Predicting Medicine Inpatient Discharges at Massachusetts General Hospital

by

Keren S. Starobinski

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

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2020

Author . . . .
Department of Electrical Engineering and Computer Science
February 5, 2020

Certified by . . . .
Retsef Levi
J. Spencer Standish (1945) Professor of Operations Management
MIT Sloan School of Management
Thesis Supervisor

Accepted by . . . .
Katrina LaCurts
Chair, Master of Engineering Thesis Committee

Predicting Medicine Inpatient Discharges at Massachusetts General Hospital

by

Keren S. Starobinski

Submitted to the Department of Electrical Engineering and Computer Science
on February 5, 2020, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

At Massachusetts General Hospital, inpatients often experience significant non-clinical delays in patient care, and frequently wait in the Emergency Department or in inpatient-floor hallways before receiving bed assignments. Such delays result in overcrowding in the Emergency Department, heightened dissatisfaction among patients, and an increase in overall patient length-of-stay.

Delays in bed assignments primarily occur because of the discrepancy between the timing of admissions, which generally occur throughout the day, and the timing of discharges, which typically occur in the afternoon. Furthermore, although bed managers know about scheduled admissions in advance, there is no standardized protocol that allows bed managers at the Admitting Department to identify which patients are ready to leave the hospital.

In this project, we develop a discharge prediction tool that identifies medicine and neurology inpatient discharges that will occur within the next 24 hours. The goal is to use this tool to enable a more proactive bed-management process at MGH, provide the hospital staff with a methodical way to identify discharges, and ameliorate overcrowding challenges in the Emergency Department.

The model was trained using the data of 60,993 inpatients who were hospitalized sometime between May 2016 and September 2018. The prediction algorithm achieved a 0.830 mean AUC-ROC (SD 0.002), 47.6% precision (24 hours), 67.4% precision (48 hours), and 43.8% recall using a decision threshold of 0.31. For inpatients who were on cardiology floors within the Department of Medicine, the model achieved 58.3% precision (24 hours), 74.3% precision (48 hours), and 63.5% recall using 0.31 as the decision threshold. Since the model used data that is accessible in most hospital information systems, it can be applied to other hospitals as well.

Thesis Supervisor: Retsef Levi
Title: J. Spencer Standish (1945) Professor of Operations Management
MIT Sloan School of Management

Acknowledgments

There are many individuals who I would like to thank for their time and effort in supporting this thesis.

First, I would like to express my deepest gratitude to Professor Retsef Levi for his invaluable guidance, support, and insight throughout this project. His enthusiastic commitment to improving healthcare operations is truly inspiring.

I am also very grateful to the remarkable team at Massachusetts General Hospital. Thank you to Dr. Peter Dunn and Bethany Daily for their leadership and constructive feedback throughout this project. I would also like to thank Dr. Kyan Safavi for his clinical insights and devotion, and Cecilia Zentano for her mentorship and advice. A special thank you to Dr. Marjory Bravard for improving my understanding of the Department of Medicine, and to Christina Stone R.N. and Patricia Gaudette R.N. for answering my questions about case management hospital operations.

In addition, I am thankful to Martin Copenhaver for his technical guidance, and to Taghi Khaniyev, a postdoctoral fellow in the Operations Management group at the MIT Sloan School of Management, who was an instrumental collaborator throughout this project.

Finally, I would like to express my wholehearted gratitude to my family. Thank you to my parents, Shira and David Starobinski, and to my sister, Noga, for their unconditional love and encouragement throughout this journey. None of this would have been possible without them.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

1 Introduction **15**

1.1 Background . . . 15

1.1.1 Massachusetts General Hospital . . . 15

1.1.2 The MGH-MIT Collaboration . . . 16

1.1.3 Departments of Medicine and Neurology . . . 16

1.1.4 Medicine Inpatients at MGH . . . 17

1.1.5 Capacity Challenges at MGH . . . 20

1.2 Project Overview . . . 22

1.2.1 Problem Statement . . . 22

1.2.2 Goals . . . 22

1.2.3 Technical Approach . . . 23

1.3 Model Performance . . . 25

1.4 Thesis Outline . . . 27

2 Literature Review **29**

2.1 Discharge Prediction . . . 29

2.2 Length-of-Stay Prediction . . . 31

2.3 Natural Language Processing in the Clinical Domain . . . 33

3 Current State **37**

3.1 Medicine Inpatient Floors and Staff . . . 37

3.1.1 Department of Medicine Inpatient Floors . . . 37

3.1.2 Department of Neurology Inpatient Floors . . . 42

3.2 Admission Sources 43
3.3 Discharge Process 45

4 Methodology **49**
4.1 Data Sources Overview 49
4.1.1 Static Patient Data 50
4.1.2 Dynamic Patient Data 54
4.2 Features Generation 65
4.2.1 Static Data 65
4.2.2 Dynamic Data 66
4.3 Predictive Classification Model 72
4.4 Evaluating the Model 74

5 Model Performance **81**
5.1 ROC Curve 81
5.2 Precision and Recall 83

6 Recommendations and Conclusions **89**
6.1 Recommendations 89
6.2 Areas for Improvement and Opportunities for Further Study 90

A Overview of the Prediction Tool's Framework **93**

B Data Sources Overview **95**

C Sentiment Analysis Classifier Overview **109**

List of Figures

1-1 A histogram of the LOS of medicine inpatients in 2017. The mean, standard deviation, and median are, respectively, 7.38 days, 7.62 days, and 5 days. 18

1-2 The percentage of medicine inpatients cared by each patient service from May 2016 until September 2018. 19

4-1 An overview of the methodology used to develop the model. 49

4-2 The distribution of medicine inpatient ages in 2017. The mean age is 63 years old and the standard deviation is 17.5 years. 52

4-3 A screenshot of 4Next's Summary page embedded within EPIC (June, 2019). On this page, a case manager is able to view which facilities or home-with-services accepted or denied a patient, as well as which facilities or home-with-services are pending placements. 55

4-4 A screenshot of the communication messages between two facilities and a case manager on the 4Next application (June, 2019). 56

4-5 A screenshot of 4Next's Search & Refer tab embedded within EPIC (June, 2019). This tab allows a case manager to search, send a referral, or provide clinical data or comments to a facility or home-with-service. 56

4-6 The number of medicine inpatient discharges per day-of-the-week in 2018. Notice that on weekends (Saturday and Sunday), the number of discharges is noticeably lower than the number of discharges on weekdays (Monday through Friday). Also notice that the number of discharges on Monday is lower compared to the number of discharges on other weekdays. This is because multiple federal holidays occur on Monday, including Memorial Day, Labor Day, Columbus Day, and Martin Luther King Jr. Day. 57

4-7 In the left table, each patient service is assigned to a number that corresponds to a location in the one-hot encoded vector. In this case, medicine corresponds to the first entry, cardiology corresponds to the second entry, oncology corresponds to the third entry, and neurology corresponds to the last entry. The resulting one-hot encoded vector for each patient service is represented in the right table. For instance, the one-hot encoded vector for oncology would be $[0, 0, 1, 0]$. 65

4-8 An example of the trigger-resolution features for 4Next data. The left table corresponds to sample 4Next data. The right table corresponds to the trigger-resolution features that were generated using the sample 4Next data. 67

4-9 A visualization of the discharge prediction model's architecture. A fully-connected dense layer takes static data as input (1), while the LSTM takes dynamic data as input (2). The output of the LSTM goes through a fully-connected layer (3), and is then concatenated with the transformed static data (4). This is then used to output a prediction score (5). 73

4-10 A visualization of the input shapes of the discharge prediction model. n represents the number of feature vectors, d represents the number of dynamic features, and s represents the number of static features. Notice that the following hyperparameters were used (i) 16 hidden neurons for the fully-connected layer that takes static data as input; (ii) 32 hidden neurons for the fully-connected layer following the LSTM; (iii) 50 hidden neurons for the LSTM; and (iv) 7 previous days of inpatient data. 75

4-11 This figure shows a simplified example of the dynamic data that is fed into the LSTM before it has been preprocessed. Each row corresponds to a day the patient was at the hospital. Since 7 was found to be the optimal number of previous days of inpatient data that should be used as input into the LSTM model, the table above has 7 rows. In this particular example, the patient has been in the hospital overnight for 4 days, so the first three rows are assigned to zero. 75

4-12 An example of a plot that illustrates the precision and recall values for various thresholds. A threshold of 0.28 is chosen since it minimizes the absolute difference between precision and recall. 78

5-1 The ROC Curve for the 24-hour medicine discharge prediction model. The ROC was computed using 10-fold cross validation and inpatient data from May 2016 until September 2018. 82

5-2 The ROC Curve for the 24-hour medicine-minus discharge prediction model. The ROC was computed using 10-fold cross validation and inpatient data from May 2016 until September 2018. 84

5-3 The recall and precision values for various thresholds in the medicine model. The threshold that minimizes the absolute difference between recall and precision is 0.31. 85

5-4 The recall and precision values for various thresholds in the medicine-minus model. The threshold that minimizes the absolute difference between recall and precision is 0.29. 86

A-1 This diagram illustrates the steps necessary to train the discharge prediction model. 94

C-1 A visualization of the dispo-phrase sentiment-analysis classifier's architecture. A dispo phrase is fed into an Embedding Layer (1). The resulting output is inputted into an LSTM that analyzes the sequence of words in the dispo phrase (2). Afterwards, the fully-connected layer (3) takes the output from the LSTM and transforms it into a score that indicates whether the dispo phrase contains positive sentiment (i.e., the patient will leave within 24 hours) or not. 111

List of Tables

1.1 The mean AUC-ROC, precision (24 hours, 48 hours) and recall metrics for two versions of the medicine discharge prediction model: one model (Medicine) uses additional patient data that was not in Zanger's model, a trigger-resolution framework, and an LSTM-based architecture; while the second model (Medicine (-)) does not. 27

1.2 The mean AUC-ROC, precision (24 hours, 48 hours) and recall metrics for Zanger's enhanced surgical discharge prediction model and the medicine discharge prediction model. 27

3.1 The medicine inpatient floors at MGH along with their corresponding floors types and patient services. 38

3.2 The regionalized teams in the Albert Service. Note that there are 36 total beds on Ellison 16, but only 10 of them were designated under the General Internal Medicine division in December 2019. The remaining 26 beds were assigned to the Hematology-Oncology division. 40

3.3 The firms within the Bigelow Service. The Bigelow B team covers Bigelow 9 patients that are not covered by the Gray team. 41

3.4 The number of private rooms and beds on cardiology inpatient floors. 41

3.5 The number of private rooms and beds on oncology floors. Note that 26 of the 36 total beds on Ellison 16 were designated under the Oncology division in December 2019. The remaining 10 beds were assigned to the General Internal Medicine division. 42

3.6 The number of private rooms and beds on the neurology inpatient floors. 43

4.1 The distribution of medicine inpatients' admission sources during 2017. 'Other' includes court, law enforcement, and hospice. 51
4.2 The distribution of the top ten diagnosis groups for medicine inpatients in 2017. 53
4.3 The distribution of LDAs medicine patients' had during 2017. 'Other' includes negative pressure wound therapy, biliary drain, penrose drain, epidural catheter, and external fecal management system. 59
4.4 The distribution of microbiology tests medicine inpatients' had during 2017. 'Other' includes mycobacterial, throat, genital, and Group B Strep Infection (GBS). 62

5.1 The precision (24 hours, 48 hours) and recall metrics for the medicine model, the medicine-minus model (marked with '-'), and Zanger's enhanced surgical model. 83

Chapter 1

Introduction

This thesis aims to address the bed capacity challenges at Massachusetts General Hospital (MGH) by developing a discharge prediction tool for medicine and neurology patients. The ultimate goal is to use this tool to identify likely discharges within 24 hours in order to improve MGH's discharge management process. Since the challenges at MGH are typical to all academic medical centers in the United States, the work and results described in this thesis have high promise to generalize more broadly.

This chapter provides an overview of the project, highlights the key results, and presents the outline for the remainder of this document.

1.1 Background

The following subsections provide background information about MGH, the MGH-MIT Collaboration, and the Departments of Medicine and Neurology. The inpatient population of this study and the capacity challenges at MGH are also described.

1.1.1 Massachusetts General Hospital

Massachusetts General Hospital was established in 1811 to provide medical care to the physically and mentally ill [1]. It is consistently ranked as one of the top hospitals in the nation by *U.S. News & World Report* and serves as a teaching hospital for

Harvard Medical School. MGH is also the largest hospital-based research enterprise in the United States. With an annual budget of over $900 million, the hospital conducts research in more than 30 departments, centers, and institutes [2].

Each year, MGH admits approximately 50,000 inpatients, handles nearly 1.5 million outpatient visits, and carries out over 42,000 operations [3]. In addition, MGH is one of the founding members of Partners HealthCare, a nonprofit integrated healthcare delivery system that consists of community hospitals, home care and long-term care services, and other health care entities [4]. More than 26,000 employees work at MGH, making the 1000-plus bed medical center the largest nongovernment employer in Boston, Massachusetts [5].

