022 to January 2023. The number of changes may be reflected in our data, as we found higher variability in the models trained on the data collected early during the epidemic.

Limitations and future work

This study contains a variety of limitations. One hindrance to the generalizability of this ML framework for clinical decision support is the lack of variation in the study data samples. Due to the emergency of COVID-19, all biomedical data in this study were taken at the individual's time of admission to one healthcare facility, and most patients were of the Han Chinese ethnicity group. The limited ethnicity coverage could result in potential sampling bias, and the conclusions from this study might only apply to the specific demographic groups.

While this study builds a ML framework and demonstrates its feasibility in COVID-19 CDSS, it is necessary to further validate the findings (e.g., key influential clinical and biomedical features) with larger scale multi-center studies across different regions and different phases of the pandemic to increase patients' representativeness [43]. We plan to identify additional studies and/or collect new data when possible to enhance data representation, especially across more demographic groups. Findings from these future studies could further evaluate the robustness of the ML workflow (e.g., whether there are further variations in influential feature sets), and lead to new clinical discoveries and insights on the pathological mechanisms of COVID-19

prognosis in different demographic groups.

Furthermore, the generalizability of this study is limited by the differences in datasets. As mentioned earlier, our results witnessed differences in performances between the two sets, resulting in speculation that patient-level data may have higher quality in the Omicron wave than the original variant. It is also possible that the difference in cohort sizes and prevalence of severe cases influenced the models' performance metrics. Nevertheless, the two cohorts were obtained from the same hospital, and we had carefully designed the inclusion criteria to make the two cohorts as comparable as possible, especially with regard to key risk factors such as gender and age. Future research could improve this limitation by performing analysis on datasets with fewer differences in size and prevalence.

Another limitation is that we were not able to evaluate ML model predictability for other major SARS-CoV-2 variants, such as Alpha and Delta. Retrospective studies are needed to comprehensively evaluate the robustness of the developed ML models across different phases of COVID-19 with different dominant variants and subvariants. These evaluations may be able to support or provide alternative hypotheses for the higher variation in predictability in the earlier data as it was previously discussed.

Considering the importance of interpretability in ML for clinical decision support, we acknowledge a limitation in our interpretation methods. For the RF models, we calculated feature importance using the mean decrease in Gini impurity. Gini impurity measures have been shown to be biased towards features with a high number of possible split points [44]. This can often result in continuous features to be favored over binary features when ranking their importance using an impurity measure. We

acknowledge that this may have introduced bias to the random forest feature rankings of the fusion models, as the biochemical features were continuously valued whereas the clinical features were binary. This potential favoritism of the biochemical features over the clinical features in our fusion models limits the clinical interpretation of our feature ranking results.

Future research could overcome this limitation by using another method of feature ranking. One potential method proposes an alternative way of constructing decision trees during the training phase [45]. In this study, statistical methods are used to pre-select the most informative and unbiased features for constructing the trees. As reported by the study, this results in more accurate decision trees and a reduction in the dimensionality of training data. Exploring these and other methods for debiasing feature rankings from random forest classifiers would provide an opportunity to improve on this limitation.

There are a variety of future research directions suggested by this study. Notwithstanding improvements made on the limitations previously discussed, we acknowledge the existence of other emerging ML techniques worthy of investigation. Considering the desire for explainable ML models, techniques with easily understandable decision processes would be of the highest interest. Additionally, given the tentative promise of LR as a predictive tool for COVID-19 CDSS, other regression techniques such as the Lasso or Ridge regressions may be useful. Since the Lasso and Ridge regressions are also interpretable, an analysis of the performance and feature importances of these techniques may provide even more insight.

While we found the individual ML techniques to be predictive of severe COVID-19,

investigating the power of advanced ensemble techniques may provide new analyses. Explorations of ensemble techniques involving LR, RF, kNN, and SVM are showing signs of predictive power for cardiovascular diseases using similar underlying datasets to this study [46, 47]. These studies using ensemble techniques suggest that predictive power is enhanced even when individual classifiers are not as robust [46, 47]. Our study may find improved predictive power of severe COVID-19 by utilizing more advanced ensemble methods beyond RF.

Further studies beyond investigating COVID-19 may shed light on the predictive power of ML techniques from individual-level data collected from patients with other respiratory illnesses. This is especially useful to healthcare systems inundated with patients infected with various influenza strains or respiratory syncytial virus (RSV). Using similar ML techniques and leveraging the power of transfer learning, our developed ML pipeline can be further applied to other diseases with similar underlying datasets (e.g., clinical and biochemical). Connecting back with the goal of aiding healthcare resource optimization, a potential application of this work is to simulate burdens on the health system of an unexpected inflow of patients, some of whom are severe patients and therefore need intensive care.

Another potential direction is the incorporation of more data modalities, such as patient-level medical imaging (including X Ray and CT scans) and multi-omics data. Due to the higher dimensionality of imaging modality in relation to the biochemical and clinical modalities, more advanced ML techniques such as a deep convolutional neural network (CNN) would need to be applied to handle such modalities.

2.5 Conclusion

The ability to accurately differentiate patient types is crucial to clinicians treating patients with a novel pathogen like COVID-19. During such emergencies, CDSS tools can provide clinicians supplemental data and insights alongside their workflow. We found that ML may be a useful component of the CDSS. Our ML pipeline showed potential in its ability to accurately predict COVID-19 severity types even across different variants of the illness. This opens the opportunity to study the applicability and generalizability of using ML in CDSS for other known or emergent severe illnesses.

CHAPTER 3: MODELING EMERGENCY PREPAREDNESS IN THE HEALTHCARE SYSTEM

3.1 Introduction

Public health emergency preparedness is a vital aspect of the healthcare system. Emergencies like the COVID-19 pandemic motivate efforts to optimize resource allocation and prioritize healthcare personnel safety [48]. Critical care resources become especially valuable as the influx of critically ill patients overwhelms intensive care units (ICUs) during health emergencies [2]. The high demand for ICU resources necessitates hospitals and healthcare systems to address the challenges with capacity and staffing levels [49].

Mathematical modelers have met this challenge with a variety of tools, one of which being the Markovian queueing theoretic model [50, 49, 51, 52, 53, 54, 55]. Queueing theory (QT) is the mathematical study of lines or "queues" in a service system, originating from early telephone communication research done by A. K. Erlang [56]. Since 2004 [50], modelers have turned to the queueing system as a means to estimate the necessary number of ICU beds to meet community needs. These models have shown promise in accurately projecting the number of beds needed in the ICU to meet community demands [50]. After finding QT to be a useful tool in projecting ICU resources, researchers began studying models of various ICU structures using QT. One study in 2015 compares the allocation of different beds for both high level

critical patients and intermediate level critical patients to improve the distribution of resources [49].

Soon after the emergence of SARS-CoV-2 causing the COVID-19 pandemic [1], modelers turned to the queueing theoretic model to assess the preparedness of the healthcare system [52, 53, 55]. Some used this framework to assess the number of beds needed to care for critically ill patients infected with COVID-19 [53, 55]. Others estimated the number of ventilators needed by incorporating disease transmission models of COVID-19 to project the additional arrivals [52].

One aspect of the ICU that has not been modeled with QT is the burden placed on healthcare providers. Clinicians, nursing staff, technicians and environmental service workers face significant additional workplace risks during emergencies like the COVID-19 pandemic [57]. During times of community disease spread, these healthcare providers face exposure to pathogen and are at risk of contracting illness [57]. The combination of high ICU utilization and decreased healthcare worker staffing levels during an emergency can result in increased adverse patient outcomes [58] and increased risk to the remaining staff [59].

We propose a novel queueing theoretic framework for incorporating healthcare personnel into the simulated ICU. In this modeling framework, we aim to capture the trade-off between ICU capacity and healthcare provider availability. Using this model, we seek to aid in providing recommendations for resource allocation when constrained by an operational budget. In this work, we demonstrate this model's ability to simulate various systemic trends and a diverse set of healthcare facility types under multiple distinct scenarios.

3.2 Methods

Model foundation

QT is the mathematical study of lines or "queues" in a service system. The core components of a queueing system are it's arrival process, number of servers, departure process, and queue discipline [60]. The arrival process describes the process by which individuals arrive to the system. Arrivals may be scheduled or appear at random. The number of servers describes the resources each individual uses before departing. Some specific examples include the number of cashier lanes at a grocery store or number of beds in a hospital ward. The departure process describes the process by which individuals leave the system. This is often based on the average length of the service time. Lastly, the queue discipline describes the rules by which individuals wait in line or queue when there is not an available server. These rules determine the size of the queue, which can be finite or infinite. They also determine how individuals of the queue progress through the system as servers become available, with many using the "first in, first out" rule. When applicable, queue disciplines may allow for "reneging" or individuals abandoning the queue without service [61].

We describe a standard model of QT to provide the foundation of our model. These standards are used in related works using QT for modeling the ICU [50, 49, 51, 52, 53, 54, 55]. Similar to these works, we made the following assumptions. We use a continuous-time Markovian birth-death process to describe our arrival and departure processes, respectively. The "births" or arrivals, are patients who are recommended for admission to the ICU. This random process is a Poisson process with parameter λ set equal to average number of patients admitted in a 24 hour period. To find

the times between arrivals, a random variable can be drawn from an exponential distribution with the same parameter λ as described in the equation below [61]:

$$p(t) = \lambda e^{-\lambda t}. \tag{3.1}$$

The number of servers n corresponds to the number of ICU beds. When patients are recommended for admission, they are treated in an ICU bed if one is available. If the unit is at capacity, patients wait in the queue for an available bed. Following the "first in, first out" rule, the patient who is at the front of the line is the first to be transferred to the ICU once a bed is available. We assume an infinite queue size.

The "deaths" or departures from the ICU follow a Markov process similar to that of the arrivals. Times between departures are drawn from an exponential distribution like Equation (3.1) with parameter μ_i described by

$$\mu_i = \begin{cases} i\mu; & i \in [0, n) \\ n\mu; & i \in [n, \infty). \end{cases} \tag{3.2}$$

The individual departure rate μ is the inverse of the average service time or patient length of stay (LOS). We assume this value is the mean period between the patient's admission and discharge, either through recovery or mortality.

For a standard queueing system with n servers, the expected server utilization or "traffic density" is calculated by the ratio

$$\text{density} = \frac{\lambda}{n\mu}. \tag{3.3}$$

For this queueing system, there exists a unique stationary probability distribution for patients waiting in the queue if and only if $\frac{\lambda}{n\mu} < 1$ [60]. Otherwise, when this ratio is greater than or equal to 1, this indicates the system is overburdened and we expect the queue length to grow.

To acknowledge the importance of treating critically ill patients in a timely fashion, our queue discipline allows for reneging. Patients may abandon the queue for a number of reasons such as recovering in a general ward within the hospital, transferring to another facility, or suffering mortality [62]. We assume times between reneging follow an exponential distribution according to parameter ρ_k calculated by

$$\rho_k = k\rho; \quad k \in [0, \infty), \tag{3.4}$$

where k is the number of patients waiting in the queue, and ρ is the probability an individual will abandon the queue. Critically ill patients reneging for any of the reasons mentioned earlier is the result of an overburdened system. Later in this study, we analyze systems largely through their resulting accumulated abandonment estimates.

While this system provides a strong modeling foundation, we extend this model to account for the change in arrivals occurring during a community epidemic. These additional features include time dependent arrival rates, patient type dependent departure rates, and patient-to-healthcare provider transmission.

Epidemic dependent arrival rates

The static arrival rates previously described cannot capture the sudden influx of critically ill patients suffering from an infectious disease. To incorporate these arrivals, we introduce an additive time dependent component to the arrival rate λ. Towards this aim, we simulate the spread of a novel pathogen through the community. We aim to simulate a pathogen with characteristics similar to infectious respiratory illnesses like respiratory syncytial virus (RSV), influenza, and ancestral SARS-CoV-2 (COVID-19). We simulate the spread using a deterministic compartmental model of ordinary differential equations. We provide an overview of this model here; however those interested in or unfamiliar with this class of disease modeling may explore these additional articles [63, 64, 65]. Consistent with other compartmental models of SARS-CoV-2 [66, 67] and influenza [68], the compartments of the simulated population are Susceptible, Exposed, Infected, and Recovered/Removed.

We used the following system of equations to simulate the transmission of the

Table 3.1: SEIR compartment transition rate parameters and descriptions

Parameter	Description (units)
β	transmission coefficient (dimensionless)
α	rate of transitioning from exposed to infected (inverse of latent period in days)
γ	rate of transitioning from infected to recovered (inverse of recovery period in days)
N	community population size (number of individuals)

novel pathogen through the community

$$\begin{aligned}
\frac{dS}{dt} &= -\frac{\beta SI}{N} \\
\frac{dE}{dt} &= \frac{\beta SI}{N} - \alpha E \\
\frac{dI}{dt} &= \alpha E - \gamma I \\
\frac{dR}{dt} &= \gamma I
\end{aligned}$$

where $S + E + I + R = 1$. The rate parameters directing the flow between each compartment are described by Table 3.1.

The SEIR model is normalized by N for the community population size. We parametrize this model with values considering parameter ranges found in published literature. The basic reproduction numbers for these pathogens can vary spatially and temporally. We chose a basic reproduction number of 2.73 to represent our hypothetical pathogen, falling in the range for the three pathogens mentioned previously [69, 70, 71, 72, 73]. We chose a mean incubation period (α) of 3 days [68, 74, 75, 76], and a mean recovery period (γ) of 5 days [68, 77, 78]. We assume that individuals in the Exposed class are not infectious. For this deterministic SEIR model, the basic reproduction number R_0 is given by [79]

$$R_0 = \frac{\beta}{\alpha}. \tag{3.5}$$

Using this formulation for R_0, we solve for the transmission coefficient β. The initial conditions of the system are set as $S(0) = 0.99, E(0) = 0.01, I(0) = 0$, and

$R(0) = 0$, meaning 1% of the population has been exposed to the pathogen and the remaining 99% are susceptible.

Once the infection wave is simulated, at any given time during the Markov chain we may access the estimated number of infected individuals. Assuming that a proportion of these infections are critical (κ) yields the expected number of epidemic-related patients that will arrive to the ICU. This dynamic arrival rate is referred to as $\lambda_{\text{epidemic}}(t)$. In order to integrate the normalized SEIR model into the queueing framework, we calculate $\lambda_{\text{epidemic}}(t)$ by the following:

$$\lambda_{\text{epidemic}}(t) = \kappa N \cdot I(t)$$

where N represents the size of the community population.

The baseline stream of non-epidemic related patients are still assumed to arrive to the ICU at their static rate, $\lambda_{\text{baseline}}$. To model both of these patient types, our arrival rate $\lambda(t)$ is described by

$$\lambda(t) = \lambda_{\text{baseline}} + \lambda_{\text{epidemic}}(t). \tag{3.6}$$

Thus in our adapted model inter-arrival times are drawn according to an exponential distribution with parameter $\lambda(t)$ defined by

$$p(t) = \lambda(t)e^{-\lambda(t)t}.$$

When an arrival occurs, we distinguish the patient type according to the following

probability

$$\text{Arriving patient type} = \begin{cases} 0 & \text{w.p. } \frac{\lambda_{\text{baseline}}}{\lambda(t)} \\ 1 & \text{w.p. } \frac{\lambda_{\text{epidemic}}(t)}{\lambda(t)}, \end{cases} \tag{3.7}$$

where 0 indicates the patient is non-infectious and 1 indicates the patient is infectious. This arrival is placed in a bed if one is available, otherwise they wait in the queue.

As previously mentioned, our SEIR model is normalized by N for the community population size. This approach allows us to hold the disease dynamics constant for different scales of population. This choice is consistent with frequency-dependent (FD) transmission modeling as opposed to density-dependent transmission modeling [80]. By choosing the FD transmission approach, we avoid confounding effects of the rate $\lambda_{\text{epidemic}}(t)$ on our model outcomes.

By determining simulated patient types, we may see clearly the impact of the surging infectious patients on the ICU system. We will be able to discuss the patient composition of the queue, the ICU, and those who abandon the queue. This classification will also be necessary in our methods of implementing dependent departure rates and patient to healthcare provider transmission.

Patient type dependent departure rates

Distinguishing patient types allows us to model different recovery times for each type. In general, literature reviewing patient data found COVID-19 [77] and influenza [81] patients had a longer average LOS in the ICU when compared to the overall average LOS [82]. This poses the necessity to differentiate these types for the most accurate parametrization of the departure rates.

When a patient is placed in a bed, their patient type (0 or 1) is recorded. This patient is either an arrival to the system or is next in line from the queue. Recall that the patient type 0 is assumed non-infectious while the patient type 1 is assumed infectious. Each type is assumed to have an average length of stay of either μ_1 for non-infectious or μ_2 for infectious. Let i denote the number of non-infectious patients and let j denote the number of infectious patients. Then the overall departure rate $\mu_{i,j}$ is

$$\mu_{i,j} = \begin{cases} i\mu_1 + j\mu_2; & i + j \in [0, n) \\ (n-j)\mu_1 + j\mu_2; & i + j \in [n, \infty), \end{cases} \tag{3.8}$$

where i and j are non-negative integers. Thus in our adapted model inter-departure times are drawn according to an exponential distribution with parameter $\mu_{i,j}$. When a departure occurs, the patient type is determined by the following probability

$$\text{Departing patient type} = \begin{cases} 0 & \text{w.p. } \frac{i}{i+j} \\ 1 & \text{w.p. } \frac{j}{i+j}. \end{cases} \tag{3.9}$$

After this event occurs, the result is used to update i and j for the next time step.

Patient-to-healthcare provider transmission

Specialized healthcare providers, namely nurses and clinicians, are essential to the function of the ICU. High utilization of the ICU during an epidemic poses risks to healthcare providers. Particularly, providers are at risk of contracting a disease when there is high prevalence of infectious patients. When staffing levels are reduced

due to illness, the functionality of the ICU decreases. To model this, we introduced healthcare providers as a secondary and susceptible resource to our queueing theoretic system.

We implemented a minimum safe staffing ratio between patients and providers. For simplicity one provider type is selected at a time, either nurses or clinicians, since guidelines for these minimums is significantly different for these roles [83]. In our model, we simulate transmission between patients and clinicians, but this framework is adaptable to other provider roles. The maximum number of patients in the ICU is given by n, the bed capacity of the unit. Each staff member is assigned up to r patients per day. We assume a provider will work three 12 hour shifts in a week and an absolute minimum of 5 providers is necessary to avoid overworking the staff while maintaining constant 24/7 coverage. We arrive at the following calculation for the total number of clinicians on staff H based on ICU capacity and the minimum safety coverage:

$$H = \max\left(5, \left\lceil \frac{14}{3} \left\lceil \frac{n}{r} \right\rceil \right\rceil\right). \tag{3.10}$$

Conversely, this calculation can be rewritten to determine the maximum number of patients that can be safely treated dependent on the number of available clinicians. This effective ICU capacity n^* is expressed by

$$n^* = \min\left(n, \left\lfloor \frac{3}{14} \lfloor rH \rfloor \right\rfloor\right). \tag{3.11}$$

If a clinician is unable to work their expected schedule with no substitute, then the

number of available clinicians is reduced by one. At each time step, Equation (3.11) is computed using the number of available clinicians. If the minimum is not met, then the effective capacity of the ICU n^* is decremented to meet the requirement.

Enforcing the minimum required providers for the ICU is necessary as we introduce their susceptibility to an infectious disease. Staying within the continuous-time Markov chain (CTMC) framework, we introduce the rate at which providers become infected η_j calculated by

$$\eta_j = j\eta; \quad j \in [0, \infty), \tag{3.12}$$

where j is the number of infectious, epidemic-related patients in the ICU and η is the probability that an individual patient infects a provider. Times between patient-to-provider transmission events are assumed to be exponentially distributed with parameter η_j. The overall provider recovery rate ν_g is calculated by

$$\nu_g = g\nu; \quad g \in [0, H], \tag{3.13}$$

where g is the number of inactive providers and ν is equal to the inverse of the recovery period. Thus inter-recovery times for infected healthcare providers is drawn from an exponential distribution with parameter ν_g.