1.1.2 The MGH-MIT Collaboration

The MGH-MIT Collaboration is a research partnership between MGH and the Massachusetts Institute of Technology (MIT) Sloan School of Management. The collaboration's objective is to improve the effectiveness and efficiency at MGH by applying state-of-the-art methods and operations research techniques to the hospital's most critical operational challenges. The members of the collaboration consist of administrative leaders and clinicians from MGH, and faculty, post-doctorate fellows and students from MIT.

The collaboration's previous initiatives to address MGH's inpatient capacity challenges include evaluating a 'Just-In-Time' bed assignment policy in the neuroscience units [6] [7], analyzing interventions to reduce patient wait-times in the Department of Medicine [8], and developing a discharge prediction tool for the surgical inpatient population [9] [10]. This thesis builds upon the methods developed in the surgical discharge prediction tool and extends them in order to create a new discharge prediction model for medicine and neurology inpatients.

1.1.3 Departments of Medicine and Neurology

The Department of Medicine at MGH provides medical care to over 800,000 patients per year. It is the largest department by inpatient volume, and is comprised of 9 research units, 20 primary-care outpatient locations, 10 clinical divisions[1], and over 10 physical inpatient floors. The 10 clinical divisions provide consultative services, while the General Internal Medicine, Cardiology, and Hematology-Oncology divisions additionally provide services for inpatient care. These, respectively, correspond to the medicine, cardiology, and oncology patient services.

The Department of Neurology at MGH provides consultative services, diagnostic testing, and treatment for a broad range of neurological conditions [11]. It is part of MGH Neuroscience, and handles more than 50,000 inpatient admissions and office visits each year [12]. Patients within the Department of Neurology (which corresponds to the neurology patient service) are treated on one dedicated ICU unit and on two inpatient floors. Both the medicine and neuroscience units consistently operate near capacity and experience overcrowding challenges.

Chapter 3 provides further information about the Department of Medicine's and the Department of Neurology's inpatient floors and hospital staff.

1.1.4 Medicine Inpatients at MGH

The focus of this thesis is to predict the discharge likelihood for inpatients who (i) are receiving care on MGH's medicine and neuroscience inpatient floors; and (ii) are being treated by the medicine, cardiology, or oncology patient service within the Department of Medicine (DOM), or the neurology patient service within the Department of Neurology in MGH Neuroscience. For simplicity, throughout the thesis these inpatients are referred to as *medicine inpatients*.

Admissions

Medicine inpatients are admitted to MGH for a wide-variety of reasons. They usually have diverse, multi-system medical conditions that are difficult to diagnose and distinct, psychosocial complexities that require highly specialized care. As a result, the

[1]The 10 clinical divisions are Cardiology, Endocrinology, Gastroenterology, General Internal Medicine, Hematology-Oncology, Infectious Diseases, Nephrology, Palliative Care & Geriatric Medicine, Pulmonary & Critical Care Medicine, and Rheumatology, Allergy & Immunology.

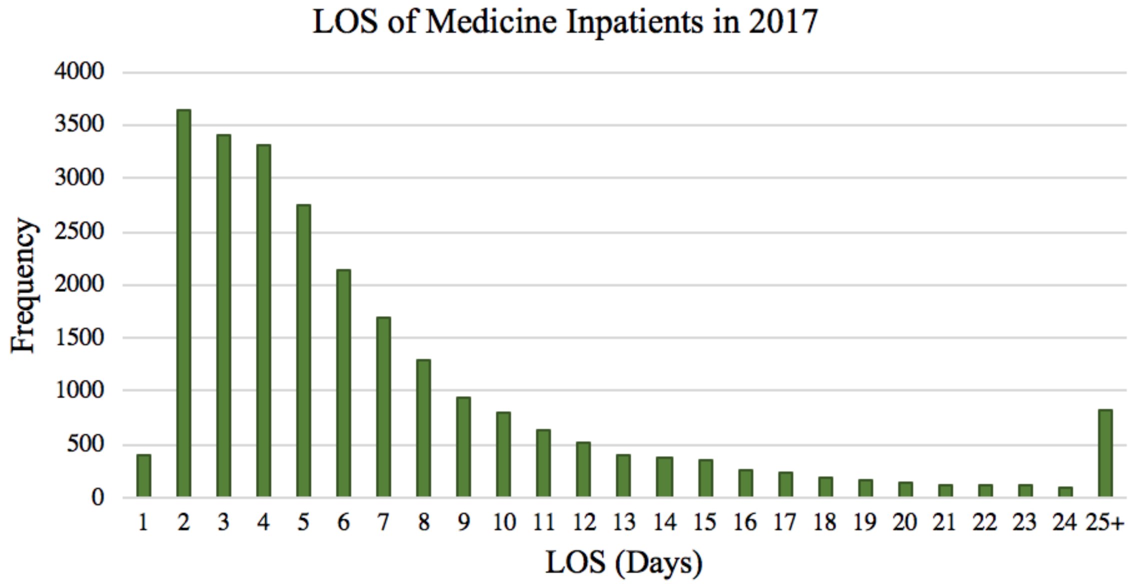

Figure 1-1: A histogram of the LOS of medicine inpatients in 2017. The mean, standard deviation, and median are, respectively, 7.38 days, 7.62 days, and 5 days.

length-of-stay (LOS) of medicine inpatients can range from one day to over a year (Figure 1-1). By contrast, the treatment plans and LOS of surgical inpatients are generally more predictable than those of medicine inpatients. This is because most surgical admissions are typically scheduled in advance and patients' diagnoses and comorbidities are usually known prior to admission [8].

The Admitting Department is responsible for working with the hospital floor care-teams to admit, transfer, and discharge surgical and medicine patients on a case-by-case basis. Medicine patients arrive to inpatient floors from various sources, such as the Emergency Department (ED), intensive care unit (ICU), home, or from other inpatient units. Since the demand for medicine beds frequently exceeds the allocated capacity, medicine inpatients are not only treated on medicine floors, but also on surgical or mixed[2] floors. While on these floors[3], medicine inpatients are cared by the

[2] Mixed floors care for medicine and surgical inpatients.

[3] The floors at MGH that care for medicine inpatients are: Blake 6, Bigelow 9, Bigelow 11, Bigelow 14, Ellison 6, Ellison 7, Ellison 8, Ellison 10, Ellison 11, Ellison 12, Ellison 14, Ellison 16, Ellison 19, Phillips 20, Lunder 7, Lunder 8, Lunder 9, Lunder 10, Phillips 21, Phillips 22, White 6, White 7, White 8, White 9, White 10, and White 11.

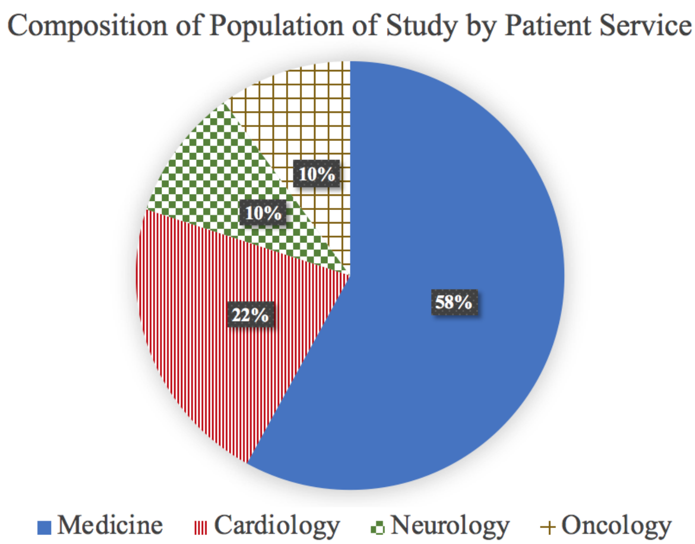

Figure 1-2: The percentage of medicine inpatients cared by each patient service from May 2016 until September 2018.

medicine, cardiology, oncology, or neurology patient service (Figure 1-2).

Medicine patients are especially challenging to assign to inpatient floors since they often have behavioral issues or infection precautions that require them to either be (i) in a room alone; or (ii) in a room with a patient that has as similar ailment. These patients must wait for a bed to become available in a private room, a whole semi-private room[4], or a semi-private room with a patient who has a similar infection. In addition, medicine patients within the DOM's General Internal Medicine division who are expected to require more physician resources must wait for a bed to become available on a *regionalized* floor. Regionalized floors are settings in which physicians are in close proximity to their patients (see Section 3.1.1 for more details).

Discharge Process

Discharge planning is initiated upon admission and fine-tuned throughout a patient's stay. This process involves the collaboration of the patient, the patient's family, physicians, nurses, and case managers, and in some cases, also social workers, physical therapists, and other supporting functions. Compared to surgical patients, medicine

[4]A semi-private room contains two beds that are separated by a curtain.

patients are usually more challenging to discharge since many have psychosocial complexities or behavioral issues, suffer from substance abuse, or do not have health insurance. Approximately 50% of medicine inpatients are discharged home, while the remainder receive home-service arrangements or are discharged to a care facility[5].

1.1.5 Capacity Challenges at MGH

Over the last ten years, the Emergency Department and medicine floors at MGH have experienced strains in capacity. Visits to the Emergency Department have increased by 18% and inpatient bed occupancy has regularly been between 95 and 100 percent [13]. Other capacity challenges at MGH include the surge in the number of *ED Boarders*[6], as well as the increase in *Code Help*[7] and *Capacity Disaster*[8] instances in the ED [13]. For instance, in November FY '20 approximately 77% of admitted patients boarded in the ED, and the average number of boarders was 7% higher than in November FY '19. Since most ED admissions are to medicine floors, ED congestion is highly affected by the bed availability on medicine units.

Patient wait-times for bed assignments are a result of both intraday delays and interday (or multi-day) delays. Intraday delays primarily occur because of temporary capacity constraints at peak hours and the discrepancy between the timing of admissions, which generally occur throughout the day, and the timing of discharges, which typically occur in the afternoon. Interday delays are typically the result of long-term, capacity constraints in certain parts of the hospital. These persistent capacity issues can delay patient transfers between different parts of the hospital for several days (e.g., patients transferring from the ICU to an inpatient floor). Interday delays can affect intraday delays and vice-versa.

Multiple efforts have been initiated to ease patient congestion and to improve ca-

[5]Source: Clinical-4Next table located in Partners Enterprise Data Warehouse (EDW).

[6]An ED Boarder is a patient who has been admitted to the hospital, but has waited in the ED for two or more hours to transfer to an inpatient bed.

[7]Code Help is a Massachusetts mandated policy that requires all admitted patients to move out of the ED within 30 minutes when the ED's maximum occupancy is reached or exceeded.

[8]Capacity Disaster is defined as the state in which Code Help has been activated for two or more hours and ED Boarders are present.

pacity management at MGH, including identifying inappropriate admissions, reducing patient readmissions, and increasing the number of early discharges on inpatient floors. Specifically, in 2011, MGH attempted to prevent overcrowding in the ED with a $500 million expansion [14]. Five years later, MGH formed the ED Boarder Service with the goal of providing inpatient care for ED Boarders waiting to get admitted to inpatient medicine floors. The Boarder Service is staffed by medicine physicians and nurse practitioners who initiate patients' diagnostic and treatment plans while admitted patients are still in the ED.

More recently, in early 2019, MGH announced a $1 billion project to build a new, 12-story, two-tower facility for inpatient and outpatient care. The aim of this project is to accommodate the growing volume of patients, replace outdated beds, and modernize the medical practices at the hospital [15]. In the spring of 2020, a Capacity Coordination Center will also be formed to more efficiently manage capacity challenges at the hospital. The center will consist of a multi-disciplinary team of physicians, nurse supervisors, Admitting Department leadership, and case managers. One of the center's initiatives is to create a hospital-wide capacity dashboard that will include a list of pre-identified inpatients who are likely to get discharged. This feature will use the medicine inpatient prediction tool described in this thesis.

Bed capacity challenges have become a serious concern not only for MGH, but also for hospitals across the United States. In the last decade, the ratio of the number of occupied beds to available bed capacity has increased across US hospitals [16], and from 2000 to 2014 the number of emergency room visits in the US climbed by over 40 percent [13]. The median emergency room wait time before admission has also risen for patients in the US to about five-and-a-half hours in 2017 [17]. Due to the increasing number of patients from abroad and the aging baby boomer generation, capacity challenges are expected to become worse [13].

1.2 Project Overview

1.2.1 Problem Statement

Throughout the past decade, MGH has faced severe bed capacity problems. Medicine inpatients often experience significant non-clinical delays in patient care, and frequently wait in the Emergency Department or in inpatient-floor hallways before receiving bed assignments [13] [18] [19]. Such delays result in overcrowding in the ED, heightened dissatisfaction among patients, and an increase in overall patient LOS [13].

Delays in bed assignments primarily occur because of the discrepancy between the timing of admissions, which generally occur throughout the day, and the timing of discharges, which typically occur in the afternoon. Furthermore, although bed managers know about scheduled admissions in advance, there is no standardized protocol that allows bed managers at the Admitting Department to identify which patients are ready to leave the hospital. Currently, bed managers become aware of prospective discharges by receiving self-reports from medicine floors. This type of communication is manual and subjective, and prevents bed managers and the hospital staff from proactively monitoring the discharge process.

Even though this project focuses on Massachusetts General Hospital, other hospitals in the United States have similar bed-management practices and capacity constraints.

1.2.2 Goals

The goal of this thesis is to develop a discharge prediction tool that relies on electronic medical data to identify medicine inpatients that are likely to be discharged within the next 24 hours. The aim is to use this tool to enable a more proactive bed-management process at MGH, provide the hospital staff and the Capacity Coordination Center with a methodical way to identify medicine inpatient discharges and their barriers to discharge, and ameliorate overcrowding challenges such as Code Help and Code

Disaster in the ED.

1.2.3 Technical Approach

In order to develop the discharge prediction model, inpatient data recorded in MGH's electronic medical record system, EPIC, was extracted. EPIC is used by clinicians to monitor an inpatient's clinical progress. It includes inpatients' progress notes, past medical histories, administered medications, immunizations, diagnoses, medical imaging outcomes, demographics, and laboratory results. The newly developed model uses both *static* patient data and *dynamic* patient data from EPIC. Static patient data is information that remains relatively constant during a patient's hospitalization, such as age or admission-source location, while dynamic data is information that is subject to change throughout a patient's stay, such as lab results or administered medications.

Patient data relevant to the medicine population was extracted from Zanger's surgical model [9] [10]. Additionally, the following patient data was extracted:

- Clinical patient data relevant to the medicine population. This includes
  - Patient diagnosis at admission
  - Microbiology orders (e.g., stool, sputum)
  - Additional laboratory results (e.g., ALP, AST)
  - Additional administered medications (e.g., nebulizers, phenobarbital)
  - Isolation status of a patient (i.e., whether a patient is in contact, droplet, or airborne isolation)
  - Echocardiogram imaging orders
  - Catheterization laboratory and electrophysiology (EP) procedure orders

  Why these were added: Medicine inpatients are admitted to the hospital for a wider variety of reasons than surgical inpatients. In order to monitor the clinical progression of all medicine inpatients, additional clinical data was included.

- Non-clinical patient data such as case-managers' electronic communication on 4Next[9] (e.g., whether a referral was sent to a facility or whether a facility accepted a patient).

  Why this was added: In contrast to surgical inpatients, medicine inpatients often have psychosocial complexities, suffer from substance abuse, or do not have health insurance. Thus, even if a medicine inpatient is clinically ready for discharge, the inpatient may remain in the hospital (e.g., a facility may refuse to accept a non-insured patient). In order to capture such situations in the model, non-clinical patient data was added.

- Physician, nursing, physical therapy, social worker, and case manager notes data.

  Why these were added: The treatment plans of medicine inpatients are generally more variable than those of surgical patients. Since physician, nursing, physical therapy, social worker and case-manager notes record medical treatment plans or discharge-related events that occur throughout a patient's hospitalization, data from these notes was extracted as well.

In order to decide which features to incorporate into the model, an iterative process was used:

- Run the prediction tool.
- Identify the medicine inpatients whose discharge status' were incorrectly predicted by the model.
- Search for data that was not captured for these inpatients.
- Extract this data from the relevant database and incorporate it into the model.

Raw patient data was processed using a *trigger-resolution* framework that is based on Zanger's framework of *clinical milestones* and *barriers-to-discharge* [9] [10]. In

[9] 4Next is a web-based application developed by Partners HealthCare that is used by case managers to contact various post-acute care facilities and home-with-services.

Zanger's model [9] [10], clinical milestones are events that correlate with a patient's progression towards discharge or recovery (e.g., stool occurrence), whereas barriers-to-discharge are events that may delay a patient's discharge from the hospital (e.g., abnormal vital signs).

The medicine model applies a trigger-resolution framework which uses clinical milestones and barriers-to-discharge. A clinical milestone corresponds to a binary (i.e., 0 or 1) feature that indicates whether the milestone has been completed. A barrier-to-discharge corresponds to two binary variables: (1) a trigger feature that marks whether the barrier-to-discharge has occurred; and (2) a resolution feature that signals whether the barrier has been resolved. To extract features from physician and case-managers' notes, the model leverages natural-language processing techniques.

Forming the feature vectors using the trigger-resolution framework outlined above allowed us to train the discharge prediction model for the medicine inpatients using a multi-input, fully-connected layer and long short-term memory (LSTM) neural-network architecture. Unlike the multi-layered perceptron model used in Zanger's surgical model, an LSTM is capable of making predictions based on time series data and of learning long-term dependencies. This is especially crucial for the medicine inpatient population, which is clinically diverse and whose inpatients have an LOS that can range from one day to over a year.

The discharge prediction model was trained on 60,993 medicine inpatient discharges from two-and-a-half years worth of data.

1.3 Model Performance

The discharge prediction model was trained and tested using the data of MGH medicine inpatients who were hospitalized sometime between May 2016 and September 2018. This population included 60,993 inpatients who stayed in the hospital overnight for a total of 421,343 days. The model's performance was assessed using 10-fold cross-validation. One metric that was used to assess the performance was the AUC-ROC. The AUC-ROC represents how well the model is able to distinguish between patients who will get discharged within 24 hours and patients who will remain at the hospital for at least 24 more hours. An AUC-ROC of 1 represents a perfect classification, whereas an AUC-ROC of 0.5 represents an uninformative classifier. The model achieved a mean AUC-ROC of 0.830 (SD 0.002).

Although AUC-ROC is commonly used to evaluate the performance of classification algorithms, it does not necessarily indicate how well the model will perform operationally. As a result, the model was also evaluated with respect to recall and precision metrics. Recall represents how often the model is able to identify patient discharges, while precision represents how often the model is correct when it predicts that a patient will get discharged within 24 hours.

Additionally, the model was evaluated with respect to the precision of discharge within 48 hours. This represents how often a patient is predicted to be discharged within 24 hours by the model, but is actually discharged within 48 hours. This metric is motivated by the fact that there exist situations when a patient may be clinically ready for discharge, but ends up staying at the hospital for non-clinical reasons, such as a physician decision or a post-hospital care constraint[10]. The 48-hour precision metric takes these situations into account.

The discharge prediction model achieved 47.6% precision (24 hours), 67.4% precision (48 hours), and 43.8% recall using a decision threshold of 0.31[11]. For medicine inpatients who were on cardiology floors, the model achieved 58.3% precision (24 hours), 74.3% precision (48 hours), and 63.5% recall using 0.31 as the decision threshold.

In order to assess the effect of (1) adding patient data that was not in Zanger's model; (2) generating features using the trigger-resolution framework; and (3) training the discharge prediction model using an LSTM-based architecture, a *modified* medicine model that does not use the previously mentioned techniques was assessed. The AUC-ROC, precision (24 hours and 48 hours), and recall metrics of this modified model were lower than the non-modified medicine model's metrics (Table 1.1).

[10]Note that there may also be cases when the model predicts that a patient will get discharged within 24 hours, but the patient is actually discharged within 48 hours for *clinical* reasons.

[11]The discharge prediction model outputs a score between 0 and 1. A decision threshold of 0.31 means that scores above 0.31 indicate that a patient will get discharged within 24 hours and scores below 0.31 indicate the a patient will remain in the hospital.

Model	Mean AUC-ROC (SD)	Precision (24 hours)	Precision (48 hours)	Recall
Medicine	0.830 (0.002)	47.6%	67.4%	43.8%
Medicine (-)	0.783 (0.004)	40.2%	59.3%	38.3%

Table 1.1: The mean AUC-ROC, precision (24 hours, 48 hours) and recall metrics for two versions of the medicine discharge prediction model: one model (Medicine) uses additional patient data that was not in Zanger's model, a trigger-resolution framework, and an LSTM-based architecture; while the second model (Medicine (-)) does not.

Model	Mean AUC-ROC (SD)	Precision (24 hours)	Precision (48 hours)	Recall
Surgery	0.912 (0.006)	62.1%	83.1%	62.5%
Medicine	0.830 (0.002)	47.6%	67.4%	43.8%

Table 1.2: The mean AUC-ROC, precision (24 hours, 48 hours) and recall metrics for Zanger's enhanced surgical discharge prediction model and the medicine discharge prediction model.

Additionally, the model was compared to an enhanced version of Zanger's surgical model, which used the additional patient data, trigger-resolution framework, and LSTM-based architecture described previously. The medicine model's AUC-ROC is around 0.08 lower than that of Zanger's enhanced model, while the precision and recall are approximately 15% and 20% less than the enhanced surgical model's metics (Table 1.2). Notice, however, that the medicine model's 24-hour precision and recall results for the cardiology service are similar to those of the enhanced surgical model.

1.4 Thesis Outline

The remainder of this thesis is organized as follows. Chapter 2 presents a review of relevant studies in the existing literature. Chapter 3 provides an overview of the medicine inpatient floors (Section 3.1), admission sources (Section 3.2), and discharge planning process (Section 3.3). Chapter 4 details the methodology used in developing the discharge prediction algorithm by reviewing the data sources used for prediction tool (Section 4.1), describing the raw data transformation (Section 4.2), and outlining the training (Section 4.3) and evaluation (Section 4.4) processes for the predictive classification algorithm. Chapter 5 evaluates the model's AUC-ROC (Section 5.1), and precision and recall metrics (Section 5.2). Chapter 6 presents recommendations for using this tool to improve the bed-management process (Section 6.1) and provides suggestions for further study (Section 6.2).

Chapter 2

Literature Review

This chapter reviews the existing literature related to hospital discharge prediction models (Section 2.1), length-of-stay predictions (Section 2.2), and natural language processing in the clinical domain (Section 2.3).

2.1 Discharge Prediction

There are several publications that use machine learning methods to predict the likelihood of a patient's discharge. For example, Levin et al. (2012) developed a real-time model that predicts the probability of an ICU patient getting discharged within a 6-hour interval, up to 72 hours in advance [20]. In order to train the model, Levin et al. used the data of 2,178 pediatric intensive-care unit patients during a 16-month time period. The model was developed using logistic regression and achieved a 46% recall rate in predicting discharges within 24 hours.

Barnes et al. (2015) studied patients from an academic medical center, with the goal of developing an algorithm that predicts the probability of a patient getting discharged by 2pm or midnight of the same day [21]. In order to train the model, Barnes et al. used the data of 8,852 hospitalizations throughout a 34-month period. The model uses random forests and data features such as demographic information, diagnosis data, observation status[1], elapsed length-of-stay, and day-of-the-week. The

[1] Observation status is assigned to patients who are cared for within the unit and must be moni-

prediction algorithm achieved 60.0% sensitivity and 66.0% specificity in predicting discharges by 2pm, and 66.1% sensitivity and 68.3% specificity in predicting discharges by midnight.

Zanger (2018) developed a discharge prediction model to identify surgical inpatient discharges at MGH [9] [10]. This study covered 15,553 surgical inpatients admitted to MGH between May 2016 and August 2017. In developing the surgical model, Zanger collaborated with clinicians at MGH to derive clinical milestones, events that correlate with a patient's progression towards discharge or recovery, and barriers-to-discharge, events that may delay a patient's discharge from the hospital. Zanger used these clinical insights to develop a multi-layered perceptron classifier that correctly predicted 90% of surgical patients discharged within 24 hours on non-holiday weekdays.

This thesis uses a trigger-resolution framework that is based on Zanger's framework of clinical milestones and barriers-to-discharge in order to process raw patient data. In addition, since medicine patients are admitted to MGH for a wider-variety of reasons than surgical patients, and usually have psychosocial complexities or behavioral issues, we also use natural-language processing techniques to extract features from physician and case-managers' notes. These unstructured, free-text notes contain meaningful information about the discharge process of patients.

In contrast to previous work, this thesis uses additional data sources such as clinical notes, isolation status data, and data from a software that manages the communication between case managers and post-discharge facilities. This data is readily accessible in most hospital information systems, and can therefore be applied to discharge prediction models in other hospitals as well. Furthermore, this thesis makes use of a multi-input, fully-connected layer and LSTM neural-network based classifier. Unlike regression, random forest, and multi-layered perceptron models, an LSTM is capable of making predictions based on time series data and of learning long-term dependencies.

tored and evaluated before they are eligible for safe discharge.

2.2 Length-of-Stay Prediction

The goal of a length-of-stay prediction model is to predict the LOS of each patient at the time of admission. Although the objective of such a model is different than that of a discharge prediction algorithm, the techniques used to predict LOS can be adapted to predict discharge and vice versa.

Gustafson (1968) is one of the earliest examples in the literature to develop a model that estimates patients' length-of-stay [22]. The algorithm was developed using Bayesian and regression statistical methods, and was trained using symptomatic and demographic data from 8 herniotomy patients. Gustafson suggests replicating these techniques on a larger and more representative sample of patients, but does not do so in this study.

Tu et al. (1993) analyzed 713 ICU patients following cardiac surgery in order to develop a model that predicts whether a patient will stay at the ICU for longer than 2 days. The model was developed using a neural network, and was trained using patient data such as age, gender, type of surgery, urgency of surgery, and comorbidities. The algorithm achieved an AUC-ROC of 0.7094 on the training set and 0.6960 on the test set.