Complete adapted model and scenario design

Altogether the arrivals, departures, abandonments, provider infections and provider recoveries form the state space for our model. A visualization of the system architecture is displayed in Figure 3.1. To run a complete simulation of the CTMC, at a

given time t all rates $(\lambda(t), \mu_{i,j}, \rho_k, \eta_j, \nu_g)$ are updated. Then a time step δ is drawn from $\text{Exp}(\sigma)$ where

$$\sigma = \lambda(t) + \mu_{i,j} + \rho_k + \eta_j + \nu_g.$$

Once the current time is updated to $t + \delta$, the event that occurs is chosen from the following probability

$$\text{Event at time } t+\delta = \begin{cases} \text{Arrival} & \text{w.p.} \frac{\lambda(t)}{\sigma} \\ \text{Departure} & \text{w.p.} \frac{\mu_{i,j}}{\sigma} \\ \text{Abandonment} & \text{w.p.} \frac{\rho_k}{\sigma} \\ \text{HCP infection} & \text{w.p.} \frac{\eta_j}{\sigma} \\ \text{HCP recovery} & \text{w.p.} \frac{\nu_g}{\sigma}. \end{cases}$$

The event chosen occurs and all statuses are updated accordingly. Each run of the full simulation models 365 days. 200 sample paths are completed before finding the mean trajectory of the system. This model represents a single facility in which we track the following trends over time: the length of the queue, ICU utilization and capacity, queue abandonments, and available healthcare providers.

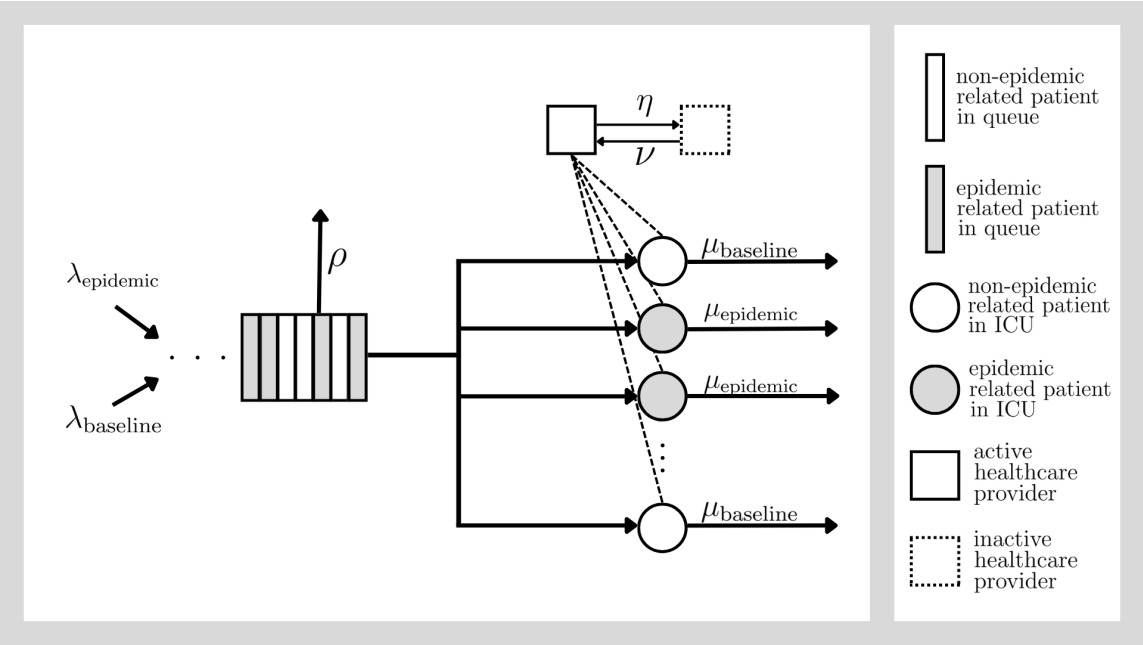

Figure 3.1: This visualization of the queueing system displays the flow of patients through our simulated ICU. Critically ill patients are requested for admission either due to an epidemic related illness ($\lambda_{\text{epidemic}}$) or non-epidemic related illness ($\lambda_{\text{baseline}}$), such as surgery or trauma. Patients are placed in an ICU bed whenever possible. If no beds are available, patients wait in an unlimited length queue. Patients may leave the queue without ICU treatment (ρ) due to recovery in a general ward, transferring to another facility, or suffering mortality. Patients in the ICU depart at a rate relative to their patient type (μ_{epidemic} or μ_{baseline}). Epidemic-related patients are considered infectious and may infect healthcare providers (η). If enough healthcare providers are ill and unable to work, the capacity of the ICU is reduced until they return (ν).

Optimal resource allocation

The complete adapted model described above can be parametrized to represent a wide range of single facility types. Through analysis of the queue abandonments, different facility structures can be compared to determine their emergency preparedness. To generate different facilities with similar amounts of resources, we propose the following constraints and meta-parameters to our model.

For a given healthcare facility, suppose it has an operational budget U for the ICU. All monetary values are assumed to be in USD. Let's determine the optimal allocation of healthcare providers and ICU beds during an emergency. Let a represent the annual salary of a healthcare provider and b represent the annual cost of one ICU bed. We also assume a percentage c of the operational budget is required for constant costs. Recall earlier we noted that only one provider type, clinicians or nurses, is selected to be susceptible to patient transmission due to the difference in interactions with patients. We also note that annual salary between clinicians and nurses is markedly different. For these reasons, we will select only one provider type to optimize over at a time.

In our study, we chose to optimize over the number of clinicians and will include nurse staff as another constant percentage cost d within the operational budget. Thus the budget amount with which we will optimize over $\overline{U}$ is given by

$$\overline{U} = (1 - c - d)U.$$

We set the constraint to be

$$\overline{U} = aH + bn,$$

where H is the number of providers and n is the number of ICU beds. Thus we define the complete resource optimization problem as

$$\begin{aligned} \arg\min_{H,n} \quad & \text{expected queue abandonments}(H, n) \\ \text{s.t.} \quad & \overline{U} = aH + bn. \end{aligned} \tag{3.14}$$

We perform a grid search over H and n along with stochastic simulations to estimate solutions to this optimization problem. Each single-scenario is simulated over 200 sample paths. This "multi-facility" simulation allows us to capture degrees of burden between these systems to find local and global minima. In addition, by extending the limit of our grid to include combinations outside of the budget constraint, we quantify the benefit of additional funds in terms of relative queue abandonment reduction. This gradient array allows us to identify recommendations for a given scenario, thus improving a facility's emergency preparedness.

To demonstrate the capabilities of this model, we designed the following case scenarios. For each one, systems are assessed by the number of queue abandonments occurring during the simulated 365 days. We aim to compare the abandonment rates and resource recommendations between a simulated high resource setting and a low resource setting. Each facility will be simulated under three distinct scenarios. A graphical representation of these scenarios is presented in Figure 3.3. The first scenario, Scenario A, will represent a pre-epidemic baseline. During this phase, there

is no wave of infectious patients. This will allow us to capture the typical behavior of each system. The second scenario, Scenario B, will represent an outbreak of a hypothetical novel pathogen as described earlier. Using the SEIR model calibrated to this hypothetical pathogen from published literature, we will simulate the wave of critically ill infectious patients. The third scenario, Scenario C, will represent an outbreak of a hypothetical novel pathogen with a 20% increase in transmissibility from the pathogen in Scenario B. The characteristics of the pathogen in Scenario C will be the same as Scenario B, with the exception of a 20% increase to the basic reproduction number, thus a 20% increase to the transmission coefficient β. Figure 3.2 displays the infection curves for each scenario, plotting the proportion of the population that is infected over time.

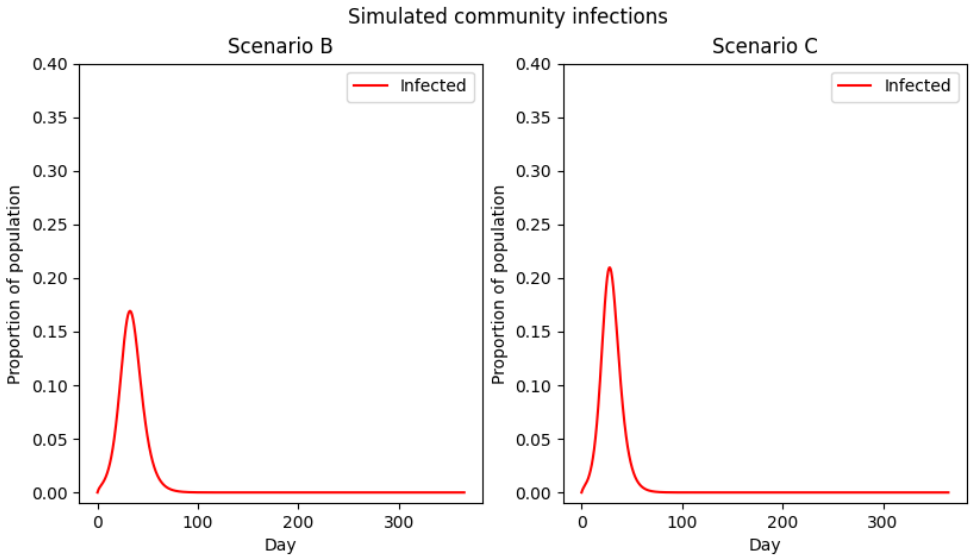

Figure 3.2: This figure shows the simulated proportion of the community population infected with a hypothetical novel pathogen over time. Scenario B represents a pathogen with a basic reproduction number $R_0 = 2.73$ (left). Scenario C represents a pathogen with a basic reproduction number $R_0 = 3.28$ (right).

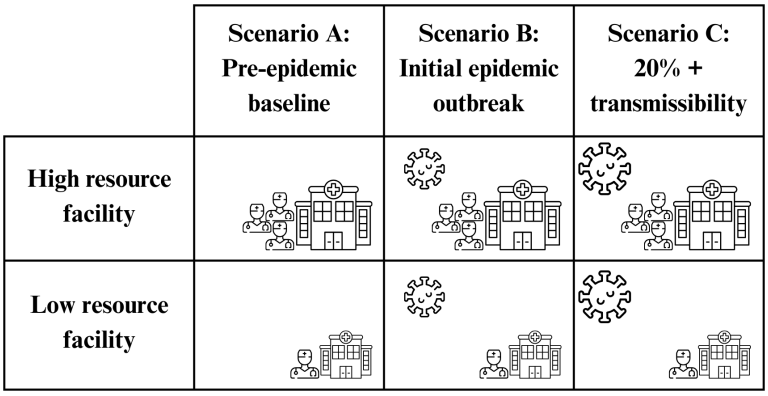

Figure 3.3: Scenario A represents a pre-epidemic baseline where there is no wave of infectious patients. Scenario B represents an outbreak of a hypothetical novel pathogen, simulating a wave of critically ill infectious patients. Scenario C represents an outbreak of a hypothetical novel pathogen with increased transmissibility. For A, B, and C, these scenarios are applied to two simulated facility settings, one with a high amount of resources and one with a low amount of resources.

Model parametrization and sensitivity analysis

This model depends on many values for parameterizing the queueing system, SEIR compartmental model, and budget constraints. We conducted a search of the literature for parameter ranges using PubMed, Google Scholar, and Elicit. Ranges and means were collected to design reasonable parameter settings for each of the case scenarios described above.

Before parameter selection, we conducted analysis of the model's sensitivity to perturbations in the baseline arrival rate $\lambda_{\text{baseline}}$, baseline departure rate μ_1, epidemic-related departure rate μ_2, queue abandonment rate ρ, patient-to-clinician transmission probability η, clinician recovery rate ν, and minimum staffing ratio $1/r$ on the

resulting accumulated queue abandonments. Sensitivity testing was conducted separately for both the high resource and low resource assumption, as the ranges for the baseline arrival rate varies significantly between the two resource levels. The high resource setting consisted of 15 clinicians with 35 ICU beds, while the low resource setting consisted of 5 clinicians and 10 ICU beds. We conducted testing assuming the initial outbreak scenario in order to analyze sensitivity on the epidemic-related parameters.

A baseline of queue abandonments was established with models using the mean values of the parameter ranges found from literature over 25 sample paths. The mean and standard deviation of the accumulated abandonment counts served as the baseline. To compare, one parameter at a time was selected to be altered, while all other parameters were held at their baseline values. At the start, we used a coarse method of testing the minimum and maximum of each parameter range. If the resulting mean queue abandonments was within one standard deviation established by the baseline model, we did not perform a fine analysis. For those parameters which did not fit this description, we ran a finer method of sensitivity testing by iterating over the entire parameter range.

We plotted the resulting queue abandonments for each parameter range to identify trends. This analysis informed our selection of parameters for the case scenarios discussed earlier. Upon completing sensitivity analysis, we chose parameters for modeling the six scenarios according to Table 3.2.

Table 3.2: Table of model parameters

Facility Parameters							
Parameter	Symbol	High resource	[Range]	Low resource	[Range]	Units	Source
ICU Budget	U	$2.6 \cdot 10^7$	$[1.7 \cdot 10^7, 3.5 \cdot 10^7]$	$5.0 \cdot 10^6$	$[3.0 \cdot 10^6, 5.0 \cdot 10^6]$	USD	[84]
HCP annual salary	a	$4.0 \cdot 10^5$	$[2.0 \cdot 10^5, 4.0 \cdot 10^5]$	$2.0 \cdot 10^5$	$[2.0 \cdot 10^5, 4.0 \cdot 10^5]$	USD	[85]
ICU bed annual cost	b	$1.5 \cdot 10^5$	$[1.0 \cdot 10^5, 1.5 \cdot 10^5]$	$1.0 \cdot 10^5$	$[1.0 \cdot 10^5, 1.5 \cdot 10^5]$	USD	[86]
Constant costs	c	32%	[30%, 35%]	32%	[30%, 35%]	% of budget	[84]
Other staff (not being modeled) costs	d	28%	[26%, 31%]	28%	[26%, 31%]	% of budget	[85]
Scenario A: Baseline/Pre-epidemic							
Parameter	Symbol	High resource	[Range]	Low resource	[Range]	Units	Source
Arrival rate	$\lambda_{\text{baseline}}$	7.2	$[4, 15]$	2	$[0.5, 7]$	patients per day	[87, 88, 89]
Departure rate	μ_1	1/3.4	$[0.059, 0.29]$	1/3.4	$[0.059, 0.29]$	reciprocal of mean length of stay in days	[82]
Queue abandonment	ρ	0.12	$[0.01, 0.17]$	0.12	$[0.01, 0.17]$	% transferred to another facility	[62]
Coverage ratio (HCP:patient)	$1/r$	1/9.3	$[1/11, 1/8]$	1/9.3	$[1/11, 1/8]$	1 clinician to 9.3 patients	[83]
Scenario B: Hypothetical epidemic outbreak							
Parameter	Symbol	High resource	[Range]	Low resource	[Range]	Units	Source
Arrival rate	$\lambda_{\text{baseline}}$	7.2	$[4, 15]$	2	$[0.5, 7]$	patients per day	[87, 88, 89]

Departure rate	μ_1	1/3.4	$[0.059, 0.29]$	1/3.4	$[0.059, 0.29]$	reciprocal of mean length of stay in days	[82]
Infectious departure rate	μ_2	1/8	$[1/15, 1/4]$	1/8	$[1/15, 1/4]$	reciprocal of mean length of stay in days	[77]
Queue abandonment	ρ	0.12	$[0.01, 0.17]$	0.12	$[0.01, 0.17]$	% transferred to another facility	[62]
Patient-HCP pathogen transmission	η	0.008	$[0.005, 0.011]$	0.008	$[0.005, 0.011]$	dimensionless	[57]
HCP recovery	ν	1/5	$[1/8, 1/4.1]$	1/5	$[1/8, 1/4.1]$	reciprocal of mean recovery period in days	[77, 90]
Coverage ratio (HCP:patient)	$1/r$	1/9.3	$[1/11, 1/8]$	1/9.3	$[1/11, 1/8]$	1 clinician to 9.3 patients	[83]
Basic reproduction number	R_0	2.73	$[1.5, 6.49]$	2.73	$[1.5, 6.49]$	average secondary infections per infective	[69, 70, 71, 73]
Community transmission rate	β	0.5467	$[0.300, 1.298]$	0.5467	$[0.300, 1.298]$	dimensionless	[69, 70, 71, 73]
Pathogen incubation rate	α	1/3	$[1/4, 1/1.2]$	1/3	$[1/4, 1/1.2]$	reciprocal of mean incubation period in days	[76, 90]
Pathogen recovery rate	γ	1/5	$[1/8, 1/4.1]$	1/5	$[1/8, 1/4.1]$	reciprocal of mean recovery period in days	[77, 90]
Community population	N	91,176	N/A	28,818	N/A	number of people in community	[91]
Proportion of critically ill infected patients	κ	0.11%	[0.013%, 3.2%]	0.11%	[0.013%, 3.2%]	% of infected individuals requiring critical care	[92, 93]

Scenario C: Hypothetical epidemic outbreak with 20% increased transmissibility							
Parameter	Symbol	High resource	[Range]	Low resource	[Range]	Units	Source
Arrival rate	$\lambda_{\text{baseline}}$	7.2	$[4, 15]$	2	$[0.5, 7]$	patients per day	[87, 88, 89]
Departure rate	μ_1	1/3.4	$[0.059, 0.29]$	1/3.4	$[0.059, 0.29]$	reciprocal of mean length of stay in days	[82]
Infectious departure rate	μ_2	1/8	$[1/15, 1/4]$	1/8	$[1/15, 1/4]$	reciprocal of mean length of stay in days	[77]
Queue abandonment	ρ	0.12	$[0.01, 0.17]$	0.12	$[0.01, 0.17]$	% transferred to another facility	[62]
Patient-HCP pathogen transmission	η	0.0096	$[0.005, 0.011]$	0.0096	$[0.005, 0.011]$	dimensionless	[57] + 20%
HCP recovery	ν	1/5	$[1/8, 1/4.1]$	1/5	$[1/8, 1/4.1]$	reciprocal of mean recovery period in days	[77, 90]
Coverage ratio (HCP:patient)	$1/r$	1/9.3	$[1/11, 1/8]$	1/9.3	$[1/11, 1/8]$	1 clinician for 9.3 patients	[83]
Basic reproduction number	R_0	3.28	$[1.5, 6.49]$	3.28	$[1.5, 6.49]$	average secondary infections per infective	[69, 70, 71, 73]
Community transmission rate	β	0.6560	$[0.300, 1.298]$	0.6560	$[0.300, 1.298]$	dimensionless	[73] + 20%
Pathogen incubation rate	α	1/3	$[1/4, 1/1.2]$	1/3	$[1/4, 1/1.2]$	reciprocal of mean incubation period in days	[76, 90]
Pathogen recovery rate	γ	1/5	$[1/8, 1/4.1]$	1/5	$[1/8, 1/4.1]$	reciprocal of mean recovery period in days	[77, 90]

Community population	N	91,176	N/A	28,818	N/A	number of people in community	[91]
Proportion of critically ill infected patients	κ	0.11%	$[0.013\%, 3.2\%]$	0.11%	$[0.013\%, 3.2\%]$	% of infected individuals requiring critical care	[92, 93]

Computing methods

This model was built in the Python coding language, version 3.8.8. A complete list of packages used to support this model can be found in the Appendix. Local machines are capable of completing high volumes of single-scenario simulations and a few runs of multi-scenario simulations. However, a large number of runs for multi-scenario simulations requires a longer computing time than what is reasonable for the average local machine. Therefore we used our institution's high performance computing cluster to decrease runtime when finding our results. A public link to our GitHub repository for this study is available here: https://github.com/hnwestpage/ICU-simulator

3.3 Results

Sensitivity analysis

We conducted analysis of the model's sensitivity to perturbations in the baseline arrival rate, baseline departure rate, epidemic-related departure rate, queue abandonment rate, patient-to-clinician transmission probability, clinician recovery rate, and minimum staffing ratio on the resulting accumulated queue abandonments. Sensitivity testing was conducted separately for both the high resource and low resource assumption, as the ranges for the baseline arrival rate varies significantly between the two resource levels. We conducted testing assuming the initial outbreak scenario (Scenario B) in order to analyze sensitivity on the epidemic-related parameters. A baseline of queue abandonments was established with models using the mean values of the parameter ranges found from literature over 25 sample paths. The mean and standard deviation of the accumulated queue abandonment counts served as the baseline.