Walczak et al. (1998) used a neural network to predict hospital length-of-stay for pediatric trauma patients [23]. The neural network was trained using over 8,000 patients' data, such as age, gender, heart rate, temperature, hemoglobin, hematocrit, and various other blood analysis results. In contrast to Tu et al.'s model, the data used in this research is available within the first ten minutes of a patient's arrival to the hospital. Nevertheless, both Tu et al.'s model and Walczak et al.'s model indicate that neural networks have the potential to improve clinical planning and patient care.

Liu et al. (2010) used a linear regression model to predict the length-of-stay of patients [24]. Training was based on 155,474 hospitalizations between 2002 and 2005 at 17 Northern California Kaiser Permanente hospitals. Since electronically linked laboratory and comorbidity data variables greatly improved the performance of this study's model, the authors encourage a more widespread use of electronic medical records.

Azari et al. (2012) presented a multi-tiered data mining approach to predict patient length-of-stay [25]. The training set was initially preprocessed into groups using K-means clustering, and the samples close to each cluster were then used to predict the LOS of patients. This study found that preprocessing the training set using the K-means clustering was more effective than not doing so.

Morton et al. (2014) compared the results of several supervised machine learning algorithms to predict long-term versus short-term length-of-stay of in-hospital diabetic patients [26]. Patient data such as age, gender, race, and admission type were used to train the following models: multiple linear regression, support vector machine (SVM), multi-task learning, and random forests. The SVM method had the best performance for predicting short-term LOS in hospitalized diabetic patients, followed closely by random forests.

Harutyunyan (2017) used a multitask-learning technique to model the risk of mortality, extended length-of-stay (longer than 7 days), physiologic decline, and phenotype classification [27]. The model was trained using 60,000 ICU hospitalizations recorded in the Medical Information Mart for Intensive Care (MIMIC-III) database and achieved an AUC-ROC of 0.84 for the extended length-of-stay prediction algorithm.

Zebin et al. (2019) applied an autoencoder deep neural-network model in order to identify short-stays (7 days or less) and long-stays (greater than 7 days) of hospitalizations within 24 hours of admission [28]. Patient admission records, demographics, diagnosis codes, body temperatures, blood pressures, and heart rates were used from the MIMIC-III dataset in order to train the model. The prediction algorithm achieved 80.2% precision and and 77.6% recall for short-stays, and 75.2% precision and 77.9% recall for long-stays. The overall classification accuracy of the model was 77.7%.

2.3 Natural Language Processing in the Clinical Domain

Natural Language Processing (NLP) is a field of artificial intelligence that allows computers to analyze, manipulate and understand human language. Applying NLP techniques to the clinical domain is particularly challenging, since clinical texts contain a significant number of abbreviations, acronyms, and medical terms. In addition, patient records are unstructured and often have misspellings due to the limited amount of time physicians and nurses have to write these documents. Clinical notes also contain sensitive patient-specific information that raise privacy and security concerns.

One of the earliest uses of natural language processing in the clinical domain was presented by Pratt and Pacak in 1969. In order to automatically encode pathology data, Pratt and Pacak developed algorithms for morphological[2], syntactic[3] and semantic analysis[4] [29]. The long-term goal of this work was to expand and apply these techniques to all types of medical text.

Yang and Chute (1994) developed a linear least-squares fit (LLSF) method for text categorization[5] and text retrieval[6] in the clinical domain [30]. This study used a training set of manually grouped documents to learn word-category associations for text categorization (e.g., a diagnosis text summary and the corresponding ICD-9-CM code[7]), and a training set of queries and their related terms in documents for text retrieval. This paper found that using the LLSF approach is an effective way to categorize and retrieve text.

[2]Morphological parsing is the process of determining the morphemes (the smallest meaningful unit in a language) of a given word.

[3]Syntactic analysis is the process of analyzing natural language using grammatical rules.

[4]Semantic analysis refers to the process of understanding the meaning of a text, given its word-and-sentence structure.

[5]Text categorization refers to the task of labeling natural language texts using categories from a predefined set.

[6]Text retrieval is defined as obtaining specific information from non-structured, text-based data sources.

[7]International Classification of Disease, 9th Revision, Clinical Modifications (ICD-9-CM) is a system of assigning codes to diagnoses and procedures in the United States.

Pakhomov et al. (2005) compared two text-based classification algorithms to identify patients diagnosed with congestive heart failure and other related conditions [31]. Manually labeled clinical notes from the Mayo Clinic were used to train a naive bayes model and a perceptron classification model. A positive-labeled note indicated that a document contained evidence of a patient having congestive heart failure, whereas a negative-labeled note indicated that a document did not include such information. The naive bayes model achieved a recall of 95% and accuracy of 57%, and the perceptron classifier achieved a recall of 86% and an accuracy of 65%.

Savova et al. (2010) built an open-source NLP system capable of performing information extraction[8] from electronic medical records' unstructured clinical-text [32]. This system, called clinical Text Analysis and Knowledge Extraction System (cTAKES), combined rule-based and machine-learning techniques to extract meaningful information from medical documents. The system achieved a sentence-boundary detector accuracy of 94.9%, a tokenizer[9] accuracy of 94.9%, and a part-of-speech tagger accuracy of 93.6%. The authors suggest that cTAKES annotations can be used as a foundation to process clinical free-text.

McCoy (2016) used natural-language text-processing to improve stratification of risk for suicide after medical or surgical hospital discharge [33]. This study analyzed the notes of 845,417 hospital discharges that occurred between January 1, 2005 and December 31, 2013. 3,000 subjective terms (e.g. 'glad', 'gloomy') were identified in order to characterize each clinical note as having positive or negative sentiment. This study concluded that clinical notes improved the ability to estimate risk for suicide and accidental death.

In contrast to previous literature, this thesis applies natural language processing to improve the performance of an inpatient discharge prediction tool. In particular, we use regular expressions to identify discharge disposition phrases (dispo phrases) in physician, nursing, physical therapy, and social work notes. A dispo phrase be-

[8]Information Extraction is the task of automatically retrieving specific information (e.g. named entities) from unstructured sources.

[9]Tokenization is the process of breaking up a sequence of strings into words, keywords, phrases, symbols and other elements called tokens. Tokens can be individual words, phrases or even whole sentences.

gins with the character sequence “Dispo:” or “Disposition:” and signifies a patient’s progress towards discharge. For example, the dispo phrase “Dispo: home likely tomorrow” suggests that a patient will get discharged within 24 hours, whereas the dispo phrase “Dispo: pending improvement of cellulitis” implies that a patient will remain in the hospital for the near future. By applying sentiment analysis techniques, we can use dispo phrases as another feature in our model to predict whether or not a patient will be discharged within 24 hours. We discuss dispo phrases in more detail in Section 4.1.2 and in Appendix C.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

Current State

This chapter provides an overview of the medicine inpatient floors and their hospital staff (Section 3.1), the admission sources (Section 3.2), and the discharge planning process for medicine inpatients at MGH (Section 3.3). The information in this chapter was gathered through interviews with administrative and clinical hospital staff, and through the documentation of work processes in the ED, Admitting Department, and medicine inpatient units.

3.1 Medicine Inpatient Floors and Staff

This section describes the inpatient floors and hospital staff within the Department of Medicine (Section 3.1.1) and the Department of Neurology (Section 3.1.2). A comprehensive list of the medicine inpatient floors along with their corresponding floors types and patient services is provided in Table 3.1.

3.1.1 Department of Medicine Inpatient Floors

Medicine inpatients within the Department of Medicine are cared by the General Internal Medicine, Cardiology, or Hematology-Oncology divisions. A description of the inpatient floors and hospital staff within each of these three divisions is provided below.

Floor	Floor Type	Patient Service(s)
Blake 6	Mixed	Medicine
Bigelow 9	Medicine	Medicine
Bigelow 11	Medicine	Medicine
Bigelow 14	Surgical	Medicine
Ellison 6	Surgical	Medicine
Ellison 7	Surgical	Medicine
Ellison 8	Mixed	Cardiology
Ellison 10	Medicine	Cardiology
Ellison 11	Medicine	Cardiology
Ellison 12	Medicine	Medicine
Ellison 14	Surgical	Medicine
Ellison 16	Medicine	Medicine, Oncology
Ellison 19	Mixed	Medicine
Lunder 7	Mixed	Neurology
Lunder 8	Mixed	Neurology
Lunder 9	Medicine	Oncology
Lunder 10	Medicine	Oncology
Phillips 20	Medicine	Medicine
Phillips 21	Surgical	Medicine
Phillips 22	Mixed	Medicine
White 6	Surgical	Medicine
White 7	Surgical	Medicine
White 8	Medicine	Medicine
White 9	Medicine	Medicine
White 10	Medicine	Medicine
White 11	Medicine	Medicine

Table 3.1: The medicine inpatient floors at MGH along with their corresponding floors types and patient services.

General Internal Medicine

Patients within the General Internal Medicine division reside on both *regionalized* floors and *non-regionalized* floors. Regionalized floors usually consist of a team of attending physicians, residents, and nurses who care for patients local to the floor (i.e., the care team treats patients only on that floor and not anywhere else). Non-regionalized floors are typically staffed by nurses that are local to the floor and by hospitalists that go around to multiple units to care for patients. Many of these units are specialized in other clinical specialities such as surgery. The placement of General Internal Medicine patients on non-regionalized floors occurs because the demand for medicine beds frequently surpasses the allocated capacity, and other services (e.g., surgery) usually have some degree of extra capacity on their designated units.

Not all General Internal Medicine patients can be assigned to non-regionalized settings. In particular, patients with higher clinical acuity (called *Level 1* patients) must be placed on regionalized floors so that their physicians are regularly close by. The remaining patients (called *Level 2* patients) can be assigned to both regionalized and non-regionalized units.

There are two types of hospital staff structures within the General Internal Medicine division: (1) the Albright Service; and (2) the Bigelow Service. The Albright Service is staffed by hospitalists who care for patients on both regionalized and non-regionalized floors, while the Bigelow Service is staffed by attending physicians who supervise residents and only care for patients on regionalized floors. In what follows, a more detailed description of these services is provided.

- *Albright Service*

  The Albright Service consists of the Yellow, Blue, Red, Purple, and Gray teams that care for patients on regionalized floors (Table 3.2), and the Green team that treats patients on non-regionalized floors[1]. These teams[2] are staffed by

[1] The non-regionalized floors include Blake 6, Ellison 6, Ellison 7, Ellison 8, Ellison 14, Ellison 19, Phillips 22, White 6, and White 7. Most of the beds on these floors are assigned to surgical services.

[2] The Albright Service also consists of the Orange, Bronze, and Gold teams that care for patients with special social needs.

Team	Floor(s)	# of Private Rooms	Total Bed Capacity
Yellow	Phillips 20 Phillips 21	20 20	20 20
Blue	Ellison 12	2	36
Red	Ellison 16	2	36
Purple	Bigelow 14	5	27
Gray	Bigelow 9	18	18

Table 3.2: The regionalized teams in the Albert Service. Note that there are 36 total beds on Ellison 16, but only 10 of them were designated under the General Internal Medicine division in December 2019. The remaining 26 beds were assigned to the Hematology-Oncology division.

hospitalists who care for approximately ten patients at a time and who rotate among Albright teams on a regular basis.

Note that Ellison 12, Ellison 16, and Bigelow 14 contain a small number of private rooms. As a result, medicine patients who have infection precautions or behavioral issues may be challenging to assign to these floors, since such patients must wait for a bed to become available in a private room, a whole semi-private room, or a semi-private room with a patient who has a similar infection.

- *Bigelow Service*

  The Bigelow Service only cares for patients in a regionalized setting. It is staffed by medical residents and their supervising attending physicians, and occasionally also by a nurse practitioner who works with residents and is responsible for coordinating patient discharges.

  The Bigelow Service consists of 3 firms. Each firm is comprised of 3 teams that are responsible for treating patients on 2 floors. Two of the teams in each firm are dedicated to one specific floor, while the third *flex* team cares for patients on either floor in the firm (Table 3.3). A flex team was established in each firm so that it could help balance the workload when one team was experiencing a high patient volume and the other team was experiencing a low patient volume.

Firm	Floors	Team Name	Flex Name	# of Private Rooms	Bed Capacity
1	White 8 Bigelow 9	Bigelow A Bigelow B	X	6 18	26 18
2	White 9 White 10	Bigelow C Bigelow D	Y	5 4	25 20
3	Bigelow 11 White 11	Bigelow E Bigelow F	Z	6 4	26 24

Table 3.3: The firms within the Bigelow Service. The Bigelow B team covers Bigelow 9 patients that are not covered by the Gray team.

Floor	# of Private Rooms	Bed Capacity
Ellison 10	2	36
Ellison 11	2	36

Table 3.4: The number of private rooms and beds on cardiology inpatient floors.

Note that most of the floors in the Bigelow Service contain a small number of private rooms. Therefore, it is particularly challenging to place medicine patients with infection precautions or behavioral issues on Bigelow Service floors.

Cardiology

The Cardiology division within the Department of Medicine is responsible for treating patients on Ellison 9, Ellison 10, and Ellison 11, the first of which is an intensive care unit and is out of the scope for this project. Ellison 10 is staffed by General Internal Medicine and Emergency Department residents, while Ellison 11 is staffed by nurse practitioners. Both floors are supervised by a cardiologist. Electrophysiology patients are usually treated on Ellison 10, and interventional cardiology patients are typically cared for on Ellison 11. The number of private rooms and beds on Ellison 10 and Ellison 11 are listed in Table 3.4.

Note that there are a total of four private rooms in the Cardiology division. For this reason, it is especially challenging to place patients who require private rooms on cardiology floors.

Floor	# of Private Rooms	Bed Capacity
Ellison 16	2	36
Lunder 9	32	32
Lunder 10	32	32

Table 3.5: The number of private rooms and beds on oncology floors. Note that 26 of the 36 total beds on Ellison 16 were designated under the Oncology division in December 2019. The remaining 10 beds were assigned to the General Internal Medicine division.

Hematology-Oncology

The Hematology-Oncology division within the Department of Medicine treats patients on Lunder 9, Lunder 10, and part of Ellison 16. Lunder 9 and Lunder 10 are staffed by residents and their supervising physician attendings, while Ellison 16 is staffed by attendings and nurse practitioners. Bone marrow transplant, leukemia, and lymphoma patients are usually treated on these floors. The number of private rooms and beds on Lunder 9, Lunder 10, and Ellison 16 are listed in Table 3.5.

As mentioned previously, Ellison 16 contains a small number of private rooms. As a result, oncology patients who require private rooms are usually challenging to assign to Ellison 16.

3.1.2 Department of Neurology Inpatient Floors

Patients within the Department of Neurology are treated on one dedicated ICU and on two general care floors: (1) Lunder 7; and (2) Lunder 8 (Table 3.6). These floors are staffed by physicians, residents and nurses. The Department of Neurology serves patients with a wide range of diagnoses, including amyotrophic lateral sclerosis, epilepsy, and encephalitis.

Floor	# of Private Rooms	Bed Capacity
Lunder 7	32	32
Lunder 8	32	32

Table 3.6: The number of private rooms and beds on the neurology inpatient floors.

3.2 Admission Sources

At MGH, the Admitting Department is responsible for working with the hospital floor staff to admit, transfer, and discharge patients on a case-by-case basis. Bed managers within the Admitting Department work with the hospital floor staff and care providers to assign patients who need a bed to an available bed. There are usually around two to five bed managers within the Admitting Department who are responsible for the surgical units, medical units, lower volume services such as pediatrics, urology, and orthopedics, and hospital transfers[3].

Patients typically arrive to medicine inpatient floors from the Emergency Department, from locations outside the hospital (e.g., front door or outpatient facility), or from non-medicine care units within MGH. In what follows, each of these admitting channels is discussed in further detail.

Bed Assignments for ED Patients

Emergency Department physicians are responsible for diagnosing and stabilizing patients within the ED, and for determining whether an ED patient should get admitted to an inpatient medicine floor. If a physician decides to admit an ED patient to a General Internal Medicine unit, the physician must also consult with a senior medicine resident to determine the patient's *triage level* (i.e., whether the patient should be designated as a Level 1 or Level 2 patient). This information is given to the Admitting Department as part of the bed request. Once a bed is found for an ED patient, the ED physician team is required to handoff the patient to the medicine physician

[3]Hospital transfers are patients who come to MGH from another facility. Before the Admitting Department can admit a transfer patient, a senior medicine resident must approve the transfer patient's admission.

team before the patient can be transported to a medicine inpatient unit.

ED patients often wait for several hours before being transferred to medicine inpatient floors, since medicine inpatient floors frequently operate near capacity. This is one of the causes for the severe congestion in the ED, the surge in the number of ED Boarders, and the increase in the number of Code Help and Capacity Disaster instances [13]. In order to address capacity challenges in the Emergency Department and on medicine units, MGH formed the ED Boarder Service approximately four years ago. The goal of the ED Boarder Service is to provide inpatient care for ED Boarders waiting to get admitted to medicine inpatient floors. The Boarder Service is staffed by medicine physicians and nurse practitioners who initiate patients' diagnostic and treatment plans while admitted patients are still in the ED.

Nevertheless, admitting ED patients to medicine floors remains challenging. In November FY '20 approximately 77% of admitted patients boarded in the ED, and the average number of boarders was 7% higher than in November FY '19. Since approximately 70% of medicine inpatients arrive to a medicine inpatient floor from the ED, most medicine inpatients often experience significant non-clinical delays in patient care.

Bed Assignments for Patients from Outside the Hospital

Patients who arrive to medicine inpatient floors from outside the hospital generally come from the front door (e.g., home, a doctor's office) or from a different inpatient facility. In order for patients who come from the front door to get assigned to a bed on a medicine inpatient floor, a bed request must get sent to the Admitting Department. The admission process for patients who arrive from a different inpatient facility (i.e., hospital transfers) is handled by one of the bed managers in the Admitting Department. In general, patients who arrive to medicine inpatient floors from outside the hospital often must wait until there is a bed available for them. Roughly 20% of medicine inpatients are admitted to a medicine inpatient floor from outside the hospital.

Bed Assignments for Patients in a Hospital Care Unit

Patients who were transferred to medicine inpatient floors from an MGH hospital care unit typically come from the ICU, PACU or a non-medicine care floor. Once a care team on the ICU or on non-medicine care floors determines that a patient will require medicine inpatient care, it must send a bed request to the Admitting Department. Note that a patient may get transferred from a non-medicine care unit to a medicine inpatient floor if the patient's diagnosis or condition has changed. For patients in the PACU, their bed requests are usually generated when they are in the PACU to get a procedure. Patients who arrive to medicine inpatient floors from within an MGH hospital care unit typically must wait until there is a bed available for them. Close to 10% of medicine inpatients are admitted to a medicine inpatient floor from the ICU, PACU, or a non-medicine care floor.

3.3 Discharge Process

The discharge process of medicine inpatients involves the collaboration of the patient, the patient's family, physicians, nurses, and case managers, and in some cases, also social workers and physical therapists. Below is a description of the role that each of these participants has in the discharge process.

The Patient and the Patient's Family

A patient and a patient's family members can play a valuable role in the discharge process. They can provide the care team with relevant discharge information, including details about the patient's financial well-being, cultural traditions, and living circumstances. Sometimes however, a patient or the family can complicate the discharge process. For example, discharging a patient can become more difficult if family members disagree with the care team's discharge recommendation or if the patient refuses to cooperate with the care team. Since medicine patients often have psychological complexities or behavioral issues, forming agreeable discharge plans with medicine patients can be challenging.

Physicians and Nurses

Physicians and nurses are responsible for managing and maintaining care for patients. For example, they are involved in ordering diagnostic testing, determining medications, and adjusting the treatment plan. Physicians are ultimately responsible for determining whether a patient is clinically ready for discharge. However, because medicine patients often suffer from substance abuse or do not have health insurance, medicine patients may remain in the hospital for non-clinical reasons.

Case Managers

Case managers at MGH are responsible for working together with the patient, the patient's family, and care team in order to coordinate a discharge plan for the patient. When a patient initially arrives to an inpatient floor, a case manager fills out a High Risk Assessment. This is used to determine whether the patient has a high likelihood of having a complicated discharge due to clinical needs, insurance issues, social factors, or financial circumstances.

Throughout a patient's stay, a case manager evaluates whether the patient will require post-discharge care in a facility or home-with-services care. Case managers communicate with facility and home-with-services using 4Next, a Partners-owned software embedded within EPIC. 4Next allows case managers to send referrals to post-discharge services, check whether a facility or home-with-service has accepted the patient, and provide post-care services with insurance information, physician recommendations, and medication requirements of the patient.

Case managers also write progress notes that summarize the actions taken towards discharging a patient. For example, case managers may write a note describing a meeting with the patient's family or detailing a recent insurance complication. Case managers can also use various templates when writing these notes, such as a "VNA Referral" template that indicates that a VNA[4] referral was sent. Progress notes must be written at least once every seven days.

[4] VNA care provides home-with-services care such as as nursing, therapy, hospice, and palliative services to patients throughout Massachusetts.

Note that coordinating discharge plans with post-discharge services for medicine patients can be extremely complicated. This is because few facilities and home-with-services can provide the proper care for medicine patients (who often require highly specialized care) or can accept patients with behavioral issues.

Social Workers

While case managers are focused on the post-discharge logistics, social workers assess a patient's social, financial and emotional needs. They are responsible for informing the care team about these factors, and for working with the patient, the patient's family, and the care team to ensure that a patient will receive the proper care post-discharge. Social workers play a vital role in discharging medicine patients, since a considerable number of medicine patients are either financially unstable, homeless, or suffer from mental health disorders.

Physical Therapists

Physical therapists are health care professionals who treat individuals with disabilities, limited capabilities to move, or significant pain. After examining a patient, physical therapists provide a discharge recommendation for the patient. This is particularly useful for medicine patients, whose treatment and discharge plans are generally more variable than those of surgical patients.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 4

Methodology

This chapter describes the specific methods used to develop the medicine inpatient discharge prediction tool. Section 4.1 describes the static and dynamic patient data extracted from EPIC. Section 4.2 discusses the transformation of the raw data using the trigger-resolution framework. Section 4.3 details the training process for the predictive classification algorithm, and Section 4.4 outlines the metrics used to evaluate the model. An overview of the approach is presented in Figure 4-1 and in Appendix A.

4.1 Data Sources Overview

In order to determine whether a medicine inpatient is ready for discharge, inpatient data is extracted from MGH's electronic medical record system, EPIC. The extracted inpatient data from EPIC includes demographic information, diagnoses,

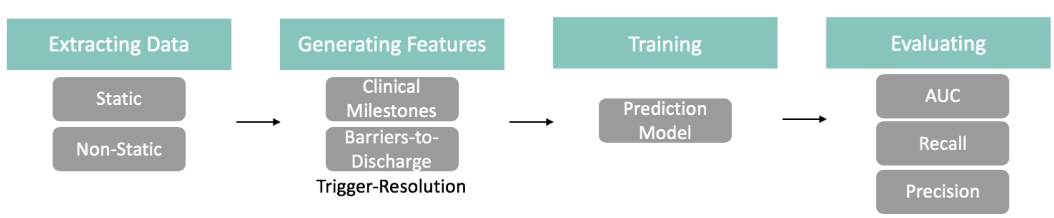

Figure 4-1: An overview of the methodology used to develop the model.

medical imaging outcomes, and laboratory results. This data can be used to determine whether a patient is progressing towards discharge or whether a patient requires further hospital care.

There are two types of patient data used by the model: static and dynamic. Section 4.1.1 describes static patient data, which is information that remains relatively constant during a patient's hospitalization, such as age or admission-source location. Section 4.1.2 describes dynamic patient data, which is information that is subject to change throughout a patient's stay, such as lab results or administered medications. These two types of data will be processed differently when generating the feature vectors to train the model.

The data extracted from EPIC can either be structured or unstructured. Structured data is information that follows a consistent order, such as laboratory results or demographic information. Unstructured data is information that is not organized in a particular format and that usually requires an additional preprocessing step. For example, providers' notes are unstructured, and necessitate natural language processing (e.g., sentiment anaysis) in order to turn them into features used by the model. A complete list of the raw data elements used by the model is provided in Appendix B.

4.1.1 Static Patient Data

This section describes the different types of static patient data that are extracted from EPIC. This information remains relatively constant throughout a patient's hospitalization.

Admission Source

A patient's admission source to the hospital can be used to estimate the level of care a patient requires. For example, a patient who arrives to MGH from home is probably in better condition than a patient who is transferred to MGH from a skilled nursing facility (SNF) [34].

Admission Source	Percentage
Self Referral	46%
Clinical Referral	34%
Outside Hospital	15%
Skilled Nursing Facility	4%
Other	1%

Table 4.1: The distribution of medicine inpatients' admission sources during 2017. 'Other' includes court, law enforcement, and hospice.

In addition, a patient's admission source can be used to evaluate the complexity of a patient's discharge process. For instance, a patient who is transferred from an SNF to MGH usually has a reserved bed in the SNF that initially cared for the patient. The discharge process of such a patient back to the SNF is considerably simplified, since a case manager does not need to send referrals to multiple facilities and the patient does not need to wait for an available bed.

The most common admission sources for medicine inpatients are: (1) self referral; (2) clinical referral; (3) outside hospital; and (4) skilled nursing facility (Table 4.1). Note that these admission sources refer to the location of the patient immediately prior to arriving to the hospital and *not* immediately prior to arriving on a medicine inpatient unit.

Demographic Information

Demographic information such as age, gender, marital status, and income impact the length-of-stay of medicine inpatients. These characteristics were added as features to the model as a way to cluster similar medicine inpatients together. Each of these features is described in more detail below.

- **Age**

  Age impacts a patient's rate of recovery and readiness towards discharge. For example, older patients take longer to recover from surgeries or procedures than younger patients [35]. A patient's age is used in order to distinguish between different age groups, such as children, adults, and seniors (Figure 4-2).