We list the resulting mean queue abandonments of the baseline low resource system, as well as the for the system with the minimum and maximum of each parameter in Table 3.3. The system displayed sensitivity on the baseline arrival rate, baseline departure rate (inverse of LOS), and infectious departure rate (inverse of LOS). All three of these parameters yielded queue abandonments outside of one standard deviation from the mean of the baseline system. The other parameters for queue abandonment, patient to HCP transmission, HCP recovery, and coverage ratio yielded means within one standard deviation of the mean of the baseline system.

We list the resulting mean queue abandonments of the baseline high resource

system, as well as the for the system with the minimum and maximum of each parameter in Table 3.4. The system also displayed sensitivity on the baseline arrival rate, baseline departure rate, and infectious departure rate. Like the low resource setting, all three of these parameters yielded queue abandonments outside of one standard deviation from the mean of the baseline system. The other parameters for queue abandonment, patient to HCP transmission, HCP recovery, and coverage ratio yielded means within one standard deviation of the mean of the baseline system.

Table 3.3: Course sensitivity analysis for the low resource setting

Parameter	Min	Max
$\lambda \in [0.5, 7]$	613	2292
$(\mu_1)^{-1} \in [2, 10]$	920	1608
$(\mu_2)^{-1} \in [4, 15]$	1067	1274
$\rho \in [0.01, 0.1]$	1191	1163
$\eta \in [0.005, 0.011]$	1192	1175
$(\nu)^{-1} \in [7, 10]$	1187	1162
$r \in [\frac{1}{8}, \frac{1}{11}]$	1172	1163
Baseline	Mean 1185	SD 45

Table 3.4: Course sensitivity analysis for the high resource setting

Parameter	Min	Max
$\lambda \in [4, 15]$	427	2676
$(\mu_1)^{-1} \in [2, 10]$	568	2744
$(\mu_2)^{-1} \in [4, 15]$	737	1066
$\rho \in [0.01, 0.1]$	943	983
$\eta \in [0.005, 0.011]$	966	963
$(\nu)^{-1} \in [7, 10]$	951	954
$r \in [\frac{1}{8}, \frac{1}{11}]$	969	948
Baseline	Mean 957	SD 49

After identifying λ, μ_1, and μ_2 as sensitive parameters in both resource settings, we performed the same testing method described earlier. We chose one parameter at a time to modify within its range from the literature, and compared the resulting projected queue abandonments to the baseline identified previously. Since the parameters identified were the same in both resource settings, we elected to conduct sensitivity analysis within the high resource setting.

In the first subplot of Figure 3.4, we find that queue abandonment increases monotonically with increased arrival rates. We notice a slight increase in slope at $\lambda = 9$, which is close to the mean value for that parameter range, 9.5. In the second subplot of Figure 3.4, we find that queue abandonment increases monotonically with increased patient LOS. The values for queue abandonment appear to form an "S" shape. In the third subplot of Figure 3.4 we also see that queue abandonment increases monotonically with increased epidemic-patient LOS. However, this increase is much less pronounced than that of the non-epidemic LOS in the second subplot.

This analysis informed our selection of parameters for the case scenarios discussed earlier. Based on these results, we chose parameters for modeling the six scenarios as detailed in Table 3.2.

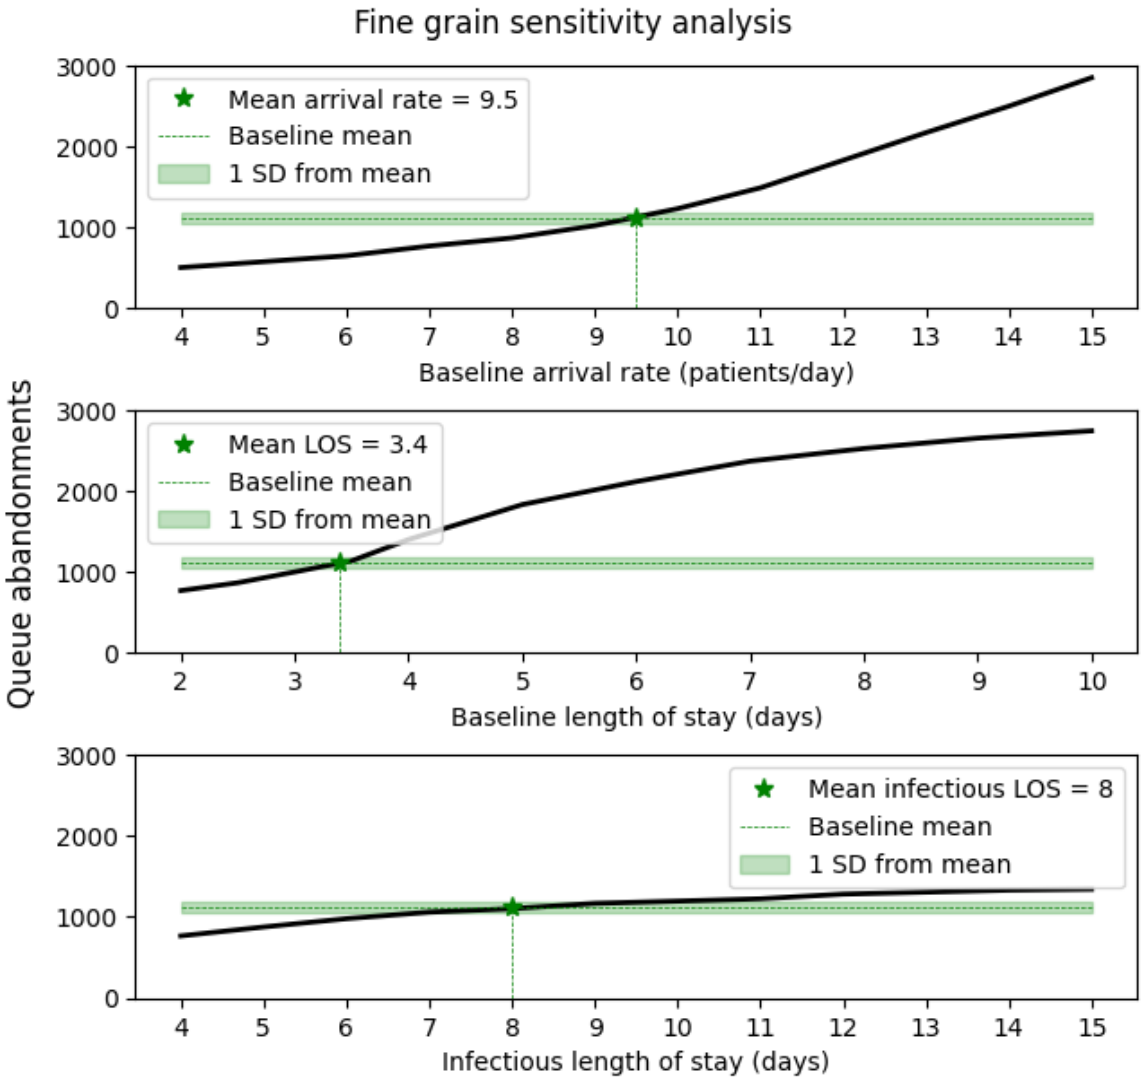

Figure 3.4: Fine-grain sensitivity analysis of baseline arrival rate, baseline departure rate, and epidemic departure rate

Low resource facility heatmaps and single models

We parametrized and simulated six total scenarios to demonstrate the queueing framework's various capabilities. These six scenarios compare a high and a low resource facility under three different conditions: a pre-epidemic baseline, an outbreak of a novel respiratory pathogen, and an outbreak of a novel respiratory pathogen with 20% increased transmissibility.

The high and low resource facilities are each constrained with an operating budget and costs related to their resource availability. An array of resource allocation combinations is simulated with each combination running 200 sample paths. The mean and standard deviation for each combination is calculated from the 200 sample paths. The mean values are arranged in the array and displayed as a heatmap. A black line representing the budget threshold of the facility is plotted over the heatmap, in which combinations that are within the budget threshold appear below the black line.

Figure 3.5 showcases the mean accumulated queue abandonments for each scenario within the low resource setting. For Scenario A (pre-epidemic baseline), the optimal resource allocation was 5 ICU clinicians and 10 ICU beds yielding a mean of 21.73 ($\pm$8.85 SD) queue abandonments. During Scenario B (initial outbreak), the optimal resource allocation was 5 ICU clinicians and 10 ICU beds yielding a mean of 177.39 ($\pm$23.07 SD) queue abandonments. In Scenario C (initial outbreak with increased transmissibility), the optimal resource allocation was 5 ICU clinicians and 10 ICU beds yielding a mean of 185.26 ($\pm$21.94 SD) queue abandonments.

Narrowing our focus to the low resource facility with 5 ICU clinicians and 10 ICU beds, we turn our attention to the following trends of the system: mean patients in

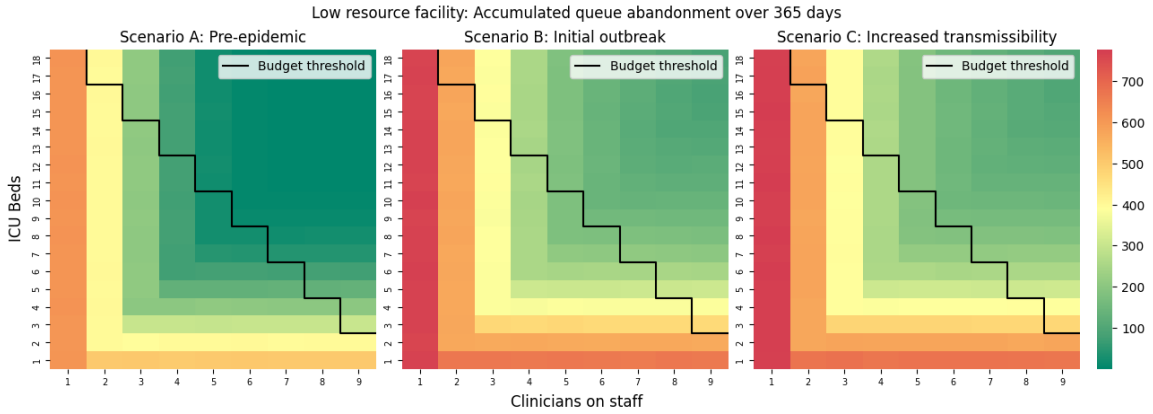

Figure 3.5: This figure displays the heatmaps of accumulated queue abandonments simulated over 365 days of a low resource facility in three case scenarios. Scenario A represents a pre-epidemic baseline in which there is no infectious disease among the community population. Scenario B represents an outbreak of a novel respiratory pathogen. Scenario C represents an outbreak of a novel respiratory pathogen with a 20% increase in transmissibility. Areas of dark green indicate the relative lowest accumulated queue abandonments, while areas of dark red indicate the highest.

queue, mean patients in ICU, mean accumulated queue abandonments, and mean number of available clinicians. Figures 3.6, 3.7, and 3.7 show the same system of 5 ICU clinicians and 10 ICU beds under the three scenarios. Each panel displays the simulated trajectories over 365 days. Light gray plots indicate the trajectories of individual sample paths, while bold plots are the mean trajectories over all 200 sample paths. Solid line mean trajectories are used to plot total numbers of patients tracked in the queueing system and the number of active ICU clinicians. Dashed line plots are used to track epidemic-related (infectious) patients within the queueing system.

Figure 3.6 Scenario A shows our low resource system under pre-epidemic conditions. This scenario establishes the baseline behavior of the system. The behavior

of all four aspects of the system are constant or linear. The average number of patients in the queue is <1 while ICU utilization remains stable at 6.75 beds or 67.5% utilization of the 10 bed capacity. We note that the parameter enforcing the maximum number of patients assigned to each clinician is set as $r = 9.3$. Since the model strictly maintains the constraint for the effective ICU capacity given by Equation (3.11), 5 clinicians may only service 9 patients. This is reflected as we see the effective capacity (depicted as a dotted black line in Figure 3.6) drop to 9 and remain there for the duration of the simulation. Given this, the actual ICU utilization maintains at 75% of the 9 bed capacity.

The number of infectious patients entering and departing the system remains at zero in the absence of a community epidemic. The mean accumulated abandonments by day 365 reaches 20.69 ($\pm$8.93 SD), consistent with the value given by the heatmap in Figure 3.5. The number of active ICU clinicians remains constant, as there are no infectious patients in the system to cause a transmission event. This constant behavior is also reflected in the consistency of the ICU capacity, as it is dependent on the number of active clinicians.

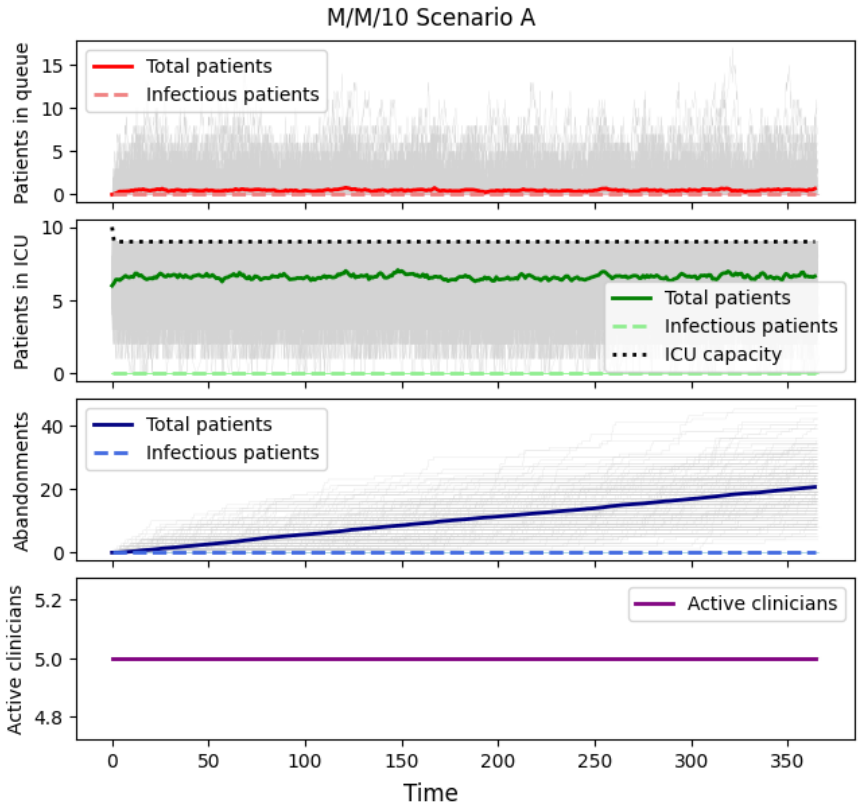

Figure 3.6: This figure displays the simulated behavior of a Markovian queueing model of an ICU with 10 beds. This low resource setting is simulated under the conditions of Scenario A, the pre-epidemic baseline case. The first subplot graphs the mean (solid red) and individual sample paths tracking the number of patients in the queue. The second subplot graphs the mean (solid green) and individual sample paths tracking the number of total patients in the ICU. The third subplot graphs the mean (solid blue) and individual sample paths tracking the number of total accumulated queue abandonments in the ICU. The fourth subplot graphs the mean (solid purple) number of active clinicians assigned to the ICU.

Figure 3.7 Scenario B shows our low resource system under epidemic conditions. This scenario shows the system burdened by the sudden influx of critically ill patients infected with a novel pathogen. The behavior of all four aspects of the system reflects the influence of the integrated SEIR model. The average number of patients in the

queue peaks at 40.4 with 29.7 patients (73.5%) being infectious on day 38. The mean queue behavior does not return to the system baseline value (<1) until day 88.

ICU utilization in Scenario B quickly exceeds capacity and remains over-saturated between days 20 through 65. During this period, infectious patients dominate the composition of the ICU accounting for 6.5 out of 8.6 (75.6%) of patients at peak on day 50. The mean ICU utilization does not return to the system baseline value (<6.75) until day 90.

Accumulated queue abandonments in Scenario B accelerates starting day 18 and returns to the baseline rate at day 78. The mean total accumulated abandonments by day 365 reaches 180.16 (± 20.79 SD), consistent with the value given by the heatmap in Figure 3.5. The mean infectious patient abandonments reaches 93.07 (± 11.48 SD) on day 76 and remains unchanged for the rest of the simulation. This shows that 51.7% of queue abandonments are attributed to infectious patients. The mean number of active ICU clinicians fluctuates from days 10 through 100, falling below a value of 4.8 from day 30 through day 67. The mean number of active clinicians reaches a minimum value of 4.73 on days 51, 53, and 58.

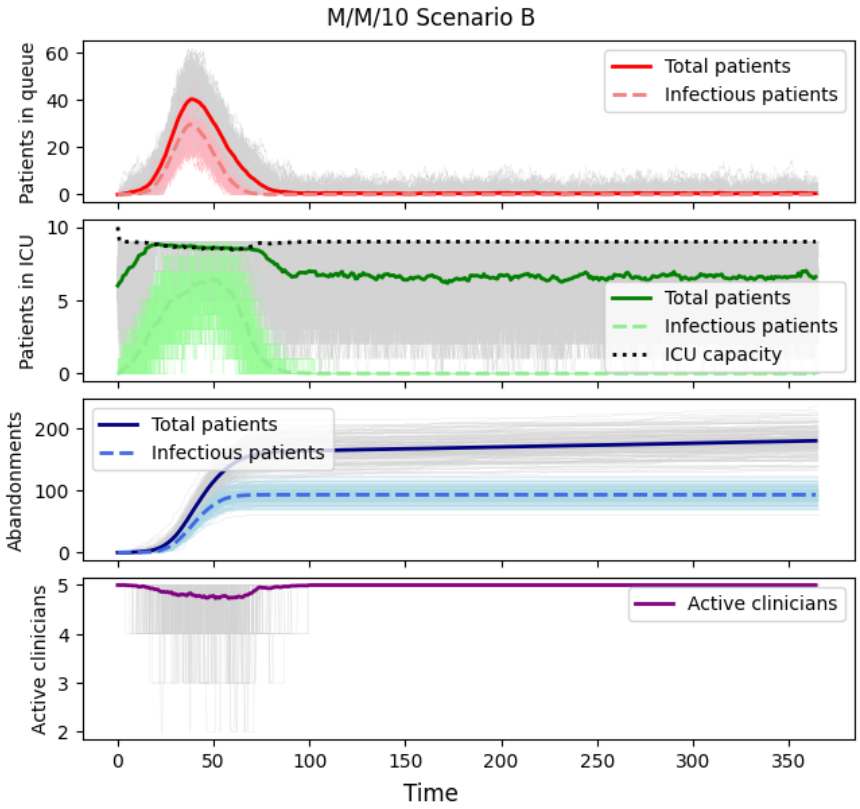

Figure 3.7: This figure displays the simulated behavior of a Markovian queueing model of an ICU with 10 beds. This low resource setting is simulated under the conditions of Scenario B, the initial outbreak case. The first subplot graphs the mean number of total patients in the queue (solid red), as well as mean number of infectious patients (dashed coral). The second subplot graphs the mean number of total patients in the ICU (solid green), as well as mean number of infectious patients (dashed light green). The third subplot graphs the mean number of total accumulated queue abandonments (solid blue), as well as mean number of accumulated abandonments who were infectious patients (dashed light blue). The fourth subplot graphs the mean (solid purple) number of active clinicians assigned to the ICU.