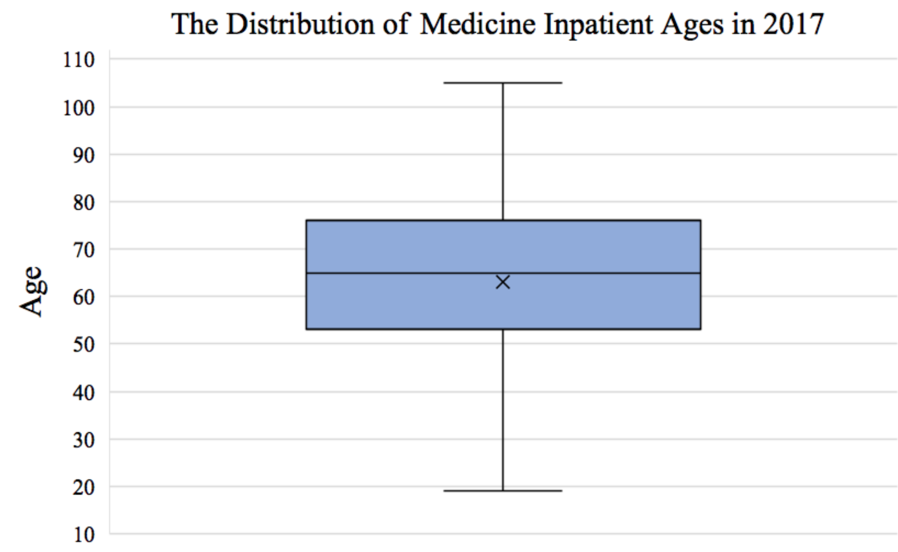

Figure 4-2: The distribution of medicine inpatient ages in 2017. The mean age is 63 years old and the standard deviation is 17.5 years.

- **Gender**

  Gender is used as a feature in order to capture gender differences in the treatment and recovery processes of patients.

- **Marital Status**

  Marital status can affect a patient's recovery after a surgery or procedure. For example, patients who are divorced, separated, or widowed have a greater chance of developing complications following cardiac surgery than married patients [36]. The intuition behind this is that patients without a spouse or partner do not have a social support system to lean on, whereas patients with a spouse or partner do. In order to capture this in the model, a feature is added to indicate whether or not the patient is married or has a partner.

- **Income**

  Income and financial well-being can impact a patient's recovery from injury and length-of-stay at the hospital [37]. The patient's income is estimated using census data, specifically the median household salary of residents with the same zip-code as the patient (the assumption is that a patient's income is equal to this

Diagnosis Group	Percentage
General Symptoms	8.6%
Chest Symptoms	6.2%
Cardiac Dysrhythmias	5.7%
Pneumonia	2.1%
Gastrointestinal Hemorrhage	2.1%
Abdomen or Pelvis Symptoms	1.9%
Heart Failure	1.8%
Acute Myocardial Infarction	1.7%
Cellulitis and Abscess	1.3%
Digestive System Symptoms	1.3%

Table 4.2: The distribution of the top ten diagnosis groups for medicine inpatients in 2017.

median salary). If a patient does not have a zip-code, this typically indicates that the patient is homeless or international.

Diagnosis

Upon admission to the hospital, a patient is assigned a principal hospital problem by a physician. A principal hospital problem is a specific description of the primary reason a patient was admitted to the hospital.

Since patients with the same hospital problem are assumed to have comparable pathways of care, features are added to the model to indicate the patient's diagnosis. However, instead of including each unique principal hospital problem to the model (there are over 4,000 of them), hospital problems that are clinically similar to each other are grouped together[1]. This divides the hospital problems into 200 diagnosis groups that cover 77% of the principal hospital problems assigned to medicine inpatients. The remaining principal hospital problems are classified as 'Other'. Table 4.2 presents the top ten diagnosis groups of medicine patients in 2017. Note that these features are the medicine-inpatient equivalent of the surgery type indicators used in Zanger's surgical discharge prediction model [9] [10].

[1]Partners Enterprise Data Warehouse, which is used to extract data from EPIC, provides a categorization of hospital problems into different groups. This categorization is used to group clinically similar hospital-problems together.

Location History

A patient's location history (i.e., the floors a patient has been on while at the hospital) can be used to group similar medicine inpatients together. In the model, the medicine floor that is caring for the patient is recorded. Additionally, features indicating whether a patient was in the ED or ICU while hospitalized are added, since such a patient is probably suffering from a more complex clinical condition than a patient who was consistently located on one or several medicine floors. Note that although the medicine floor caring for a patient can change during a patient's hospitalization, it usually remains constant.

Patient Service

The patient population of study is treated by either the medicine, cardiology, oncology, or neurology patient service. Features that indicate which patient service cared for an inpatient were added the model as a way to cluster similar medicine inpatients together. Note that although the patient service caring for a patient can change during hospitalization, it usually remains constant. In addition, notice that these features are the medicine-inpatient equivalent of the surgical service indicators used in Zanger's surgical discharge prediction model [9] [10].

4.1.2 Dynamic Patient Data

This section describes the different types of dynamic patient data that was extracted from EPIC. This information is subject to change throughout a patient's hospitalization and, unless otherwise mentioned, is structured.

4Next Referrals and Acceptances

4Next is a web-based application developed by Partners HealthCare. It is used by case managers at MGH to electronically manage a patient's post-discharge facility or home-with-services arrangements (Figure 4-3), and to simultaneously communicate

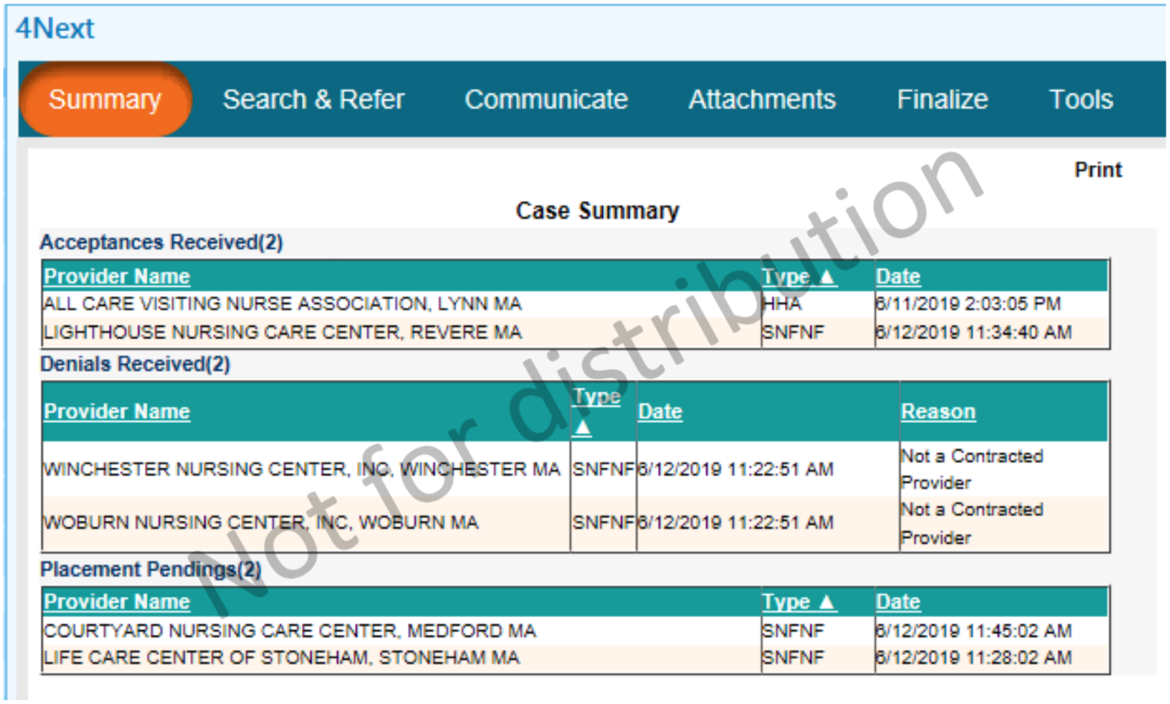

Figure 4-3: A screenshot of 4Next's Summary page embedded within EPIC (June, 2019). On this page, a case manager is able to view which facilities or home-with-services accepted or denied a patient, as well as which facilities or home-with-services are pending placements.

with various post-acute care services (Figures 4-4, 4-5). This application is updated in real-time and is currently embedded within EPIC [38].

Approximately 50% of medicine inpatients are discharged to a care facility or receive home-with-services arrangements. In order to predict the readiness of discharge for such patients, 4Next data is used to monitor the communication between case managers and facilities or home-with-services. In particular, for each patient, features were added that mark (1) whether a referral was sent to a facility or home-with-service (2) whether a facility or home-with-service accepted the patient and (3) whether the facility or home-with-service denied the patient or does not have an available bed.

This information was not used in Zanger's surgical discharge prediction model [9] [10].

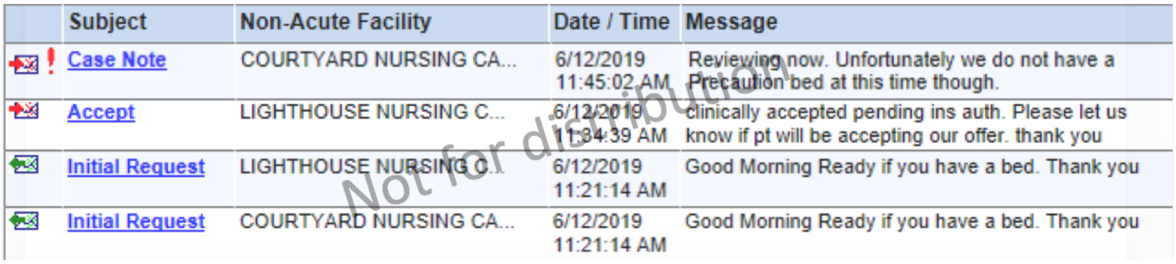

	Subject	Non-Acute Facility	Date / Time	Message
!	Case Note	COURTYARD NURSING CA...	6/12/2019 11:45:02 AM	Reviewing now. Unfortunately we do not have a Precaution bed at this time though.
	Accept	LIGHTHOUSE NURSING C...	6/12/2019 11:34:39 AM	clinically accepted pending ins auth. Please let us know if pt will be accepting our offer. thank you
	Initial Request	LIGHTHOUSE NURSING C...	6/12/2019 11:21:14 AM	Good Morning Ready if you have a bed. Thank you
	Initial Request	COURTYARD NURSING CA...	6/12/2019 11:21:14 AM	Good Morning Ready if you have a bed. Thank you

Figure 4-4: A screenshot of the communication messages between two facilities and a case manager on the 4Next application (June, 2019).

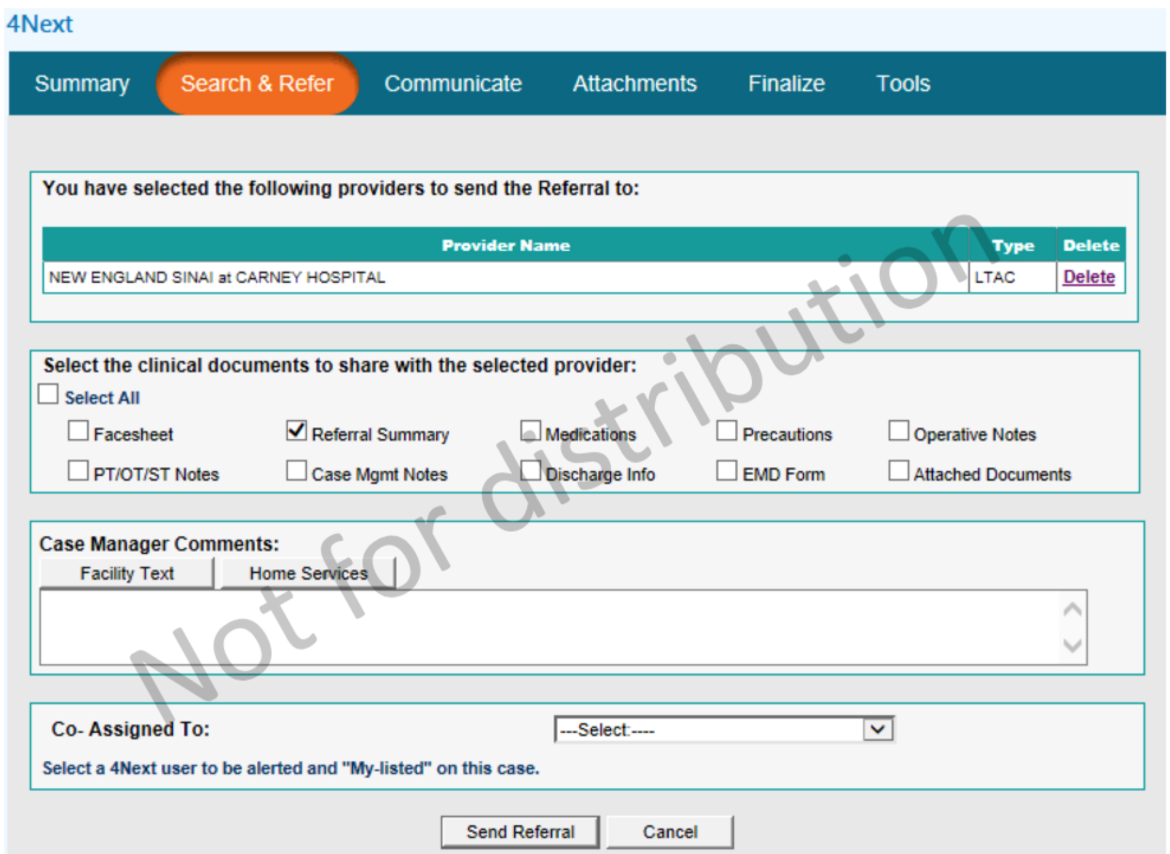

Figure 4-5: A screenshot of 4Next's Search & Refer tab embedded within EPIC (June, 2019). This tab allows a case manager to search, send a referral, or provide clinical data or comments to a facility or home-with-service.