Figure 3.8 Scenario C shows our low resource system under modified epidemic conditions. This scenario shows the system burdened by the sudden influx of critically ill patients infected with a novel pathogen 20% more transmissible than that of

Scenario B. The average number of patients in the queue peaks at 47.2 with 36.2 patients (76.7%) being infectious on day 34. The mean queue behavior does not return to the system baseline value (<1) until day 82.

ICU utilization in Scenario C quickly exceeds capacity and remains over-saturated between days 16 through 60. During this period, infectious patients dominate the composition of the ICU accounting for 6.8 out of 8.6 (79.1%) of patients at peak on day 47. The mean ICU utilization does not return to the system baseline value (<6.75) until day 90.

Accumulated queue abandonments in Scenario C accelerates starting day 15 and returns to the baseline rate at day 70. The mean total accumulated abandonments by day 365 reaches 185.13 (± 21.55 SD), consistent with the value given by the heatmap in Figure 3.5. The mean infectious patient abandonments reaches 100.91 (± 12.24 SD) on day 66 and remains unchanged for the rest of the simulation. This shows that 54.5% of queue abandonments are attributed to infectious patients. The mean number of active ICU clinicians fluctuates from days 5 through 89, falling below a value of 4.8 from day 35 through day 60. The mean number of active clinicians reaches a minimum value of 4.76 on days 46, 51, and 55.

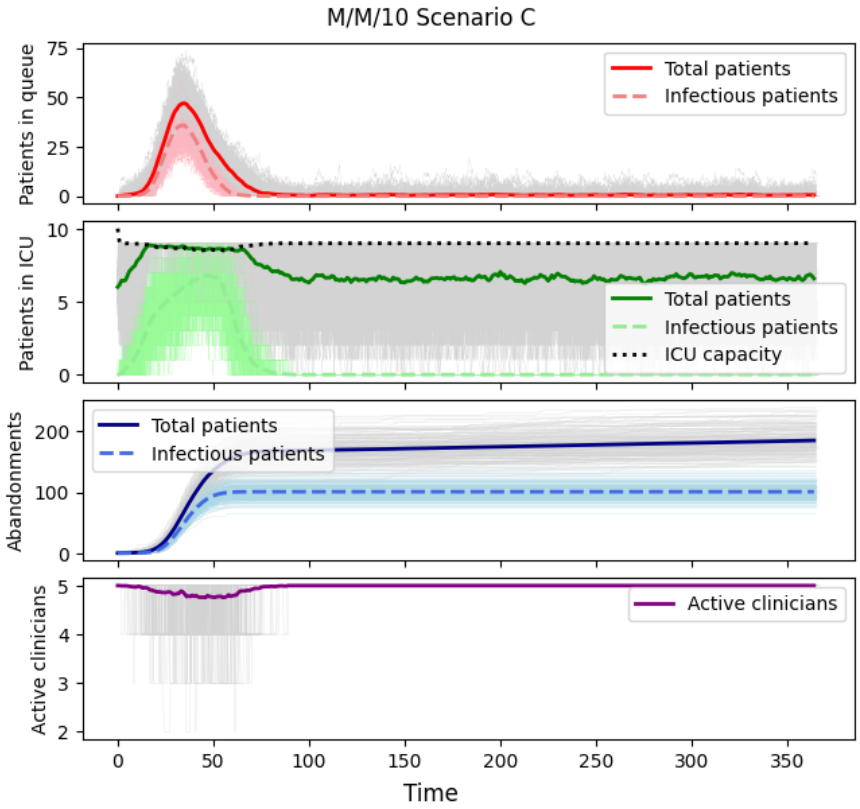

Figure 3.8: This figure displays the simulated behavior of a Markovian queueing model of an ICU with 10 beds. This low resource setting is simulated under the conditions of Scenario C, the initial outbreak case with higher transmissibility. The first subplot graphs the mean number of total patients in the queue (solid red), as well as mean number of infectious patients (dashed coral). The second subplot graphs the mean number of total patients in the ICU (solid green), as well as mean number of infectious patients (dashed light green). The third subplot graphs the mean number of total accumulated queue abandonments (solid blue), as well as mean number of accumulated abandonments who were infectious patients (dashed light blue). The fourth subplot graphs the mean (solid purple) number of active clinicians assigned to the ICU.

High resource facility heatmaps and single models

Figure 3.9 showcases the mean accumulated queue abandonments for each scenario within the high resource setting. Within Scenario A (pre-epidemic baseline), the optimal resource allocation was 15 ICU clinicians and 29 ICU beds yielding a mean of 34.63 ($\pm$13.24) queue abandonments. Within Scenario B (initial outbreak), the optimal resource allocation was 15 ICU clinicians and 29 ICU beds yielding a mean of 575.34 ($\pm$40.66) queue abandonments. Within Scenario C (increased transmissibility), the optimal resource allocation was 15 ICU clinicians and 29 ICU beds yielding a mean of of 603.35 ($\pm$43.17) queue abandonments. A summary figure of the heatmaps for both the low and high resource settings is found in Figure 3.13.

Narrowing our focus to the high resource facility with 15 ICU clinicians and 29 ICU beds, we again turn our attention to the following trends of the system: mean patients in queue, mean patients in ICU, mean accumulated queue abandonments, and mean number of available clinicians. Figures 3.10, 3.11, and 3.12 show the same system of 15 ICU clinicians and 29 ICU beds under the three scenarios. Each panel displays the simulated trajectories over 365 days. Light gray plots indicate the trajectories of individual sample paths, while bold plots are the mean trajectories over all 200 sample paths. Solid line mean trajectories are used to plot total numbers of patients tracked in the queueing system and the number of active ICU clinicians. Dashed line plots are used to track epidemic-related (infectious) patients within the queueing system.

Figure 3.10 Scenario A shows our high resource system under pre-epidemic conditions. This scenario establishes the baseline behavior of the high resource system.

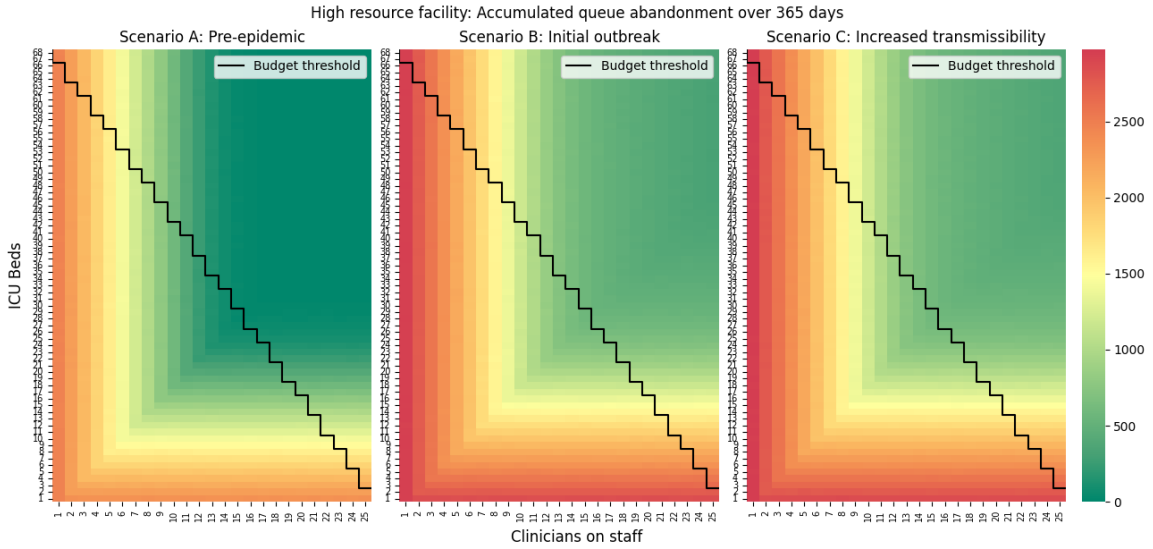

Figure 3.9: This figure displays the heatmaps of accumulated queue abandonments simulated over 365 days of a high resource facility in three case scenarios. Scenario A represents a pre-epidemic baseline in which there is no infectious disease among the community population. Scenario B represents an outbreak of a novel respiratory pathogen. Scenario C represents an outbreak of a novel respiratory pathogen with a 20% increase in transmissibility. Areas of dark green indicate the relative lowest accumulated queue abandonments, while areas of dark red indicate the highest.

The behavior of all four aspects of the system are constant or linear. The average number of patients in the queue is <1 while ICU utilization remains stable at 24 beds or 82.8% utilization. The number of infectious patients entering and departing the system remains at zero in the absence of a community epidemic. The mean accumulated abandonments by day 365 reaches 33.34 ($\pm$12.9), consistent with the value given by the heatmap in Figure 3.9. The number of active ICU clinicians remains constant at 15, as there are no infectious patients in the system to cause a transmission event.

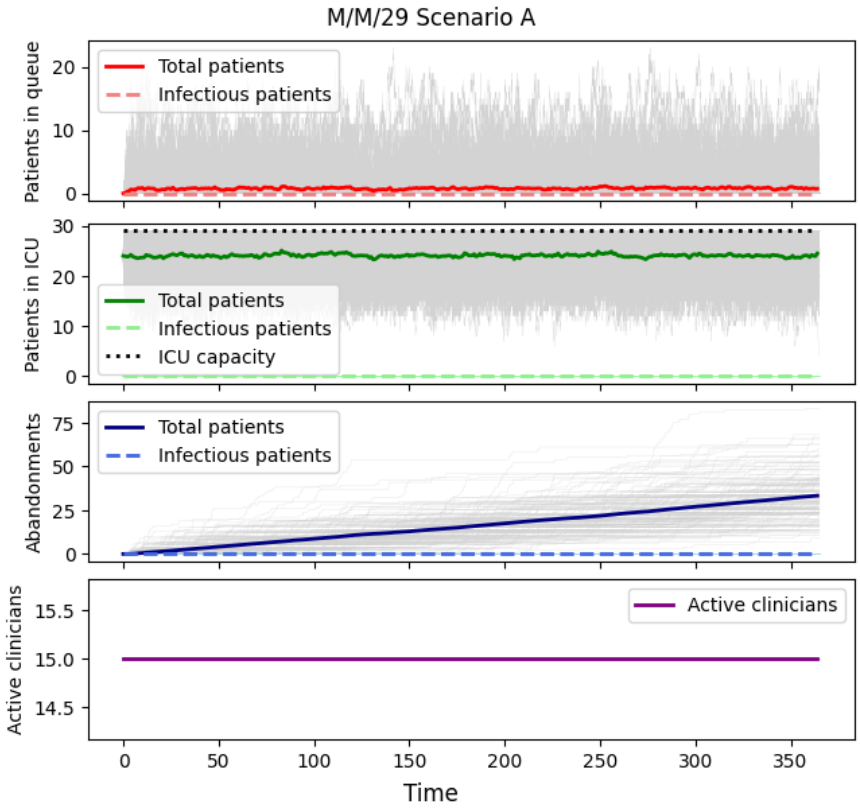

Figure 3.10: This figure displays the simulated behavior of a Markovian queueing model of an ICU with 29 beds. This high resource setting is simulated under the conditions of Scenario A, the pre-epidemic baseline case. The first subplot graphs the mean (solid red) and individual sample paths tracking the number of patients in the queue. The second subplot graphs the mean (solid green) and individual sample paths tracking the number of total patients in the ICU. The third subplot graphs the mean (solid blue) and individual sample paths tracking the number of total accumulated queue abandonments in the ICU. The fourth subplot graphs the mean (solid purple) number of active clinicians assigned to the ICU.

Figure 3.11 Scenario B shows our high resource system under epidemic conditions. This scenario shows the system burdened by the sudden influx of critically ill patients infected with a novel pathogen. The behavior of all four aspects of the system reflects the influence of the integrated SEIR model. The average number of patients in the

queue peaks at 129.1 with 88.6 patients (68.6%) being infectious on day 40. The mean queue behavior does not return to the system baseline value (<1) until day 94.

ICU utilization in Scenario B quickly exceeds capacity and remains over-saturated from days 15 to 80. During this period, infectious patients dominate the composition of the ICU accounting for 20.2 out of 27.2 (74.3%) of patients at peak on day 53. The mean ICU utilization does not return to the system baseline value (<24) until day 98.

Accumulated queue abandonments in Scenario B accelerates starting day 15 and returns to the baseline rate at day 80. The mean total accumulated abandonments by day 365 reaches 566.37 (± 42.81 SD), consistent with the value given by the heatmap in Figure 3.9. The mean infectious patient abandonments reaches 287.67 (± 20.57 SD) on day 69 and remains unchanged for the rest of the simulation. This shows that 50.8% of queue abandonments are attributed to infectious patients. The mean number of active ICU clinicians fluctuates from days 7 through 88, falling below a value of 14.8 from day 18 through day 76. The mean number of active clinicians reaches a minimum value of 14.1 on days 47, 51, and 58.

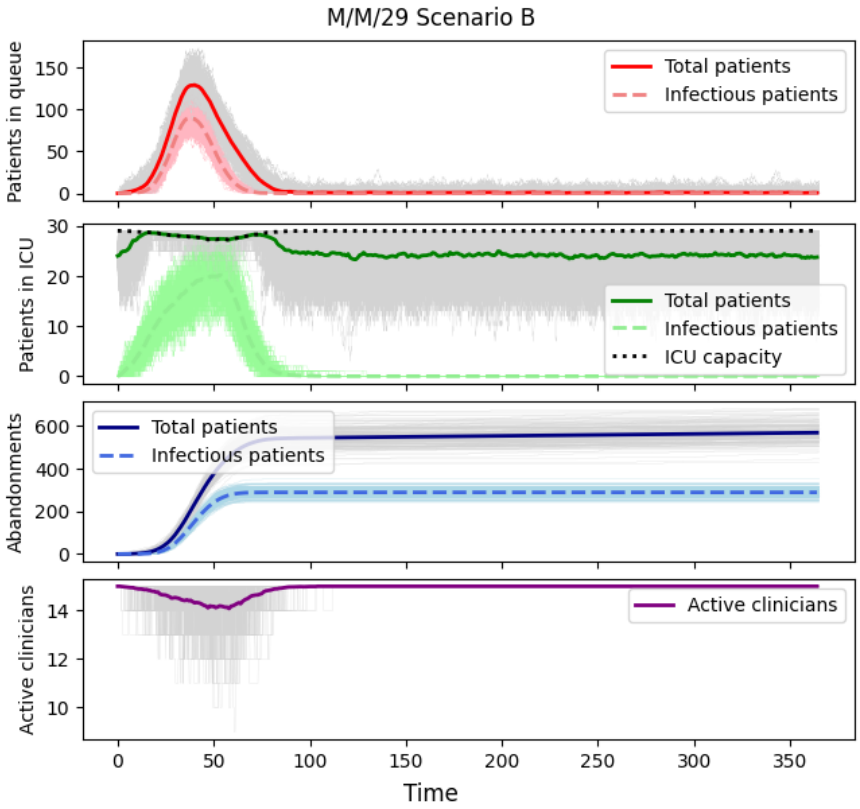

Figure 3.11: This figure displays the simulated behavior of a Markovian queueing model of an ICU with 29 beds. This high resource setting is simulated under the conditions of Scenario B, the initial outbreak case. The first subplot graphs the mean number of total patients in the queue (solid red), as well as mean number of infectious patients (dashed coral). The second subplot graphs the mean number of total patients in the ICU (solid green), as well as mean number of infectious patients (dashed light green). The third subplot graphs the mean number of total accumulated queue abandonments (solid blue), as well as mean number of accumulated abandonments who were infectious patients (dashed light blue). The fourth subplot graphs the mean (solid purple) number of active clinicians assigned to the ICU.

Figure 3.12 shows our high resource system under modified epidemic conditions. This scenario shows the system burdened by the sudden influx of critically ill patients infected with a novel pathogen 20% more transmissible than that of Scenario B. The

average number of patients in the queue peaks at 153.3 with 114.1 patients (74.4%) being infectious on day 34. The mean queue behavior does not return to the system baseline value (<1) until day 90.

ICU utilization in Scenario C quickly exceeds capacity and remains over-saturated for days 13 through 69. During this period, infectious patients dominate the composition of the ICU accounting for 22.02 out of 27.4 (80.4%) patients on average at peak on day 45. The mean ICU utilization does not return to the system baseline value (24) until day 91.

Accumulated queue abandonments in Scenario C accelerates starting day 14 and returns to the baseline rate at day 80. The mean total accumulated abandonments by day 365 reaches 599.52 ($\pm$39.99), consistent with the value given by the heatmap in Figure 3.9. The mean infectious patient abandonments reaches 322.41 ($\pm$22.27) on day 65 and remains unchanged for the rest of the simulation. This shows that 53.8% of queue abandonments are attributed to infectious patients. The mean number of active ICU clinicians fluctuates from days 9 through 81, falling below a value of 14.8 from day 16 through day 71. The mean number of active clinicians reaches a minimum value of 14.0 on days 48 and 49. A summary figure of each single scenario model discussed here can be found in Figure 3.14.

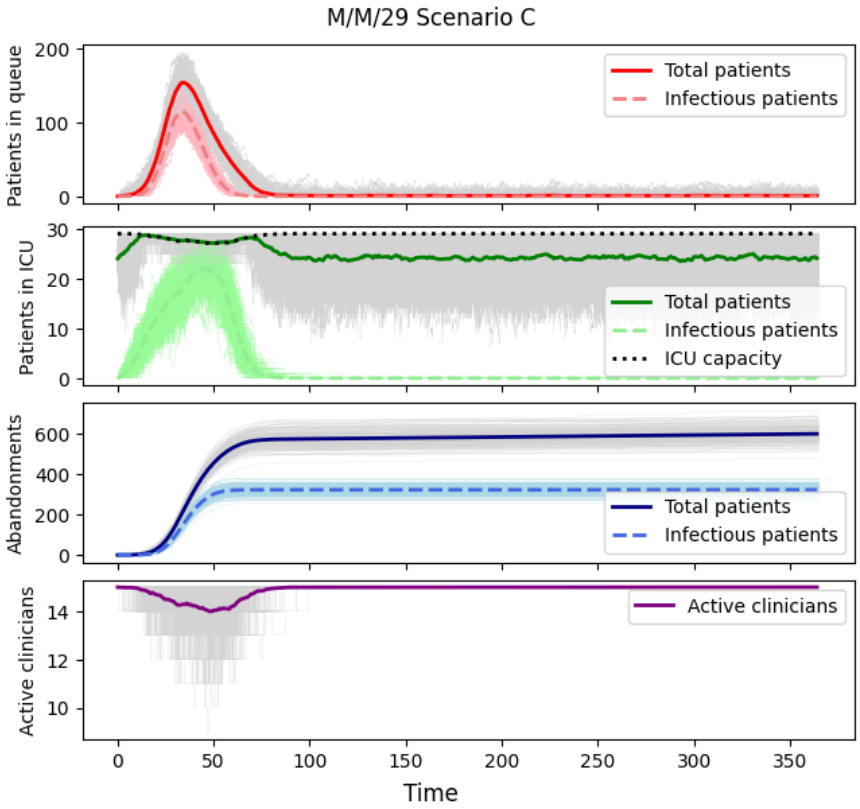

Figure 3.12: This figure displays the simulated behavior of a Markovian queueing model of an ICU with 29 beds. This high resource setting is simulated under the conditions of Scenario C, the initial outbreak case with higher transmissibility. The first subplot graphs the mean number of total patients in the queue (solid red), as well as mean number of infectious patients (dashed coral). The second subplot graphs the mean number of total patients in the ICU (solid green), as well as mean number of infectious patients (dashed light green). The third subplot graphs the mean number of total accumulated queue abandonments (solid blue), as well as mean number of accumulated abandonments who were infectious patients (dashed light blue). The fourth subplot graphs the mean (solid purple) number of active clinicians assigned to the ICU.

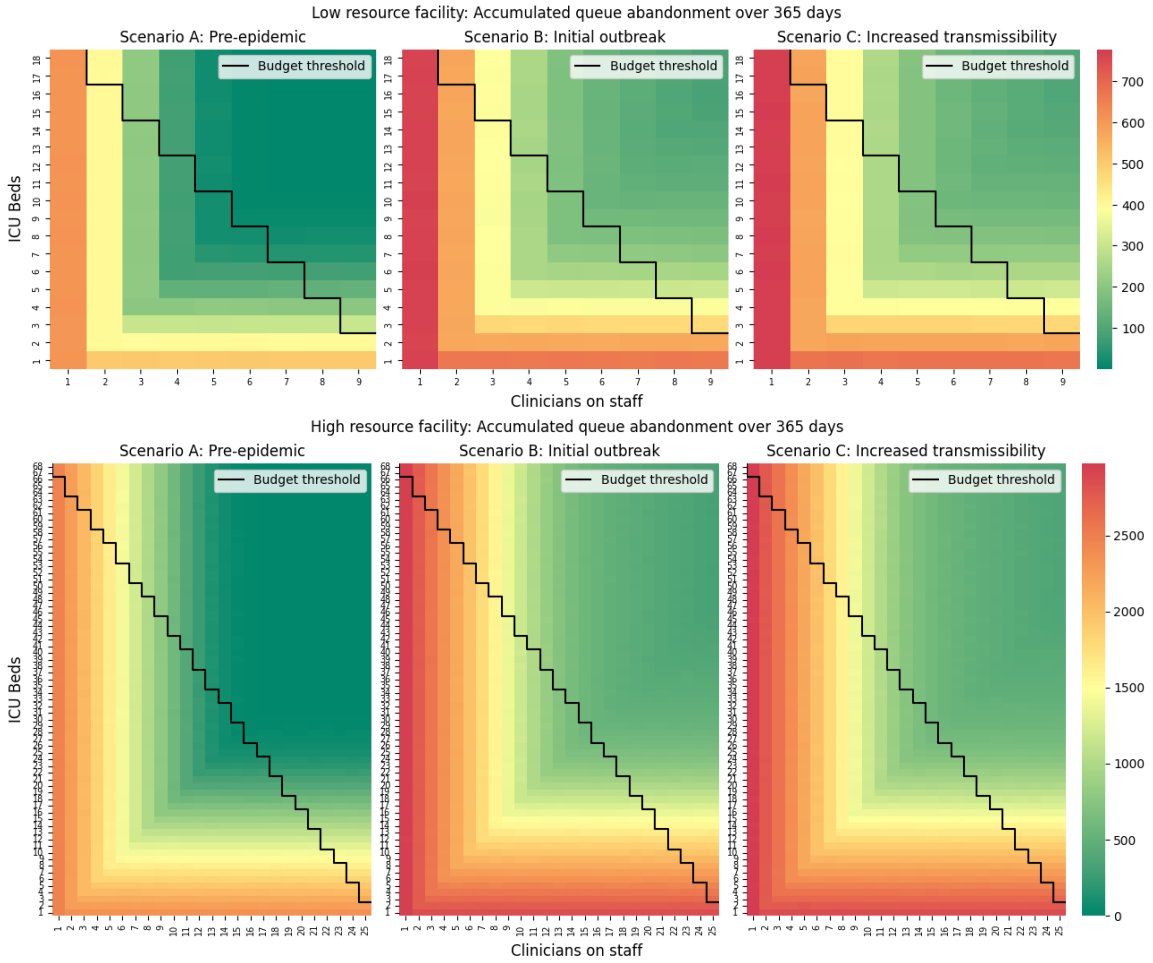

Figure 3.13: Summary of heatmaps displaying the mean accumulated queue abandonments for a diverse array of resource allocations at both the low resource setting (top) and high resource setting (bottom).

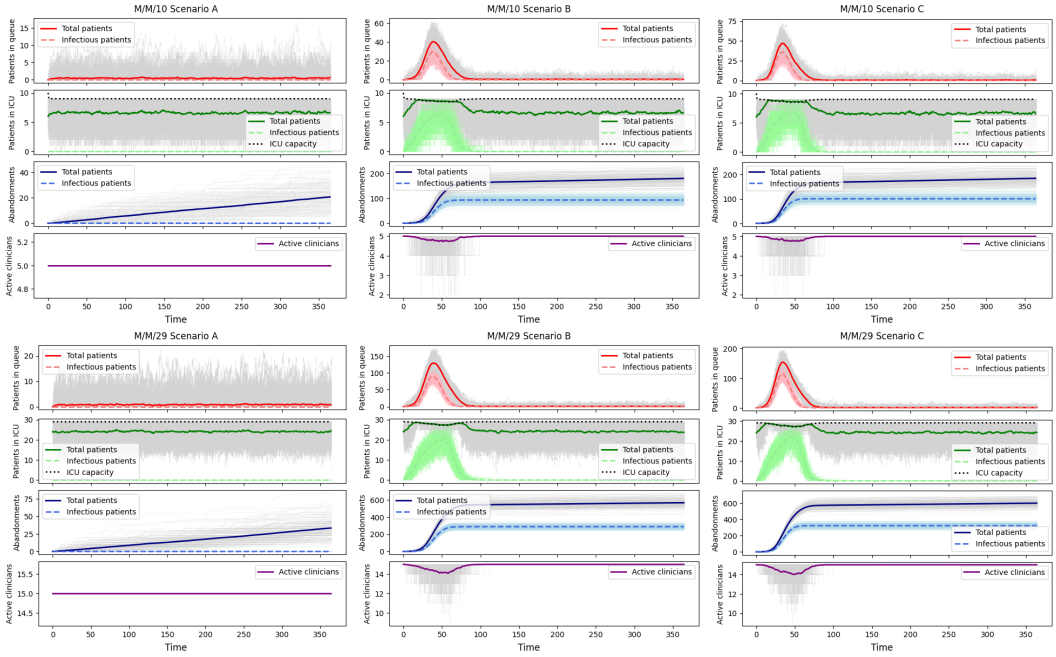

Figure 3.14: Summary of single queueing models with the top row representing the low resource setting and the bottom row representing the high resource setting. For either facility, scenario A (baseline) is left, scenario B (initial outbreak) is middle, and scenario C (increased transmissibility) is right.

3.4 Discussion

Heatmaps

We simulated and analyzed six different scenarios to demonstrate the queueing framework's various capabilities. These six scenarios compare a high and a low resource setting under three different conditions: a pre-epidemic baseline, an outbreak of a novel respiratory pathogen, and an outbreak of a novel respiratory pathogen with 20% increased transmissibility. Within this demonstration, we showed our model's ability to provide recommendations for resource allocations that minimize queue abandonment during a simulated epidemic.

In both the high and low resource settings, mean accumulated queue abandonments are noticeably higher in scenarios B and C compared to scenario A. In all three scenarios, the optimal resource allocation paring remained consistent within the resource setting, with 5 clinicians and 10 beds in the low resource facilities and 15 clinicians and 29 beds in the high resource facilities. This preference for a balance of resources over maximizing capacity is the result of the minimum staffing requirements enforced by Equation (3.11) for n^* .

We note that despite the difference between scenarios B and C being a 20% increase in transmissibility of the novel pathogen, this did not result in a 20% increase in queue abandonments. For the low resource facility the mean accumulated abandonments were 177.39 ($\pm$ 23.07) and 185.26 ($\pm$ 21.94) for scenarios B and C, respectively. Likewise, for the high resource facility the mean accumulated abandonments were 575.34 ($\pm$ 40.66) and 603.35 ($\pm$ 43.17) for scenarios B and C, respectively. In either setting, while the means for scenario C are greater than those for scenario

B, they are within the standard deviations of those means. This indicates that the mean accumulated queue abandonments are not substantially different between scenarios B and C. Further, this suggests there is a non-linear relationship between the transmissibility of the pathogen and the expected queue abandonments.

In our multi-facility simulations, we model resource allocations both within and outside of the operating budget threshold. We chose not to exclude the extreme cases that fall well-under or well-over the budget threshold for the sake of capturing the full range of possible outcomes. Given this, we interpret these extreme resource allocations as the most pessimistic or most optimistic settings. In settings where there is an imbalance of clinicians and beds, particularly too few clinicians to meet the needs of ICU capacity, we interpret this as being an ICU which lacks 24/7 coverage by a specialized clinician. Some studies have shown that this model for an ICU in practice may result in a higher risk for patient mortality [94]. Another interpretation of a facility with too few ICU clinicians is that the clinicians are overworked in order to meet the needs of the patients. This case also poses risks to the patients, as literature shows that over-worked healthcare personnel is related to a higher risk of patient mortality [58].

By modeling all possible resource configurations for a set budget, we can use the outcomes to determine the optimal resource allocation to minimize queue abandonments. Recall that queue abandonments may be interpreted as patients recovering in a general medical ward, being transferred to another facility, or suffering from mortality [62]. For a patient in critical condition, both recovery in the general medical ward and transfer to another facility are associated with an increased risk in

patient mortality [62]. We seek to avoid any of these events when possible. To use the heatmaps as a tool for emergency preparedness, we provide an example of an analysis for the low resource and high resource settings based on their suggested optimal allocation.

Simulation for the low resource resource setting found that the ICU with 5 clinicians and 10 ICU beds minimized queue abandonments subject to the operating budget constraint in all three scenarios. Recall that according to Equation (3.11) for the minimum safe staffing levels, a staff 5 clinicians can only safely treat a maximum of 9 patients. To prepare for an emergency like scenario B, we can compare the projected queue abandonments of (5,10) with its neighboring allocations. Recall that the mean accumulated queue abandonments for the (5,10) facility was 177.39 ($\pm$23.07) in scenario B. In the same heatmap, if we add one clinician to yield a total of 6 ICU clinicians with 10 beds, we reduce the projected queue abandonments to 155.27 ($\pm$ 20.20).

This reduction is noticeable, but not fully outside of the standard deviation for that expected value. The closest neighboring allocation that is less than one standard deviation below the mean is the facility 6 ICU clinicians and 11 ICU beds, meaning one additional clinician and one additional bed. In this case, we reduce the projected queue abandonments to 142.32 ($\pm$ 20.04). This a significant difference from the (5,10) facility, as it reduced the expected abandonments by 19.8% for the cost of \$300,000 according to Table 3.2.

To prepare for an emergency like scenario C, we again compare the projected queue abandonments of (5,10) with its neighboring allocations. The mean accumulated

queue abandonments for the (5,10) facility was 185.26 (±21.94) in scenario C. In the same heatmap, if we add one clinician to yield a total of 6 ICU clinicians with 10 beds, we reduce the projected queue abandonments to 165.02 (± 20.79).

Like scenario B, this reduction is noticeable, but not fully outside of the standard deviation for that expected value. By adding one clinician and one bed, yielding a total of 6 ICU clinicians and 11 ICU beds, we reduce the projected queue abandonments to 150.59 (± 20.75). This a significant difference from the (5,10) facility, as it reduced the expected abandonments by 18.7% for the cost of $300,000 according to Table 3.2.

Simulation for the high resource resource setting found that the ICU with 15 clinicians and 29 ICU beds minimized queue abandonments subject to the operating budget constraint in all three scenarios. To prepare for an emergency like scenario B, we can compare the projected queue abandonments of (15,29) with its neighboring allocations. Recall that the mean accumulated queue abandonments for the (15,29) facility was 575.34 (± 40.66) in scenario B. In the same heatmap, if we add one clinician to yield a total of 16 ICU clinicians with 29 beds, we reduce the projected queue abandonments to 561.62 (± 38.51).

This reduction is less noticeable, and well within the standard deviation for that expected value. The closest neighboring allocation that is less than one standard deviation below the mean is the facility with 16 ICU clinicians and 31 ICU beds, meaning one additional clinician and two additional beds. In this case, we reduce the projected queue abandonments to 528.34 (± 38.29). This a meaningful difference from the (15,29) facility, as it reduced the expected abandonments by 8.1% for the

cost of $700,000 according to Table 3.2.

To prepare for an emergency like scenario C, we again compare the projected queue abandonments of (15,29) with its neighboring allocations. The mean accumulated queue abandonments for the (15,29) facility was 603.45 (± 43.17) in scenario C. The closest neighboring allocation that is less than one standard deviation below the mean is also the facility with 16 ICU clinicians and 31 ICU beds, meaning one additional clinician and two additional beds. In this case, we reduce the projected queue abandonments to 557.90 (± 42.14). This a meaningful difference from the (15,29) facility, as it reduced the expected abandonments by 7.5% for the cost of $700,000 according to Table 3.2.

Single models

After analyzing the multi-facility heatmaps, we simulated the optimal resource allocation for both resource settings under the three scenarios. We note that even during the pre-epidemic baseline, we see a non-zero number of patients in the queue and patients abandoning the queue in both resource settings. This can be attributed to the expected ICU utilization levels being 75% in the low resource facility and 82% in the high resource facility. These utilization rates are high enough that due to the randomness of the system, the system becomes strained even without a community epidemic.

In general, we noticed a few consistent differences between scenario B and C in both resource settings. In short, mean projected values for peak queue length, percent of infectious patients in the queue and ICU, and queue abandonments were slightly higher in scenario C than B. Additionally, the periods of time where the system

fluctuates outside of its expected baseline behavior tend to occur earlier in scenario C than scenario B. This reflects the difference in the disease dynamics between the pathogens simulated in the two scenarios as seen in Figure 3.2. In this discussion section, we perform a pairwise comparison of the notable values given by the single-facility models.

First, we compare trends in the queue. In the low resource facility, there were more patients waiting in the queue at peak in scenario C (47.2) than scenario B (40.4). Additionally, we found that the queue reached its peak length 4 days earlier in the simulated timeline in scenario C (day 34) than scenario B (38). At this peak day, infectious patients accounted for slightly more of the total number of patients in the queue for scenario C (76.7%) than scenario B (73.5%).

In the high resource facility, we also found there were more patients waiting in the queue at peak in scenario C (153.3) than scenario B (129.1). Similar to the low resource facility, we found that the queue reached its peak length 6 days earlier in the simulated timeline in scenario C (day 34) than scenario B (day 40). At this peak day, infectious patients accounted for slightly more of the total number of patients in the queue for scenario C (74.4%) than scenario B (68.6%).

Next, we compare trends in the ICU utilization. In the low resource facility, the ICU reaches capacity 4 days earlier in scenario C (day 16) than scenario B (day 20). We see that the system returns to baseline 5 days earlier in scenario C (day 60) than scenario B (day 65). In both scenarios B and C, the ICU capacity reaches a minimum of 8.6. However, the mean composition of infectious patients in the ICU is slightly higher in scenario C (79.1%) than scenario B (75.6%).

In the high resource facility, the ICU reaches capacity 2 days earlier in scenario C (day 13) than scenario B (day 15). We see that the system returns to baseline 11 days earlier in scenario C (day 69) than scenario B (day 80). In scenarios B and C, the ICU capacity reaches a minimum of 27.2 and 27.4, respectively. The mean composition of infectious patients in the ICU is slightly higher in scenario C (80.4%) than scenario B (74.3%).

Next, we compare the trends in the accumulated queue abandonments. In the low resource facility, queue abandonments begin to accelerate 3 days earlier in scenario C (day 15) than scenario B (day 18). We see that the abandonment rate returns to baseline 8 days earlier in scenario C (day 70) than scenario B (day 78). Queue abandonments are slightly higher in Scenario C (100.91) than Scenario B (93.07). The mean composition of infectious patients is also slightly high in Scenario C (54.5%) than Scenario B (51.7%).

In the high resource facility, queue abandonments begin to accelerate 1 day earlier in scenario C (day 14) than scenario B (day 15). We see that the abandonment rate returns to baseline around the same time (day 80). Queue abandonments are slightly higher in Scenario C (322.41) than Scenario B (287.67). As seen with the low resource setting, the mean composition of infectious patients is also slightly high in Scenario C (53.8%) than Scenario B (50.8%).

Lastly, we compare trends in the mean active clinicians. In both resource settings, the minimum average number of clinicians in each resource setting is about the same for both scenarios. However, the projected times at which these minimums are reached are slightly different. In the low resource setting, the minimum number

of clinicians in scenario B (4.73 out of 5) is reached on days 51, 53, and 58. This is slightly later than when the number of active clinicians reaches a minimum in scenario C (4.76 out of 5) on days 46, 51, and 55. We see a similar pattern in the high resource setting. In this simulated facility, the minimum number of clinicians in scenario B (14.1 out of 15) is reached on days 47, 51, and 58. This is about the same if not slightly later than when the number of active clinicians reaches a minimum in scenario C (14.0 out of 15) on days 48 and 49.

Sensitivity analysis

Sensitivity analysis showed model sensitivity to the baseline arrival rate, the baseline departure rate, and the epidemic departure rate. This result aligns with the foundation of the queueing model, particularly the traffic density ratio, $\frac{\lambda}{n\mu}$. Perturbations in these three rates can drastically change this ratio, especially the departure rates. We see slightly less sensitivity in the epidemic departure rate. We hypothesize that this is due to the short time in which epidemic related individuals arrive to the ICU, contrasted by the static non-epidemic arrivals. Sensitivity to these fundamental model parameters suggests high importance in accurate parametrization.

We evaluated the sensitivity of this model on a subset of its parameters. We did not test the sensitivity on meta-parameters such as the ICU budget and cost variables. This aspect is partially captured by our multi-facility heatmaps they represent different budgets. Using these heatmaps, we can interpolate model projections based on changes in budget or hospital cost parameters by redrawing the budget threshold line. We also did not evaluate model sensitivity on the SEIR parameters, such as pathogen incubation or recovery periods. While we compared two transmission

coefficients, this is insufficient for generalization.

Limitations and future work

There are several limitations and opportunities for future work generated by this study. The first limitation to discuss is the error generated by 200 sample paths. This is referred to as the Monte Carlo simulation error caused by stochastic variability. The standard error (SE) is calculated by

$$SE = \frac{1}{\sqrt{M}}$$

where M is the number of sample paths. By setting $M = 200$, we expect that $SE = 0.070711$ or around 7.1%. This error may introduce confounding effects in the trends we identified comparing the high and low resource settings, as well as scenarios B an C. To improve this error to a lower threshold, a larger set of sample paths is necessary.

We accepted a 7.1% error for 200 runs based on another limitation of this model. To generate a high resource heatmap, it required 109.9 hours of computation time on a high performance computing cluster. This could be improved through moderate to major restructuring of the code for better parallelization. More parallelization of tasks would reduce the runtime allowing for a larger number of sample paths, thus reducing the SE.

This model directly depends on the simulated disease transmission within the community population. We introduced this dependence with a standard deterministic model of the disease spread. In future studies, this framework could be improved by

introducing a time-delay component when estimating the number of infected individuals arriving to the ICU. Currently, the model assumes a proportion of infected individuals at time t will arrive to the ICU at that same time. We acknowledge that this simplification does not reflect reality. The lack of time-delay limits the strength of our reported peak times of system stress. Introducing a time-delay between the infection curve and the epidemic related arrivals may provide a more robust projection of the timing and degree of stress on the healthcare system.

Future work can also improve this framework by exploring more advanced models of community disease transmission. Particularly, this model may be applied to one of the specific pathogens mentioned earlier, thus providing the opportunity for customization and a robust validation. Our results would be more informative when using disease transmission models that incorporated realistic pathogen specific phenomenons such seasonality, or control measures like social distancing and vaccination. This advancement would potentially provide more accurate projections and recommendations. At the same time, more advanced disease models may introduce new parameters the model may be sensitive to. Improvements on this framework may require further analysis of the model's sensitivity to inputs from the simulated disease transmission.