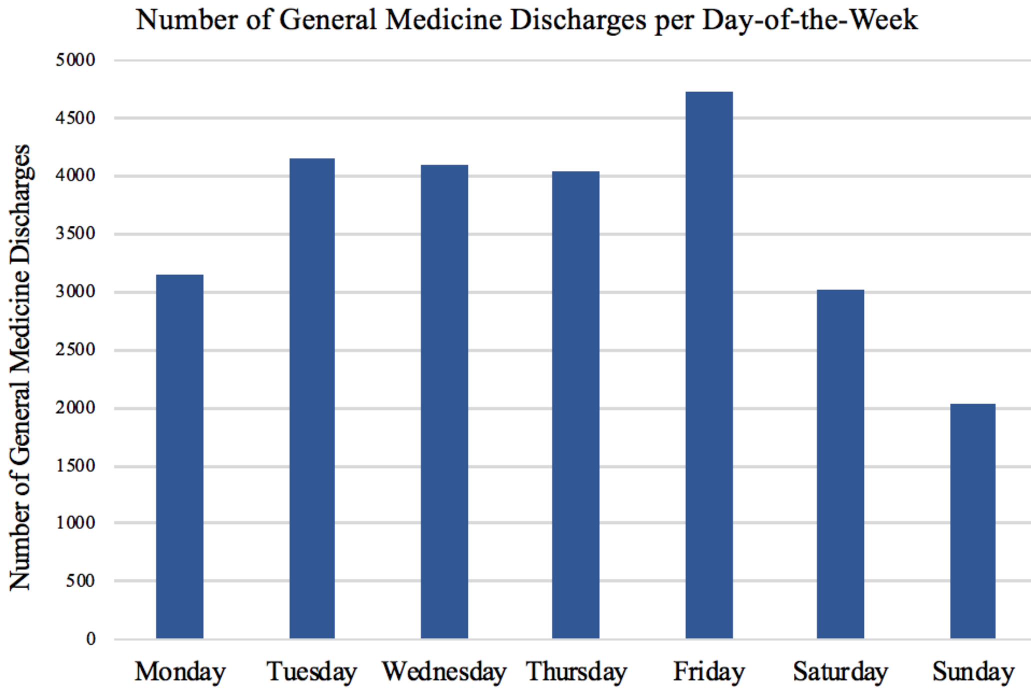

Figure 4-6: The number of medicine inpatient discharges per day-of-the-week in 2018. Notice that on weekends (Saturday and Sunday), the number of discharges is noticeably lower than the number of discharges on weekdays (Monday through Friday). Also notice that the number of discharges on Monday is lower compared to the number of discharges on other weekdays. This is because multiple federal holidays occur on Monday, including Memorial Day, Labor Day, Columbus Day, and Martin Luther King Jr. Day.

Calendrical Information

Patient-care units sometimes operate with partial staff on weekends or federal holidays. This affects a patient's pathway of care at the hospital, as well as the hospital's discharge process. For example, because more staff is available to care for patients and their discharge processes on Friday than over the weekend, more patients get discharged on Fridays than on subsequent weekend days (Figure 4-6) [39]. In order to take this into account, features were added that record the day-of-the-week in the prediction model and that indicate whether or not it is a federal holiday.

Flowsheet Records

Flowsheet records allow care providers to track specific data about a patient, such as vital signs, ongoing assessments, or lab results. In addition, flowsheets can be used to insert structured evaluations recorded by bedside nurses, case managers, physical therapists and speech-language pathologists. Since flowsheets contain data that is relevant for predicting a patient's readiness for discharge, flowsheets-related features are included in the model. For instance, several flowsheets that are monitored in the model include vital signs (blood pressure, body temperature, pulse rate, respiration rate, oxygen saturation), physical therapist evaluations (discharge recommendation, number of expected physical therapy sessions before discharge), and emesis, urine, and stool occurrence.

Labs

Laboratory tests are typically carried out in order to diagnose a patient, provide a benchmark, or monitor the effect of treatment drugs. Since the results of laboratory tests are crucial to predicting whether a patient is clinically ready for discharge, the following labs are monitored in the model: (1) Alkaline Phosphatase (ALP); (2) Alanine Aminotransferase (ALT); (3) Aspartate Aminotransferase (AST); (4) Bilirubin; (5) Creatinine; (6) Glucose; (7) Hemoglobin; (8) PT-INR; (9) Potassium; (10) Lactate Acid; (11) Sodium; (12) NT-BNP; (13) Troponin; and (14) White Blood Count. ALP, ALT, AST, Bilirubin, Lactate Acid, and NT-BNP were not included in Zanger's surgical discharge prediction model [9] [10].

LDAs

Lines, Drains, and Airways (LDAs) are medical devices used to deliver or remove air and fluids from the body. In 2017, approximately 35% of medicine inpatients had at least one LDA while hospitalized at MGH (Table 4.3). Since patients who have LDAs are usually not ready for discharge, features were added that record whether or not a patient has an LDA.

LDA	Percentage
Urinary Catheter	55%
GI Tube	17%
Chest Tube	8%
PICC	8%
Suction Drain	6%
Other	6%

Table 4.3: The distribution of LDAs medicine patients' had during 2017. 'Other' includes negative pressure wound therapy, biliary drain, penrose drain, epidural catheter, and external fecal management system.

Medications

Medications are pharmaceutical drugs used to cure, treat, diagnose or prevent diseases. Since the administration of certain medications indicates that a patient is not clinically ready for discharge, EPIC is used to identify when a medication treatment was initiated and when it was terminated. In particular, whether a patient receives one of the following medications is monitored in the model: (1) Antiemetics; (2) Antipsychotics; (3) Colonoscopy Preparation Drugs; (4) Diuretics; (5) Heparin; (6) IV Antibiotics; (7) IV Fluids; (8) Narcotics; (9) Nebulizers; (10) Nitroglycerin; (11) Phenobarbital; (12) Steroids; (13) Tamsulosin; (14) Total Parenteral Nutrition (TPN); (15) Tube Feeds; and (16) Warafin. Colonoscopy preparation drugs, IV fluids, narcotics, nebulizers, nitroglycerin, phenobarbital, steroids, tamsulosin, TPN, tube feeds, and warafin were not included in Zanger's surgical discharge prediction model [9] [10] .

Orders

Care providers place different types of orders throughout a patient's hospitalization. These orders include consult referrals, diet requests, and isolation precautions. Using EPIC, it is currently possible to capture when an order was placed, and in most cases, when an order was carried out. The types of orders that are included in the discharge prediction model are listed below.

- **Consults**

  Consult orders are electronic referrals requested by physicians in order to receive recommendations from specialists and to determine the next stage of a patient's care, including discharge if applicable. A consult order is scheduled prior to a patient's discharge, and therefore indicates that a patient is not ready to get discharged from the hospital. The timestamp of when a consult order was requested is captured and the timestamp of the corresponding consult note is used to determine when a specialist saw the patient. Several of the consult orders that are included in the model are physical therapy, occupational therapy, psychology, addiction services, and nutrition services.

- **Diet**

  Physicians typically order specific diets for a patient as part of the patient's treatment plan. Examples of diets include NPO[2], clear liquid, low fat, sodium restriction, and tube feeding. Since a patient's diet can be used to determine clinical readiness for discharge, data from EPIC is used to record when a diet order was placed as well as what type of diet was given to a patient.

- **Heart Procedures**

  Heart procedures are used to evaluate, diagnose and treat the heart of a patient. In particular, the discharge prediction model monitors when catheterization labs and electrophysiology (EP) studies are ordered and completed. These procedures were not included in Zanger's surgical discharge prediction model [9] [10].

- **Imaging**

  Imaging involves creating visual representations of the interior body in order to diagnose, monitor, and treat the patient. Examples of imaging orders include requests for computed tomography (CT) scans, ultrasounds, and magnetic resonance imaging (MRI). Since imaging technology plays a pivotal role in patient diagnosis and treatment, features were added that detect whether an imaging

[2] Nil per os (NPO) is a Latin phrase that means "nothing through the mouth".

order has been requested and whether the imaging results are available in the discharge prediction model. Echocardiogram imaging orders were not included in Zanger's surgical discharge prediction model [9] [10].

- **Isolation Status**

  In health care facilities and hospitals, isolation precautions are implemented to prevent the spread of germs. The three most-common types of isolation precautions used at MGH are contact, droplet, and airborne:

  1. <u>Contact Isolation</u>

     Contact isolation is used for germs, organisms, diseases or infections that are spread by touching the patient or items in the room. Examples of medical conditions that require contact isolation include open wounds, diarrheal illness, and VRE [40].

  2. <u>Droplet Isolation</u>

     Droplet isolation is used for germs or diseases that spread through the air by coughs or sneezes. Patients who have bacterial meningitis, influenza, or pneumonia require droplet isolation [40].

  3. <u>Airborne Isolation</u>

     Airborne isolation is used for infections or small germs that travel by air currents throughout the hospital unit and onto different floors. Several examples of infections or diseases that require airborne isolation are measles, chickenpox, and tuberculosis [40].

  Approximately 28% of medicine patients were in isolation at least once while hospitalized at MGH. Since the isolation status of a patient can impact length-of-stay [41], data from EPIC is used to record whether or not a patient is in contact, droplet, or airborne isolation.

  This information was not previously in Zanger's surgical discharge prediction model [9] [10].

Microbiology	Percentage
Urine	38%
Blood	33%
Fluid Culture Smear	8%
Fungal	7%
Stool	4%
Wound & Would Culture Smear	4%
Anaerobic	3%
Other	3%

Table 4.4: The distribution of microbiology tests medicine inpatients' had during 2017. 'Other' includes mycobacterial, throat, genital, and Group B Strep Infection (GBS).

- **Laboratory Orders**

  Laboratory orders are requested by physicians in order to diagnose, monitor and treat the patient. Since laboratory tests can be used to determine a patient's clinical readiness for discharge, EPIC data is used to capture when a physician has placed a procedure order and when the order has been completed. Several examples of laboratory orders that are included in the model are RBC, Plasma, Cryoprecipitate, and Platelets.

- **Microbiology**

  Medical microbiology involves the diagnosis, treatment, and prevention of infectious diseases. The discharge prediction model currently captures when a care provider requested a microbiology test for a patient and when the results are available. Several of the microbiology tests that are monitored include anaerobic, blood, fluid culture smear, fungal, stool, throat, wound, and urine tests (Table 4.4).

  This information was not previously in Zanger's surgical discharge prediction model [9] [10].

Notes

Digital notes are used by care providers to analyze case details, document findings, and compare a patient's previous and current status. Notes are generally unstructured, and contain a significant number of abbreviations, acronyms, and medical terms.

- **Case Management Notes**

  Upon admission to the hospital, each patient is assigned to a case manager. A case manager is responsible for working with the patient, the patient's family, care team, and insurer to coordinate a comprehensive discharge plan. Several case-management responsibilities include sending referrals to facilities, coordinating home-with-services treatments, or scheduling an ambulance to pick up the patient.

  Case managers write notes that summarize actions taken to advance discharge. For example, case managers may use progress notes to summarize events that occurred on 4next or conversations that happened with a patient's family. When writing notes, case managers can use various structured templates called SmartTexts and SmartPhrases. SmartTexts and SmartPhrases are standard templates or blocks of text that can be inserted and edited easily in many places throughout EPIC, such as in notes, patient instructions, and letters. A description of several templates that are kept track of in the discharge prediction model are described below. Note that SmartTexts and SmartPhrases were not used in Zanger's surgical discharge prediction model [9] [10].

  - *VNA Referral*

    A case manager can use a "VNA Referral" SmartPhrase to automatically insert a block of text that indicates that a VNA referral was sent for a patient. Recall that VNA Care is a nonprofit home health care organization that provides nursing care, therapy, hospice and palliative services to patients throughout Massachusetts [42]. This note can be used to keep track of a patient's discharge plan.

- *Facility Screening*

  A "Facility Screening" SmartText can be used to indicate that a facility has initiated the screening process for a patient (i.e., checking whether a patient meets certain criteria). This patient will most-likely remain in the hospital until a facility accepts the patient.

- *Facility Transfer*

  A "Facility Transfer" SmartText can be used by a case manager to indicate that a facility has accepted the patient and that the patient is in agreement with the discharge plan. This suggests that a patient will most likely leave the hospital within 24 hours.

- *Readmission*

  A case manager can use a Readmission SmartText to indicate that a patient had another hospitalization at MGH less than 30 days ago. The inclusion of this note can give a better sense of the level of care a patient requires, since readmission to the hospital may signify that the patient is suffering from a more severe condition than other patients.