Our method of sensitivity analysis may not provide enough information about the sensitivity of the framework to certain parameters.

While we performed both a course and fine grain analysis, future work may explore model sensitivity using finer parameter ranges. In particular, we tested mostly only integer values for the baseline arrival rate and patient length of stay. A finer

analysis testing rational non-integer parameter values may provide a more informative result. An addition, we did not test the framework's sensitivity to parameters associated with the SEIR model. In future works, new analyses on model sensitivity to transmission rate, latent period, and recovery period will provide information to fill this knowledge gap.

We also acknowledge a potential limitation in our method of conducting sensitivity analysis. In our course analysis, we did not identify the individual queue abandonment rate (ρ) as a sensitive parameter. This conclusion was drawn by observing the estimated maximum and minimum queue abandonments to be within one standard deviation of the mean baseline model. However during experimentation, we noticed non-monotonic behavior in queue abandonments when testing various values for ρ. We hypothesize that extreme low and extreme high values for ρ may result in equivalent values for the overall system queue abandonment rate, ρ_k.

To highlight the reasoning for this hypothesis, recall Equation (3.4) from our description of the queueing foundation, $\rho_k = k\rho$ where k is the number of patients waiting in the queue. Consider two $M/M/n$ systems with individual queue abandonment rates, $\rho_{\text{low}} < \rho_{\text{high}}$, and λ, μ, n held the same. For the system given ρ_{low}, we initially expect to observe fewer queue abandonments than those observed in the system given ρ_{high}. However depending on server utilization, the queue may build more quickly resulting in an increased overall abandonment rate, ρ_k. Similarly, for the system with ρ_{high}, we initially expect to observe more queue abandonments than the system with ρ_{low}. But depending on the same server utilization, the higher individual abandonment rate may prevent the queue length from growing as large.

In either case, we acknowledge this feedback loop between the choice of ρ and the estimated queue abandonments. This non-monotonic behavior is not yet fully understood, and requires significant further investigation. Future work can reveal the relationship between the individual and system queue abandonment rate through a thorough sensitivity analysis over ranges of not only ρ, but also $\lambda, \mu,$ and n.

As previously mentioned, we parametrized this model with values obtained through a thorough search of published literature. While we designed these resource settings and epidemic scenarios with realism as the first priority, many choices were based on the model's sensitivity to certain parameters. We found wide and varied ranges for parameters such as the arrival and departure rates, of which this model is sensitive to. Because this calibration of the model does not represent a unique hospital system, we should treat these results as a demonstration and not a direct recommendation.

In order to move towards a direct application of this framework, in the future we aim to work with published public health data of case counts and hospitalizations during the COVID-19 pandemic in the United States. Specifically using hospital data from [COVID-19 Reported Patient Impact and Hospital Capacity by Facility [95]], we aim to estimate the system parameters such as the arrival and departure rates to calibrate the model. We can also use this data to validate the model by comparing the projections with the actual values. This will also lead to a more robust quantification of the error and biases our framework may generate.

Pertaining to parameterization and model calibration, we acknowledge this framework relies on prior knowledge of the disease dynamics and pathology, which might not be available before or at the beginning of a public health emergency. In practice,

this model requires parametrization using forecasted values. Given the model's sensitivity to the projected infectious patients and their length of stay, this may render the resource allocation predictions less reliable. As previously mentioned, further uncertainty quantification and reduction of error would be required to combat this challenge.

3.5 Conclusion

The evolving threats to public health and the healthcare system necessitate accurate understandings of emergency preparedness. Healthcare resources, specifically critical care beds and healthcare personnel, must be allocated in ways that prevent adverse patient outcomes and risk to healthcare providers. We attempted to develop a queueing theoretic framework that addresses these challenges. Our expanded queueing system can provide recommendations for optimal resource allocation for a wide variety of resource settings and public health emergencies. With further validation and calibration, this flexible framework may be a useful component in the ICU's emergency preparedness toolkit.

CHAPTER 4: CONCLUSIONS

Informative tools supporting the healthcare system, particularly the ICU, are essential during a public health emergency. Developing and adapting these tools comes with many challenges. In this work, we have shown considerable promise in using ML and QT to support functions of the healthcare system.

We found that ML can be a useful tool for classifying COVID-19 patient types using individual patient-level data. This data is readily available to the clinician with the adoption of the EHR system. Further, we found evidence that ML models trained on the multiple data modalities display enhanced predictive power. The evidence we found in this study suggests that ML may be a helpful support tool even as the virus continues to mutate.

Additionally, we found that QT can be adapted to support modeling efforts towards healthcare system emergency preparedness. Building on the previous work modeling the ICU with QT, we created a highly flexible framework suited to simulating emergency conditions. This framework accounts for two levels of modeling perspectives, as it simulates both single and multiple facilities. Using values from literature creating various scenarios, our framework identified optimal allocations of ICU clinicians and beds when constrained by an operational budget.

REFERENCES

[1] R. Khandia, S. Singhal, T. Alqahtani, M. A. Kamal, N. A. El-Shall, F. Nainu, P. A. Desingu, and K. Dhama, "Emergence of SARS-CoV-2 Omicron (B. 1.1. 529) variant, salient features, high global health concerns and strategies to counter it amid ongoing COVID-19 pandemic," *Environmental research*, vol. 209, p. 112816, 2022.

[2] A. Duggal and K. S. Mathews, "Impact of icu strain on outcomes," *Current opinion in critical care*, vol. 28, no. 6, pp. 667–673, 2022.

[3] A. T. Janke, H. Mei, C. Rothenberg, R. D. Becher, Z. Lin, and A. K. Venkatesh, "Analysis of hospital resource availability and covid-19 mortality across the united states," *Journal of hospital medicine*, vol. 16, no. 4, pp. 211–214, 2021.

[4] "COVID-19 epidemiological update â 16 FebruaryÂ 2024 — who.int." `https://www.who.int/publications/m/item/covid-19-epidemiological-update-16-february-2024`. [Accessed 17-02-2026].

[5] Y. Chen, L. Ouyang, F. S. Bao, Q. Li, L. Han, H. Zhang, B. Zhu, Y. Ge, P. Robinson, M. Xu, *et al.*, "A multimodality machine learning approach to differentiate severe and nonsevere covid-19: model development and validation," *Journal of medical Internet research*, vol. 23, no. 4, p. e23948, 2021.

[6] Z. Wu and J. M. McGoogan, "Characteristics of and important lessons from the coronavirus disease 2019 (covid-19) outbreak in china: summary of a report of 72 314 cases from the chinese center for disease control and prevention," *jama*, vol. 323, no. 13, pp. 1239–1242, 2020.

[7] Q. Abbas, W. Jeong, and S. W. Lee, "Explainable ai in clinical decision support systems: a meta-analysis of methods, applications, and usability challenges," in *Healthcare*, vol. 13, p. 2154, MDPI, 2025.

[8] I. D. Mienye, G. Obaido, N. Jere, E. Mienye, K. Aruleba, I. D. Emmanuel, and B. Ogbuokiri, "A survey of explainable artificial intelligence in healthcare: Concepts, applications, and challenges," *Informatics in Medicine Unlocked*, vol. 51, p. 101587, 2024.

[9] E. C. Gök and M. O. Olgun, "Smote-nc and gradient boosting imputation based random forest classifier for predicting severity level of covid-19 patients with

blood samples," *Neural Computing and Applications*, vol. 33, no. 22, pp. 15693–15707, 2021.

[10] J. Luo, L. Zhou, Y. Feng, B. Li, and S. Guo, "The selection of indicators from initial blood routine test results to improve the accuracy of early prediction of covid-19 severity," *PLoS One*, vol. 16, no. 6, p. e0253329, 2021.

[11] Y. Xiong, Y. Ma, L. Ruan, D. Li, C. Lu, L. Huang, and N. T. C. M. M. Team, "Comparing different machine learning techniques for predicting covid-19 severity," *Infectious diseases of poverty*, vol. 11, no. 1, p. 19, 2022.

[12] "Novel coronavirus pneumonia diagnosis and treatment plan (provisional 7th edition)." `http://www.nhc.gov.cn/yzygj/s7652m/202003/a31191442e29474b98bfed5579d5af95.shtml`. [Accessed 04-03-2020].

[13] J. P. Metlay, G. W. Waterer, A. C. Long, A. Anzueto, J. Brozek, K. Crothers, L. A. Cooley, N. C. Dean, M. J. Fine, S. A. Flanders, *et al.*, "Diagnosis and treatment of adults with community-acquired pneumonia. an official clinical practice guideline of the american thoracic society and infectious diseases society of america," *American journal of respiratory and critical care medicine*, vol. 200, no. 7, pp. e45–e67, 2019.

[14] N. Shakhovska, V. Yakovyna, and V. Chopyak, "A new hybrid ensemble machine-learning model for severity risk assessment and post-covid prediction system," *Math. Biosci. Eng*, vol. 19, no. 6, pp. 6102–6123, 2022.

[15] F. Cabitza, A. Campagner, D. Ferrari, C. Di Resta, D. Ceriotti, E. Sabetta, A. Colombini, E. De Vecchi, G. Banfi, M. Locatelli, *et al.*, "Development, evaluation, and validation of machine learning models for covid-19 detection based on routine blood tests," *Clinical Chemistry and Laboratory Medicine (CCLM)*, vol. 59, no. 2, pp. 421–431, 2021.

[16] T. Rymarczyk, E. Kozłowski, G. Kłosowski, and K. Niderla, "Logistic regression for machine learning in process tomography," *Sensors*, vol. 19, no. 15, p. 3400, 2019.

[17] Y. Xiong, Y. Ma, L. Ruan, D. Li, C. Lu, L. Huang, and N. T. C. M. M. Team, "Comparing different machine learning techniques for predicting covid-19 severity," *Infectious diseases of poverty*, vol. 11, no. 1, p. 19, 2022.

[18] E. Hernández-Pereira, O. Fontenla-Romero, V. Bolón-Canedo, B. Cancela-Barizo, B. Guijarro-Berdiñas, and A. Alonso-Betanzos, "Machine learning techniques to predict different levels of hospital care of covid-19," *Applied Intelligence*, vol. 52, no. 6, pp. 6413–6431, 2022.

[19] E. Jamshidi, A. Asgary, N. Tavakoli, A. Zali, S. Setareh, H. Esmaily, S. H. Jamaldini, A. Daaee, A. Babajani, M. A. Sendani Kashi, *et al.*, "Using machine learning to predict mortality for covid-19 patients on day 0 in the icu," *Frontiers in digital health*, vol. 3, p. 681608, 2022.

[20] C. Saegerman, A. Gilbert, A.-F. Donneau, M. Gangolf, A. N. Diep, C. Meex, S. Bontems, M.-P. Hayette, V. DâOrio, and A. Ghuysen, "Clinical decision support tool for diagnosis of covid-19 in hospitals," *PLoS One*, vol. 16, no. 3, p. e0247773, 2021.

[21] P. Ranganathan, C. Pramesh, and R. Aggarwal, "Common pitfalls in statistical analysis: logistic regression," *Perspectives in clinical research*, vol. 8, no. 3, pp. 148–151, 2017.

[22] M. Schonlau and R. Y. Zou, "The random forest algorithm for statistical learning," *The Stata Journal*, vol. 20, no. 1, pp. 3–29, 2020.

[23] Z. Zhang, "Introduction to machine learning: k-nearest neighbors," *Annals of translational medicine*, vol. 4, no. 11, p. 218, 2016.

[24] M. Bansal, A. Goyal, and A. Choudhary, "A comparative analysis of k-nearest neighbor, genetic, support vector machine, decision tree, and long short term memory algorithms in machine learning," *Decision analytics journal*, vol. 3, p. 100071, 2022.

[25] M. Farhadian, P. Shokouhi, and P. Torkzaban, "A decision support system based on support vector machine for diagnosis of periodontal disease," *BMC Research Notes*, vol. 13, no. 1, p. 337, 2020.

[26] G. Medic, M. K. Kließ, L. Atallah, J. Weichert, S. Panda, M. Postma, and A. El-Kerdi, "Evidence-based clinical decision support systems for the prediction and detection of three disease states in critical care: A systematic literature review," *F1000Research*, vol. 8, p. 1728, 2019.

[27] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, *et al.*, "Scikit-learn: Machine

learning in python," *the Journal of machine Learning research*, vol. 12, pp. 2825–2830, 2011.

[28] C. Aliferis and G. Simon, "Overfitting, underfitting and general model overconfidence and under-performance pitfalls and best practices in machine learning and ai," *Artificial intelligence and machine learning in health care and medical sciences: Best practices and pitfalls*, pp. 477–524, 2024.

[29] J. Huang and C. X. Ling, "Using auc and accuracy in evaluating learning algorithms," *IEEE Transactions on knowledge and Data Engineering*, vol. 17, no. 3, pp. 299–310, 2005.

[30] M. Saarela and S. Jauhiainen, "Comparison of feature importance measures as explanations for classification models," *SN Applied Sciences*, vol. 3, no. 2, p. 272, 2021.

[31] S. Learn, "Feature importances with a forest of trees," *URL https://scikit-learn.org/stable/auto_examples/ensemble/plot_forest_importances. html*, 2024.

[32] Y. Zhang, B. Nie, J. Du, J. Chen, Y. Du, H. Jin, X. Zheng, X. Chen, and Z. Miao, "Feature selection based on neighborhood rough sets and gini index," *PeerJ Computer Science*, vol. 9, p. e1711, 2023.

[33] A. Uslu and J. Stausberg, "Value of the electronic medical record for hospital care: update from the literature," *Journal of medical Internet research*, vol. 23, no. 12, p. e26323, 2021.

[34] S. Shi, M. Qin, B. Shen, Y. Cai, T. Liu, F. Yang, W. Gong, X. Liu, J. Liang, Q. Zhao, *et al.*, "Association of cardiac injury with mortality in hospitalized patients with covid-19 in wuhan, china," *JAMA cardiology*, vol. 5, no. 7, pp. 802–810, 2020.

[35] H.-H. Yu, C. Qin, M. Chen, W. Wang, and D.-S. Tian, "D-dimer level is associated with the severity of covid-19," *Thrombosis research*, vol. 195, pp. 219–225, 2020.

[36] S. Almazeedi, S. Al-Youha, M. H. Jamal, M. Al-Haddad, A. Al-Muhaini, F. Al-Ghimlas, and S. Al-Sabah, "Characteristics, risk factors and outcomes among the first consecutive 1096 patients diagnosed with covid-19 in kuwait," *EClinicalMedicine*, vol. 24, 2020.

[37] G. Suleyman, R. A. Fadel, K. M. Malette, C. Hammond, H. Abdulla, A. Entz, Z. Demertzis, Z. Hanna, A. Failla, C. Dagher, *et al.*, "Clinical characteristics and morbidity associated with coronavirus disease 2019 in a series of patients in metropolitan detroit," *JAMA network open*, vol. 3, no. 6, p. e2012270, 2020.

[38] A. S. Yadaw, Y.-c. Li, S. Bose, R. Iyengar, S. Bunyavanich, and G. Pandey, "Clinical features of covid-19 mortality: development and validation of a clinical prediction model," *The Lancet Digital Health*, vol. 2, no. 10, pp. e516–e525, 2020.

[39] T. Jiang, J. L. Gradus, T. L. Lash, and M. P. Fox, "Addressing measurement error in random forests using quantitative bias analysis," *American journal of epidemiology*, vol. 190, no. 9, pp. 1830–1840, 2021.

[40] K. B. Johnson, W.-Q. Wei, D. Weeraratne, M. E. Frisse, K. Misulis, K. Rhee, J. Zhao, and J. L. Snowdon, "Precision medicine, ai, and the future of personalized health care," *Clinical and translational science*, vol. 14, no. 1, pp. 86–93, 2021.

[41] Y. Wu, Z. Cao, J. Yang, X. Bi, W. Xiong, X. Feng, Y. Yan, Z. Zhang, and Z. Zhang, "Innovative public strategies in response to covid-19: A review of practices from china," *Health care science*, vol. 3, no. 6, pp. 383–408, 2024.

[42] "China's actions to combat the new coronavirus pneumonia outbreak- information office of the state council of the people's republic of china." `http://www.gov.cn/zhengce/2020-06/07/content_5517737.htm`. [Accessed 12-05-2023].

[43] Y. Wu, X. Feng, M. Gong, J. Han, Y. Jiao, S. Li, T. Li, C. Shen, H.-Y. Wang, X. Yu, *et al.*, "Evolution and major changes of the diagnosis and treatment protocol for covid-19 patients in china 2020–2023," *Health Care Science*, vol. 2, no. 3, pp. 135–152, 2023.

[44] S. Nembrini, I. R. König, and M. N. Wright, "The revival of the gini importance?," *Bioinformatics*, vol. 34, no. 21, pp. 3711–3718, 2018.

[45] T.-T. Nguyen, J. Z. Huang, and T. T. Nguyen, "Unbiased feature selection in learning random forests for high-dimensional data," *The Scientific World Journal*, vol. 2015, no. 1, p. 471371, 2015.

[46] S. A. J. Zaidi, A. Ghafoor, J. Kim, Z. Abbas, and S. W. Lee, "Heartensemblenet: An innovative hybrid ensemble learning approach for cardiovascular risk prediction," in *Healthcare*, vol. 13, p. 507, MDPI, 2025.

[47] N. L. Fitriyani, M. Syafrudin, N. Chamidah, M. Rifada, H. Susilo, D. Aydin, S. L. Qolbiyani, and S. W. Lee, "A novel approach utilizing bagging, histogram gradient boosting, and advanced feature selection for predicting the onset of cardiovascular diseases," *Mathematics*, vol. 13, no. 13, p. 2194, 2025.

[48] C. E. Godshall and D. B. Banach., "Pandemic preparedness.," *Infectious Disease Clinics*, vol. 35, no. 4, pp. 1077–1089, 2021.

[49] K. S. Mathews and E. F. Long, "A conceptual framework for improving critical care patient flow and bed use," *Annals of the American Thoracic Society*, vol. 12, no. 6, pp. 886–894, 2015. PMID: 25822477.

[50] M. L. McManus, M. C. Long, A. Cooper, and E. Litvak, "Queuing theory accurately models the need for critical care resources," *Anesthesiology*, vol. 100, p. 1271, 2004.

[51] G. Dobson, H.-H. Lee, and E. Pinker, "A model of icu bumping," *Operations Research*, vol. 58, no. 6, pp. 1564–1576, 2010.

[52] S. L. Zimmerman, A. R. Rutherford, A. van der Waall, M. Norena, and P. Dodek, "A queuing model for ventilator capacity management during the covid-19 pandemic," *Health Care Management Science*, vol. 26, pp. 200–216, 2023.

[53] W. Abdellatief and A. Abdelatey, "Analysis of coronavirus patients flow in hospitals: An application of queuing theory," *International Journal of Computers and Information*, vol. 9, pp. 46–59, 2022.

[54] G. H.-H. Jen, S.-Y. Chen, W.-J. Change, C.-N. Chen, A. M.-F. Yen, and R.-E. Chang, "Evaluating medical capacity for hospitalization and intensive care unit of covid-19: A queue model approach," *Journal of the Formosan Medical Association*, vol. 120, pp. S86–S94, 2021.

[55] H. D. Meares and M. P. Jones, "When a system breaks: queueing theory model of intensive care bed needs during the covid-19 pandemic," *Medical Journal of Australia*, vol. 212, no. 10, pp. 470–471, 2020.

[56] V. Sundarapandian, *Probability, statistics and queuing theory.* PHI Learning Pvt. Ltd., 2009.

[57] X. Lai, M. Wang, C. Qin, L. Tan, L. Ran, D. Chen, H. Zhang, K. Shang, C. Xia, S. Wang, *et al.*, "Coronavirus disease 2019 (covid-2019) infection among health care workers and implications for prevention measures in a tertiary hospital in wuhan, china," *JAMA network open*, vol. 3, no. 5, p. e209666, 2020.

[58] P. Griffiths, C. Saville, J. Ball, D. Culliford, J. Jones, F. Lambert, P. Meredith, B. Rubbo, L. Turner, and C. Dallâora, "Nursing team composition and mortality following acute hospital admission," *JAMA Network Open*, vol. 7, no. 8, p. e2428769, 2024.

[59] H. E. Shimizu, D. T. Couto, E. Merchán-Hamann, and A. B. Branco, "Occupational health hazards in icu nursing staff," *Nursing research and practice*, vol. 2010, no. 1, p. 849169, 2010.

[60] L. J. Allen, *An introduction to stochastic processes with applications to biology*. CRC press, 2010.

[61] S. Bocquet, *Queueing theory with reneging*. DSTO, 2005.

[62] A. Combes, C.-E. Luyt, J.-L. Trouillet, J. Chastre, and C. Gibert, "Adverse effect on a referral intensive care unit's performance of accepting patients transferred from another intensive care unit," *Critical care medicine*, vol. 33, no. 4, pp. 705–710, 2005.

[63] I. Kiselev, I. Akberdin, and F. Kolpakov, "Delay-differential seir modeling for improved modelling of infection dynamics," *Scientific Reports*, vol. 13, no. 1, p. 13439, 2023.

[64] Y. Hou and H. Bidkhori, "Multi-feature seir model for epidemic analysis and vaccine prioritization," *Plos one*, vol. 19, no. 3, p. e0298932, 2024.

[65] D. S. Burke, "Origins of the problematic e in seir epidemic models," *Infectious Disease Modelling*, vol. 9, no. 3, pp. 673–679, 2024.

[66] S. Mwalili, M. Kimathi, V. Ojiambo, D. Gathungu, and R. Mbogo, "Seir model for covid-19 dynamics incorporating the environment and social distancing," *BMC research notes*, vol. 13, no. 1, p. 352, 2020.

[67] W. Zhang, R. Xie, X. Dong, J. Li, P. Peng, and E. D. S. Gonzalez, "Seir-fmi: A coronavirus disease epidemiological model based on intra-city movement, inter-city movement and medical resource investment," *Computers in Biology and Medicine*, vol. 149, p. 106046, 2022.

[68] F. Brauer, C. Castillo-Chavez, and Z. Feng, "Models for influenza," in *Mathematical Models in Epidemiology*, pp. 311–350, Springer, 2019.

[69] J. Reis and J. Shaman, "Retrospective parameter estimation and forecast of respiratory syncytial virus in the united states," *PLoS computational biology*, vol. 12, no. 10, p. e1005133, 2016.

[70] R. Nikbakht, M. R. Baneshi, and A. Bahrampour, "Estimation of the basic reproduction number and vaccination coverage of influenza in the united states (2017-18)," *Journal of Research in Health Sciences*, vol. 18, no. 4, p. e00427, 2018.

[71] Y. Liu, A. A. Gayle, A. Wilder-Smith, and J. Rocklöv, "The reproductive number of covid-19 is higher compared to sars coronavirus," *Journal of travel medicine*, 2020.

[72] S. Zhao, Q. Lin, J. Ran, S. S. Musa, G. Yang, W. Wang, Y. Lou, D. Gao, L. Yang, D. He, *et al.*, "Preliminary estimation of the basic reproduction number of novel coronavirus (2019-ncov) in china, from 2019 to 2020: A data-driven analysis in the early phase of the outbreak," *International journal of infectious diseases*, vol. 92, pp. 214–217, 2020.

[73] Y. Liu and J. Rocklöv, "The effective reproductive number of the omicron variant of sars-cov-2 is several times relative to delta," *Journal of travel medicine*, vol. 29, no. 3, p. taac037, 2022.

[74] P. Girardi and C. Gaetan, "An seir model with time-varying coefficients for analyzing the sars-cov-2 epidemic," *Risk Analysis*, vol. 43, no. 1, pp. 144–155, 2023.

[75] X. Xu, Y. Wu, A. G. Kummer, Y. Zhao, Z. Hu, Y. Wang, H. Liu, M. Ajelli, and H. Yu, "Assessing changes in incubation period, serial interval, and generation time of sars-cov-2 variants of concern: a systematic review and meta-analysis," *BMC medicine*, vol. 21, no. 1, p. 374, 2023.

[76] K. Zeng, S. Santhya, A. Soong, N. Malhotra, D. Pushparajah, K. C. Thoon, B. Yeo, Z. J. M. Ho, and M. C. I. Cheng, "Serial intervals and incubation periods of sars-cov-2 omicron and delta variants, singapore," *Emerging infectious diseases*, vol. 29, no. 4, p. 814, 2023.

[77] K. Sankar, N. Modi, A. Polyak, M. P. Directo, L. R. Johnson, N. Kho, S. K. Isonaka, I. Pedraza, P. Chen, and M. E. Modes, "Comparison of clinical characteristics and outcomes of critically ill adults with sars-cov-2 infection during delta and omicron variant predominance periods: a single-hospital retrospective cohort study," *BMJ Open Respiratory Research*, vol. 10, no. 1, 2023.

[78] S. Setianto and D. Hidayat, "Modeling the time-dependent transmission rate using gaussian pulses for analyzing the covid-19 outbreaks in the world," *Scientific Reports*, vol. 13, no. 1, p. 4466, 2023.

[79] B. Ridenhour, J. M. Kowalik, and D. K. Shay, "Unraveling r 0: Considerations for public health applications," *American journal of public health*, vol. 104, no. 2, pp. e32–e41, 2014.

[80] P. H. Thrall, A. Biere, and M. K. Uyenoyama, "Frequency-dependent disease transmission and the dynamics of the silene-ustilago host-pathogen system," *The American Naturalist*, vol. 145, no. 1, pp. 43–62, 1995.

[81] T. Abaziou, C. Delmas, F. Vardon Bounes, F. Bignon, L. Crognier, T. Seguin, B. Riu-Poulenc, S. Ruiz, A. Rouget, P. Cougot, *et al.*, "Outcome of critically ill patients with influenza infection: a retrospective study," *Infectious Diseases: Research and Treatment*, vol. 13, p. 1178633720904081, 2020.

[82] V. K. Moitra, C. Guerra, W. T. Linde-Zwirble, and H. Wunsch, "Relationship between icu length of stay and long-term mortality for elderly icu survivors," *Critical care medicine*, vol. 44, no. 4, pp. 655–662, 2016.

[83] H. B. Gershengorn, G. L. Anesi, V. X. Liu, D. K. Costa, E. M. Dress, A. L. Dzierba, R. Fowler, A. A. Kramer, D. Lizano, D. C. Scales, *et al.*, "Association of intensive care unit patient-to-clinician ratios with mortality across two us health systems," *Annals of the American Thoracic Society*, vol. 22, no. 9, pp. 1372–1381, 2025.

[84] "U.s hospital revenue and expense trends.." `https://www.definitivehc.com/blog/revenue-trends-at-u.s.-hospitals`, 2024. [Accessed 20-02-2026].

[85] "Occupational employment and wage statistics." `https://www.bls.gov/oes/2022/may/oes_nat.htm`, 2022. [Accessed 03-03-2026].

[86] "How much do hospitals spend on beds? is there a way to save money?." `https://www.svimedical.com/2020/03/13/how-much-hospitals-spend-on-beds/`

#:~:text=Many%20%20manufacturers%20%20customize%20hospital%20bed,adheres%20to%20their%20financial%20%20wherewithal., 2020. [Accessed 03-03-2026].

[87] M. L. McManus, M. C. Long, A. Cooper, and E. Litvak, "Queuing theory accurately models the need for critical care resources," *Anesthesiology*, vol. 100, no. 5, pp. 1271–1276, 2004.

[88] M. A. Begen, F. F. Rodrigues, T. Rice, and G. S. Zaric, "A forecasting tool for a hospital to plan inbound transfers of covid-19 patients from other regions," *BMC Public Health*, vol. 24, no. 1, p. 505, 2024.

[89] S. Murthy, A. Leligdowicz, and N. K. Adhikari, "Intensive care unit capacity in low-income countries: a systematic review," *PloS one*, vol. 10, no. 1, p. e0116949, 2015.

[90] I. M. Longini Jr, A. Nizam, S. Xu, K. Ungchusak, W. Hanshaoworakul, D. A. Cummings, and M. E. Halloran, "Containing pandemic influenza at the source," *Science*, vol. 309, no. 5737, pp. 1083–1087, 2005.

[91] H. Chen, Y. Xia, Q. Qin, J. Cheng, and C. Xiong, "Spatial equity and factors associated with intensive care unit bed allocation in china," *Archives of Public Health*, vol. 82, no. 1, p. 169, 2024.

[92] A. OâHalloran, "Influenza-associated hospitalizations during a high severity seasonâinfluenza hospitalization surveillance network, united states, 2024–25 influenza season," *MMWR. Morbidity and Mortality Weekly Report*, vol. 74, 2025.

[93] E. Mathieu, H. Ritchie, L. Rodes-Guirao, C. Appel, D. Gavrilov, C. Giattino, J. Hasell, B. Macdonald, S. Dattani, D. Beltekian, E. Ortiz-Ospina, and M. Roser, "Coronavirus (covid-19) hospitalizations," *Our World in Data*, 2020. https://archive.ourworldindata.org/20251125-173858/covid-hospitalizations.html.

[94] S. Baharoon, W. Alyafi, H. Tamim, H. Al-Jahdali, E. Alsafi, A. Al-Sayyari, and Q. Ahmed, "Continuous mandatory onsite consultant intensivists in the icu: impacts on patient outcomes," *Journal of Patient Safety*, vol. 12, no. 2, pp. 108–113, 2016.

[95] U. D. of Health Human Services, "COVID-19 reported patient impact and hospital capacity by facility," 5 2024.

APPENDIX A: MULTI-FEATURE CLINICAL DECISION SUPPORT

A.1 Tables

Table A.1: Standard Deviations of Area under the Receiver Operating Characteristic Curve (AUC)

	ML method	LR		RF		kNN		SVM	
	Training	Original	Omicron	Original	Omicron	Original	Omicron	Original	Omicron
Modality	Testing								
Biochemical	Original	0.044	0.007	0.057	0.012	0.057	0.010	0.045	0.007
	Omicron	0.026	0.022	0.029	0.027	0.034	0.023	0.033	0.022
Clinical	Original	0.061	0.006	0.054	0.012	0.060	0.011	0.043	0.007
	Omicron	0.057	0.023	0.024	0.030	0.015	0.024	0.031	0.024
Fusion	Original	0.051	0.007	0.055	0.014	0.056	0.008	0.044	0.009
	Omicron	0.058	0.021	0.043	0.019	0.027	0.023	0.044	0.020

Table A.2: Imported packages for ML pipeline

Package name	Purpose in code
csv	Allows data retrieval from .csv files
pandas	Transform data into Pandas dataframes
numpy	Mathematical computations
matplotlib	generating graphical displays
sklearn.preprocessing	Utilization of data scalars
sklearn.model_selection	Utilization of train_test_spit and GridSearchCV
sklearn.linear_model	Initializes Logistic regression
sklearn.ensemble	Initializes RandomForestClassifier
sklearn.neighbors	Initializes KNeighborsClassifier
sklearn.svm	Initializes SVM

Table A.3: Frequent hyperparameters selected during ML model training by variant and data modality

Model Hyperparameters			
Logistic Regression trained on original variant			
Hyperparameters tested	Biochemical preferred	Clinical preferred	Fused preferred
penalty	None	l1/l2	l1
Logistic Regression trained on Omicron variant			
Hyperparameters tested	Biochemical preferred	Clinical preferred	Fused preferred
penalty	None	None/l1	None/l1
Random Forest trained on original variant			
Hyperparameters tested	Biochemical preferred	Clinical preferred	Fused preferred
N_estimators (mean)	8	7	7
Max_depth (mean)	5	5	5
criterion	log_loss	entropy/log_loss	gini/log_loss
Random Forest trained on Omicron variant			
Hyperparameters tested	Biochemical preferred	Clinical preferred	Fused preferred
n_estimators (mean)	9	8	9
max_depth (mean)	8	6	6
criterion	log_loss	gini	log_loss/gini
k-Nearest Neighbors trained on original variant			
Hyperparameters tested	Biochemical preferred	Clinical preferred	Fused preferred
n_neighbors	8	9	10
weights	No preference	No preference	Uniform
metric	cosine	cosine	cosine
k-Nearest Neighbors trained on Omicron variant			
Hyperparameters tested	Biochemical preferred	Clinical preferred	Fused preferred
n_neighbors	8	9	10
weights	distance	distance	distance
metric	cosine	cosine	cosine
Support Vector Machine trained on original variant			
Hyperparameters tested	Biochemical preferred	Clinical preferred	Fused preferred
kernel	linear	sigmoid	sigmoid
Support Vector Machine trained on Omicron variant			
Hyperparameters tested	Biochemical preferred	Clinical preferred	Fused preferred
kernel	linear	sigmoid	linear/rbf

A.2 ROC plots

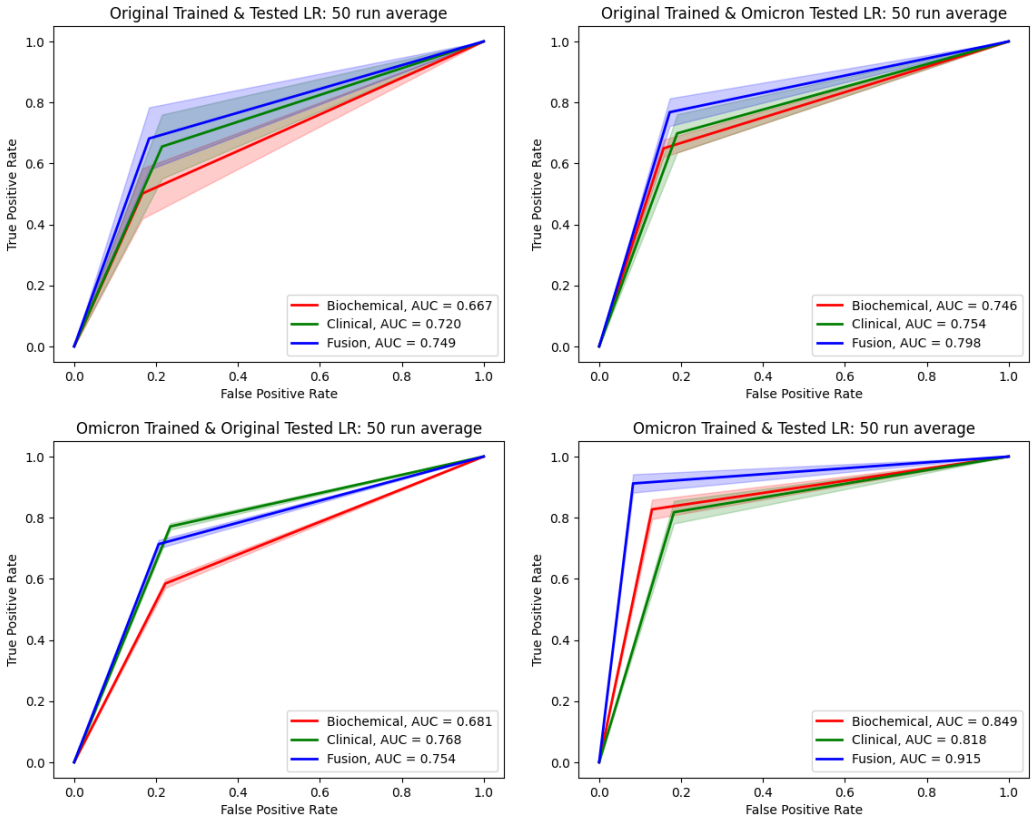

Figure A.1: ROC plots from Logistic Regression (LR) models

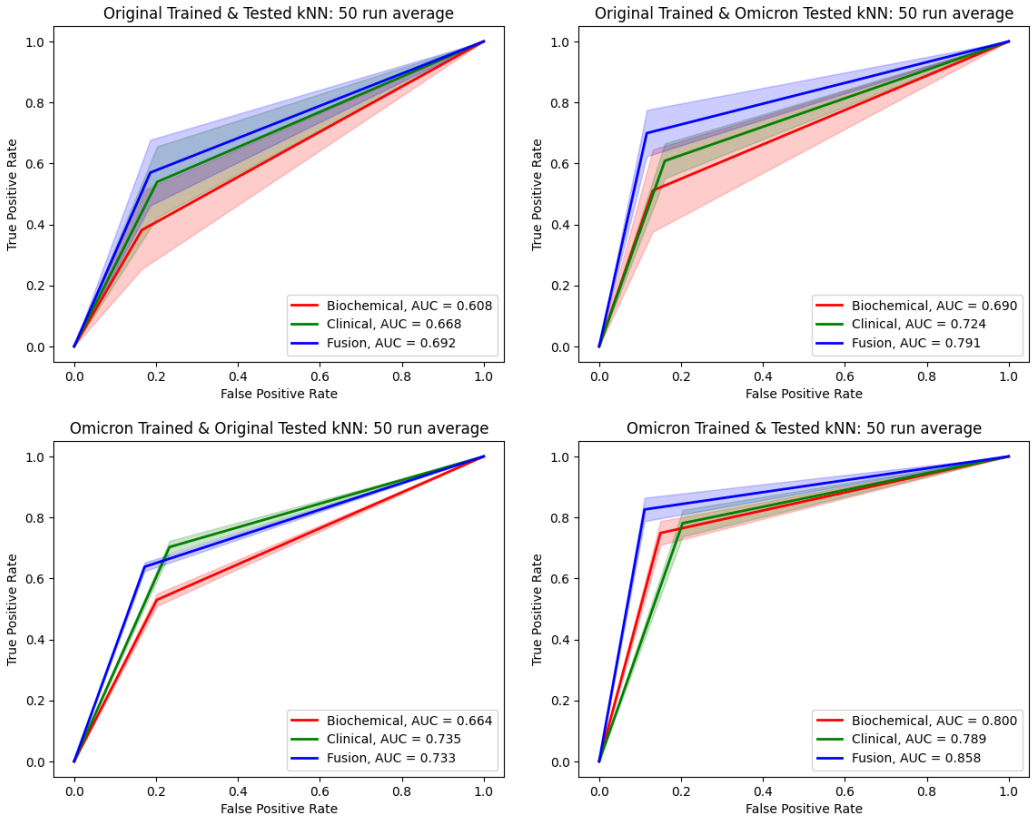

Figure A.2: ROC plots from k-Nearest Neighbors (kNN) models

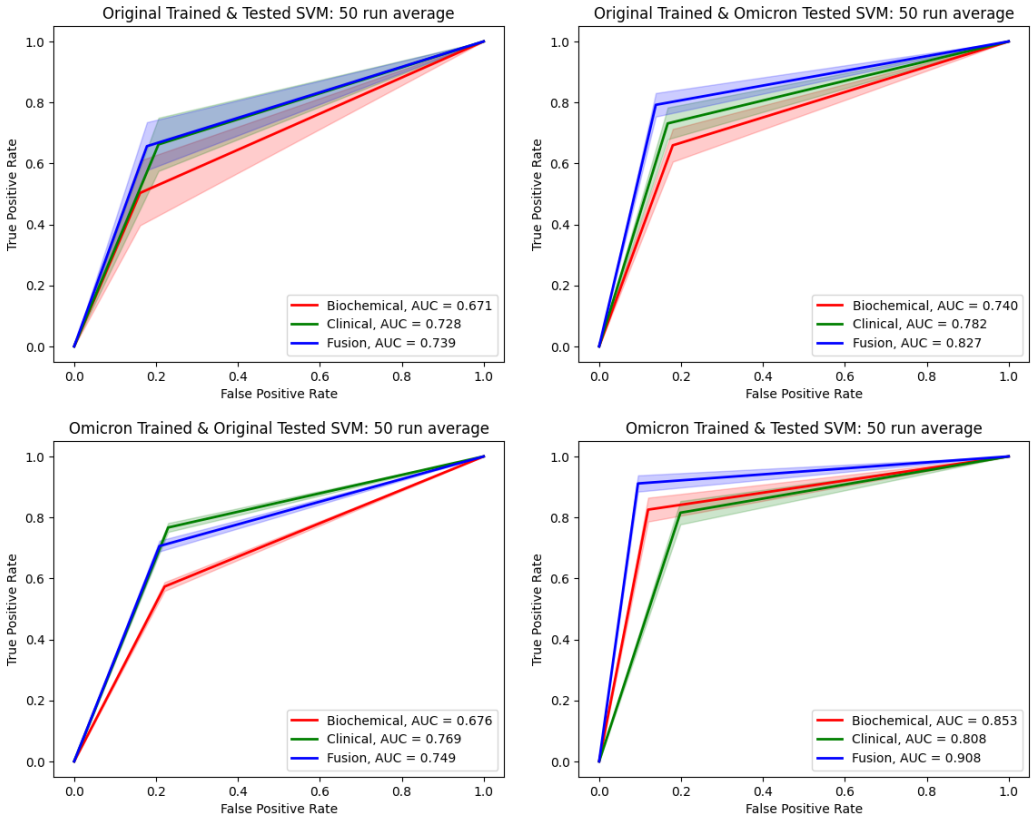

Figure A.3: ROC plots from Support Vector Machine (SVM) models

A.3 Feature Importances

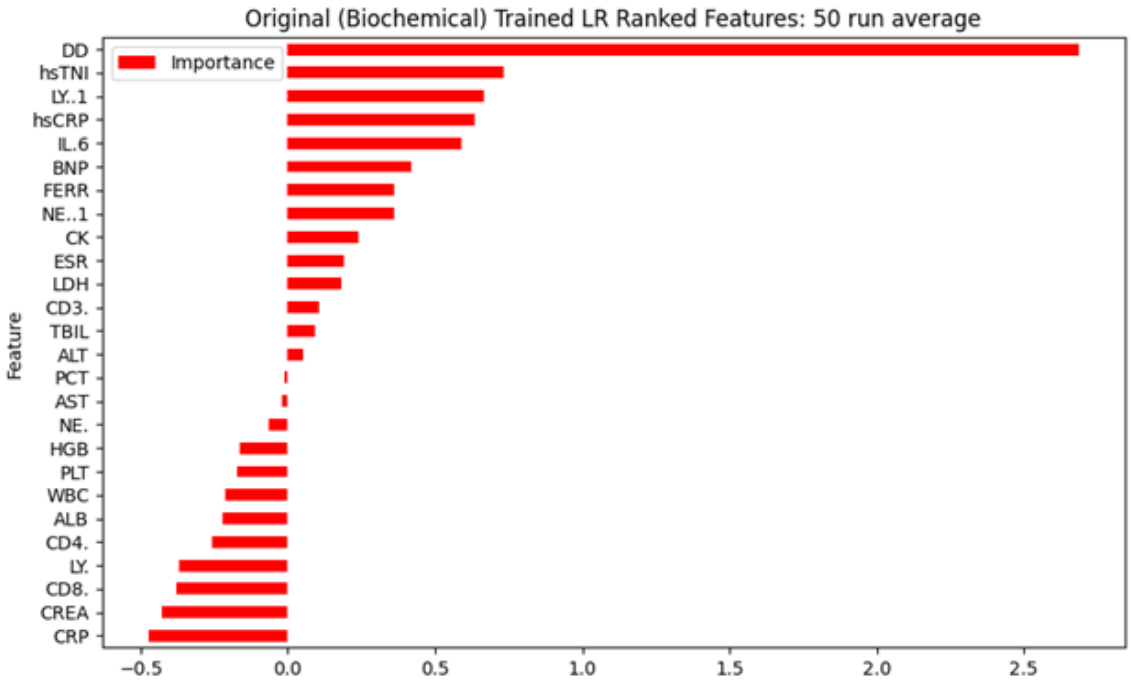

Figure A.4: Feature importances identified by LR model trained on original biochemical data only

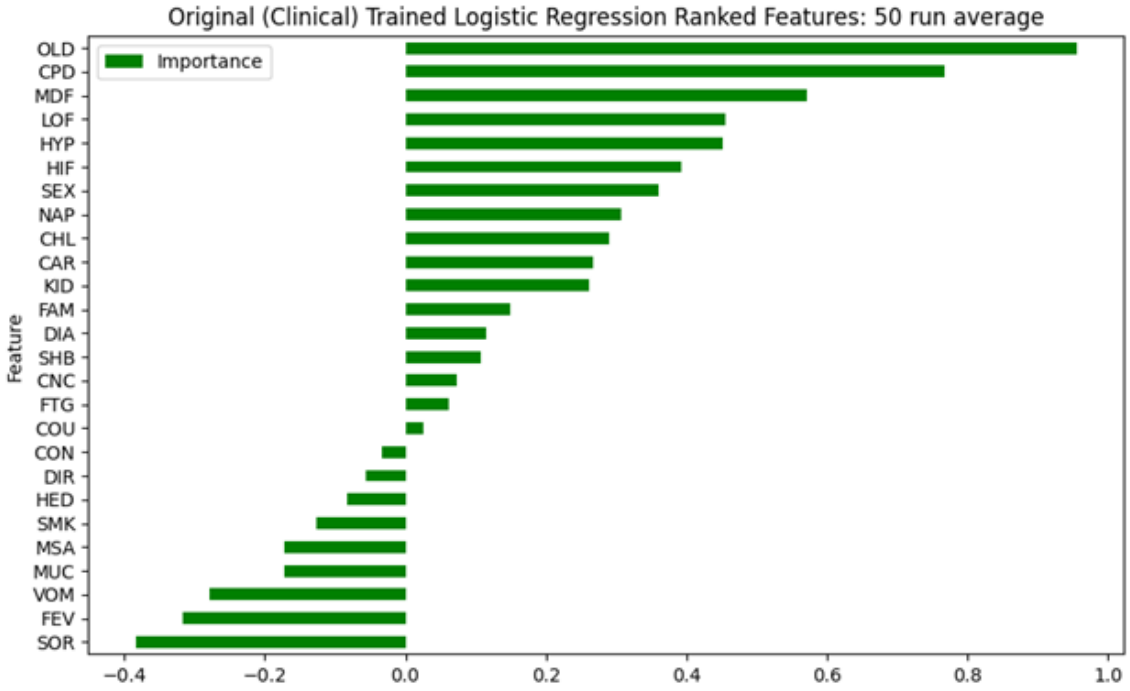

Figure A.5: Feature importances identified by LR model trained on original clinical data only

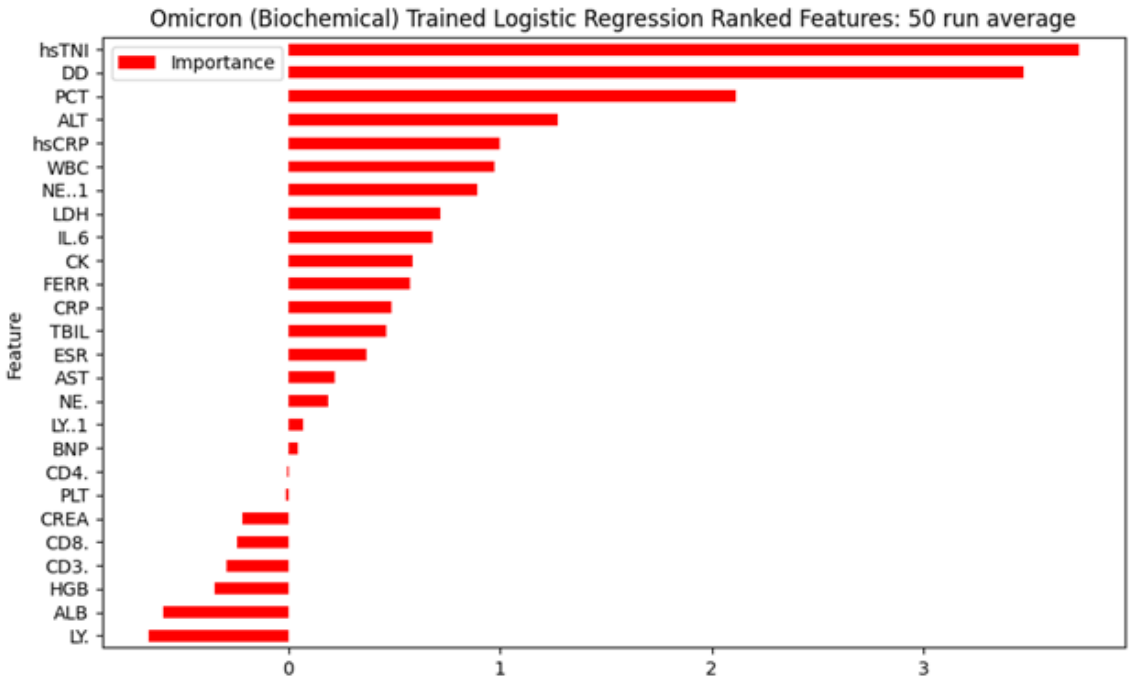

Figure A.6: Feature importances identified by LR model trained on Omicron biochemical data only

Figure A.7: Feature importances identified by LR model trained on Omicron clinical data only

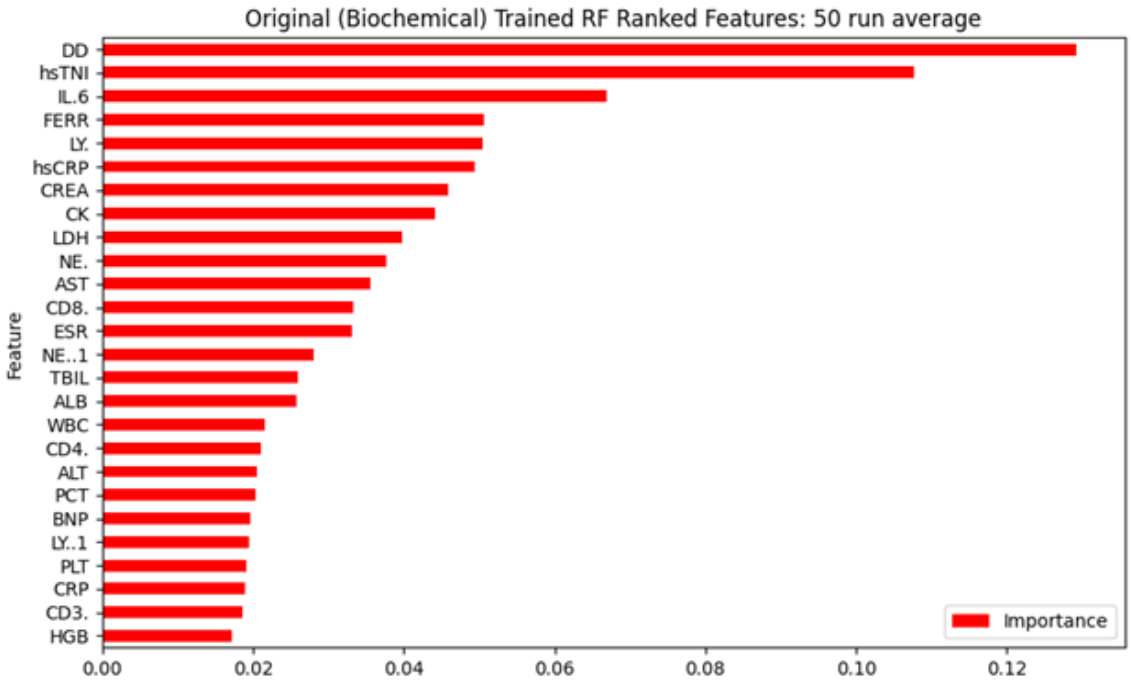

Figure A.8: Feature importances identified by RF model trained on original biochemical data only

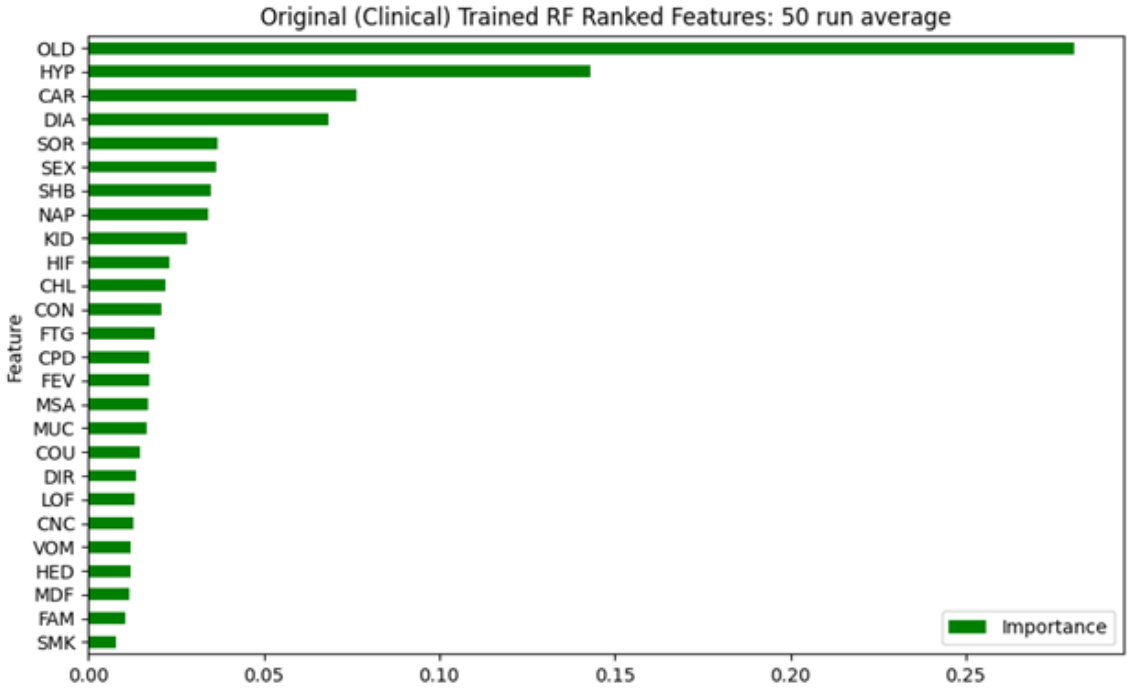

Figure A.9: Feature importances identified by RF model trained on original clinical data only

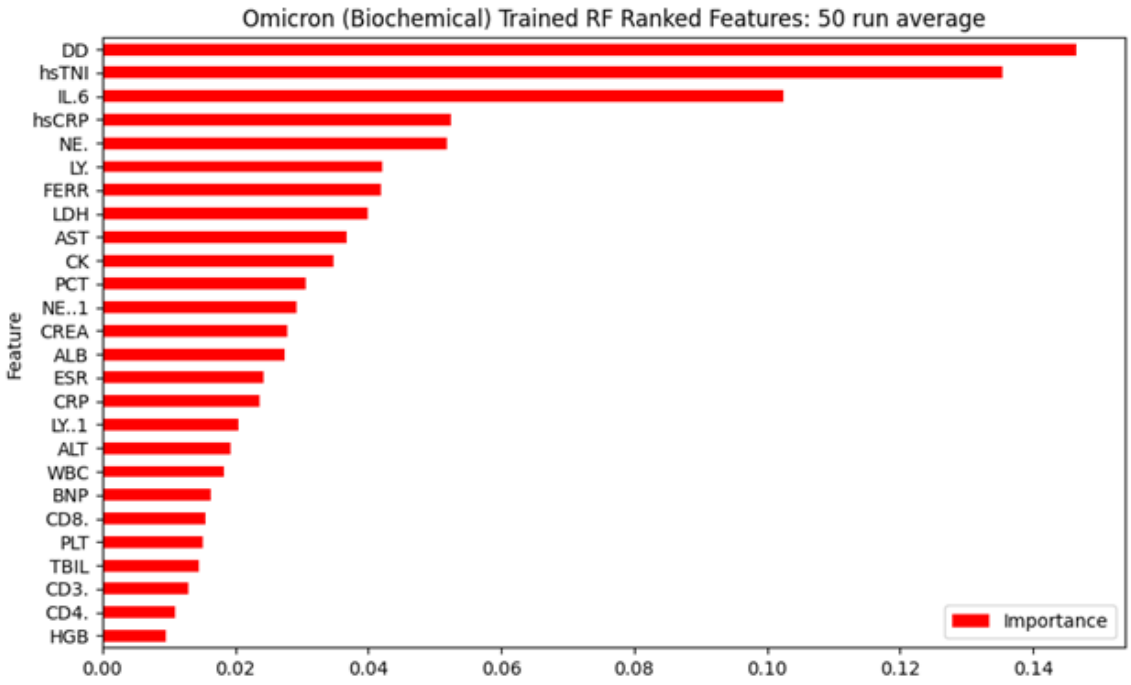

Figure A.10: Feature importances identified by RF model trained on Omicron biochemical data only

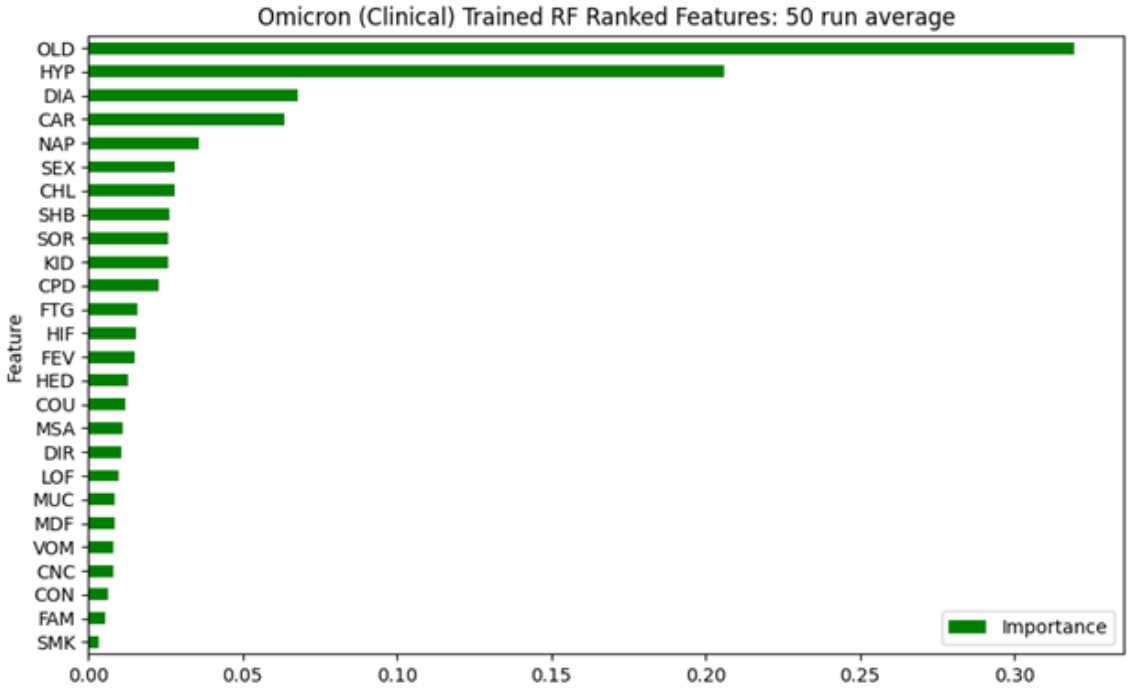

Figure A.11: Feature importances identified by RF model trained on Omicron clinical data only

APPENDIX B: MODELING EMERGENCY PREPAREDNESS IN THE HEALTHCARE SYSTEM

B.1 GitHub Link

Our code can be found at: https://github.com/hnwestpage/ICU-simulator