- **Physician, Nursing, Physical Therapy, and Social Worker Notes**

  Physician, nursing, physical therapy, and social worker notes record medically or discharge-related events that occur throughout a patient's hospitalization. Regular expressions[3] are used to identify dispo phrases within these types of notes. A dispo phrase begins with the character sequence "Dispo:" or "Disposition:" and signifies a patient's progress towards discharge. For example, the dispo phrase "Dispo: home likely tomorrow" suggests that a patient will get discharged within 24 hours, whereas the dispo phrase "Dispo: pending improvement of cellulitis" implies that a patient will remain in the hospital for the near future. Note that dispo phrases are *unstructured*, and were not used in Zanger's surgical discharge prediction model [9] [10].

[3] A regular expression is sequence of characters used for matching one more characters within a string.

Patient Service	Number
Medicine	1
Cardiology	2
Oncology	3
Neurology	4

→

	1	2	3	4
Medicine	1	0	0	0
Cardiology	0	1	0	0
Oncology	0	0	1	0
Neurology	0	0	0	1

Figure 4-7: In the left table, each patient service is assigned to a number that corresponds to a location in the one-hot encoded vector. In this case, medicine corresponds to the first entry, cardiology corresponds to the second entry, oncology corresponds to the third entry, and neurology corresponds to the last entry. The resulting one-hot encoded vector for each patient service is represented in the right table. For instance, the one-hot encoded vector for oncology would be $[0, 0, 1, 0]$.

4.2 Features Generation

A feature vector is generated for each time a medicine inpatient stayed overnight in the hospital, beginning from the day of an inpatient's admission until the day before discharge. The daily feature vector represents the state of the inpatient based on the static and dynamic data recorded in EPIC before 23:59 on that day. Note that it is possible to use a different time instead of 23:59 and to generate multiple feature vectors per day (e.g., 9:00, 12:00, 15:00, 18:00, 21:00, 00:00). The process of transforming the static and dynamic data into features that can be used in the discharge prediction model is described below.

4.2.1 Static Data

The admission source, demographic information (gender, marital status), diagnosis, location history, and patient service are transformed using one-hot encoding. One-hot encoding is a process by which categorical variables are converted into binary vectors. For example, recall that a medicine patient can be treated by the medicine, cardiology, oncology, or neurology patient service. Since there are four different categories, the patient service taking care of an inpatient can be expressed using a 4-dimensional vector that contains one 1 and three 0s (Figure 4-7).

The numeric value of a patient's age and income are used in the feature vector.

If the age of a patient is missing from the data, it is replaced using the mean age of patients who had an age recorded in EPIC [43].

4.2.2 Dynamic Data

The dynamic data (except for calendrical information, several labs, isolation status, and notes) is transformed into features using a trigger-resolution framework. Recall that in the trigger-resolution framework, a clinical milestone corresponds to a binary (i.e., 0 or 1) feature that indicates whether the milestone has occurred. A barrier-to-discharge corresponds to two binary features: (1) a trigger feature that marks whether the barrier-to-discharge has occurred; and (2) a resolution feature that signals whether the barrier has been resolved. Note that the trigger-resolution framework can be used to identify clinical milestones that have not yet happened as well as discharge barriers that have not yet been completed during a medicine inpatient's hospitalization.

Below, a description of how the trigger-resolution framework can be applied to each type of dynamic data is provided. The complete list of flowsheet records, labs, LDAs, medications, consults, microbiology tests, and procedure orders that are monitored is listed in Appendix B.

4Next Referrals and Acceptances

Patients who have not yet heard back from a home-with-service or facility will most likely remain in the hospital until a home-with-service or facility has accepted them. In order to take this into account in the model, for each post-acute care setting (i.e, home-with-services, facility) two binary features are created: (1) a trigger feature that marks that a referral was sent to the post-acute care setting; and a (2) resolution feature that indicates that the patient was accepted by the post-acute care setting (Figure 4-8).

Facility Referral	Facility Accept
01-02-2018 18:02	01-04-2018 12:21

→

Prediction Time	Facility Trigger	Facility Resolution
01-01-2018 23:59	0	0
01-02-2018 23:59	1	0
01-03-2018 23:59	1	0
01-04-2018 23:59	1	1

Figure 4-8: An example of the trigger-resolution features for 4Next data. The left table corresponds to sample 4Next data. The right table corresponds to the trigger-resolution features that were generated using the sample 4Next data.

Flowsheet Records

Flowsheet records contain text, numeric, or boolean data that is relevant for predicting a patient's readiness for discharge. The trigger-resolution framework is applied to each type of data.

- *Text*

  Several of the flowsheet records that are monitored contain structured text (e.g. cardiac rhythm, orientation level). For each of these flowsheet records, two binary features are created: (1) a trigger feature that marks whether a certain phrase has appeared in the flowsheet record during a patient's hospitalization; and (2) a resolution feature that indicates whether the most-recent flowsheet record did not contain a certain phrase. For example, the phrases used to detect barriers in the cardiac rhythm flowsheet record are "sinus" and "atrial fibrillation". The phrases that were chosen to detect barriers were determined by collaborating with clinicians at MGH.

- *Numeric*

  Flowsheet records may contain numeric values (e.g. temperature, pulse). For each of these types of flowsheet records two features are created: (1) a trigger feature that identifies whether the patient had an abnormal numeric flowsheet value during the hospitalization; and a (2) resolution feature that marks whether the most-recent flowsheet record was normal. For example, an abnormal pulse value is below 50 or above 120.

- *Boolean*

  Several flowsheet records are used to indicate whether a clinical milestone has occurred (e.g., urine, stool). For each of these types of flowsheet records, a binary feature is created that identifies whether the milestone has occurred.

Labs

The barriers (abnormal values) for each lab were determined by collaborating with clinicians at MGH. For each barrier, two binary features were created: (1) a trigger feature that indicates whether the patient had the barrier during the hospitalization; and (2) a resolution feature that marks that the barrier has been resolved. Below, a list of the labs (and their corresponding barriers) that used the trigger-resolution framework is provided.

- Creatinine Barriers:
  - The lab value is greater than 1.5 mg/dL.
  - The lab value is 30% greater than the patient's minimal creatinine level recorded so far.
- Glucose Barriers:
  - The lab value is less than 70 mg/dL.
  - The lab value is greater than 300 mg/dL.
- Hemoglobin Barriers:
  - The lab value is less than 7 g/dL.
  - The maximum lab result seen so far minus the most-recent lab result is greater than 2.
- PT-INR Barrier: the ratio is greater than 3.
- Potassium Barriers:

  - The lab value is less than 3 mmol/L.
  - The lab value is greater than 5.5 mmol/L.

- Lactate Acid Barrier: the lab value is greater than 2 mmol/L.
- Sodium Barriers
  - The lab value is less than 130 mmol/L.
  - The lab value is greater than 147 mmol/L.
- Troponin Barrier: the current lab value is above 0.01 ng/mL *and* each of the previous lab values are increasing.
- White Blood Count Barrier: the lab value is greater than 14 K/uL.

LDAs

Patients who have LDAs are usually not ready for discharge (i.e., a barrier). For each LDA that was monitored, two binary features were created: (1) a trigger feature that marks whether the patient had the LDA during the hospitalization; and (2) a resolution feature that indicates that the patient no longer has the LDA.

Medications

Medications which signal that a patient is not clinically ready for discharge (i.e., a barrier) were identified by collaborating with clinicians at MGH. For each medication that was kept track of, two binary features were created: (1) a trigger feature that marks whether the patient was given the medication during the hospitalization; and (2) a resolution feature that indicates that the patient has not received the medication in the past 24 hours.

Orders

- *Consults*: A consult order is scheduled prior to a patient's discharge, and therefore indicates that a patient is not ready to get discharged from the hospital (i.e, a barrier). For each consult in the model, two binary features are created: (1) a trigger feature that marks whether the consult was requested during the patient's hospitalization; and (2) a resolution feature that indicates that the consult saw the patient. This is detected by checking whether a consult wrote a note.

- *Diet*: A patient's diet can be used to determine clinical readiness for discharge. In order to take this into account, two binary features are created: (1) a trigger feature that marks whether a patient's diet was ever *non*-regular (e.g., diet NPO) during the hospitalization; and (2) a resolution feature that indicates that the most-recent diet of the patient is regular.

- *Heart Procedures*: A heart procedure can be used by physicians to evaluate, diagnose and treat a patient. As a result, if a heart procedure has been ordered, a patient will most likely not leave the hospital until the procedure is completed (i.e., a barrier). For each heart procedure that is monitored, two binary features are created: (1) a trigger feature that marks whether the heart procedure order was ever ordered during the hospitalization; and (2) a resolution feature that indicates that the heart procedure was completed.

- *Imaging*: An imaging order is usually scheduled before a patient is discharged. Therefore, if an imaging order has been requested this typically indicates that the patient will not leave the hospital until the imaging results are available (i.e, a barrier). Two binary features are created: (1) a trigger feature that marks whether an imagining order was ever placed during the hospitalization; and (2) a resolution feature that indicates that the results of all the imagining orders that were requested are available.

- *Laboratory Orders*: A laboratory order is scheduled prior to a patient's discharge. Therefore, a patient will usually not leave the hospital until the procedure has been completed (i.e., a barrier). For each of the laboratory orders in the model, two binary features are created: (1) a trigger feature that marks whether the laboratory order was ever scheduled during the hospitalization; and (2) a resolution feature that indicates that the laboratory order was completed.

- *Microbiology*: Microbiology tests are scheduled prior to a patient's discharge. Therefore, a patient will most likely not leave the hospital if the results from a microbiology test are not available. For each of the microbiology tests that is monitored, two binary features are created: (1) a trigger feature that marks whether the specific microbiology test was ever scheduled during the hospitalization; and (2) a resolution feature that indicates that the result of the specific microbiology test is available.

Calendrical Information

One-hot encoding is used to indicate the day-of-the-week and whether or not it is a federal holiday. The one-hot vector is 8-dimensional: 7 dimensions correspond to each day of the week, and 1 dimension corresponds to whether it is a federal holiday or not.

Remaining Labs

The trigger-resolution framework was not used for the labs ALP, ALT, AST, Bilirubin, and NT-BNP. This is because these labs do not have a strict abnormal range that can be used to detect barriers. Instead, the patient's most recent lab value is recorded as a feature.

Isolation Status

Binary features that indicate whether at the time of prediction the patient was in contact, droplet, or airborne isolation were added. This corresponds to 3 features.

Notes

- *Case Management*

  One-hot encoding is used to record whether a SmartText or SmartPhrase was selected in a note (1 if a certain SmartText was used; 0 if it was not).

- *Physician, Nursing, Physical Therapy, and Social Worker*

  Sentiment analysis[4] is used in order to determine whether or not a dispo phrase suggests that a patient will get discharged within 24 hours. In particular, a sentiment classifier was built that takes a dispo phrase as input and returns a score that indicates whether or not a patient will get discharged within 24 hours from the hospital. This model makes use of an embedding layer[5] and an LSTM. Note that the sentiment classifier is a separate model from the medicine discharge prediction tool. For further details, refer to Appendix C.

Preprocessing

After the feature vectors have been formed, they are preprocessed using the MinMax scaler from the `scikit-learn` machine learning library. This scales and translates each feature in the feature vector such that it is between zero and one [44].

4.3 Predictive Classification Model

The discharge prediction algorithm is structured as a classification problem. The model takes an inpatient's current and past data from the hospitalization as input, and outputs a score indicating whether the inpatient will be discharged within the next 24 hours. A neural-network, LSTM based classifier was selected for the prediction algorithm for the following reasons: (1) it can learn non-linear interactions given inpatients' data; and (2) it makes classification decisions based on a patient's current and past data from the hospitalization. Notice that if an inpatient has stayed at the hospital for a certain number of days, such as four, it is unclear whether the inpatient's data from the past two days, three days, or four days should be used to

[4]Sentiment analysis refers to the use of natural language processing to understand subjective information written in text.

[5]A word embedding is a natural language processing technique in which words from a vocabulary are mapped to dense vectors. An embedding layer maps a word to its dense-vector representation.

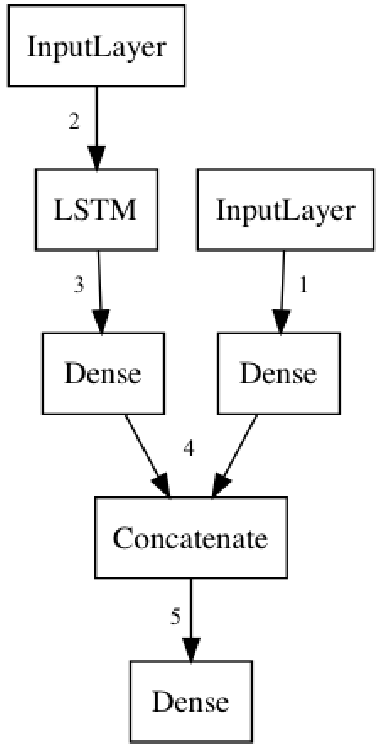

Figure 4-9: A visualization of the discharge prediction model's architecture. A fully-connected dense layer takes static data as input (1), while the LSTM takes dynamic data as input (2). The output of the LSTM goes through a fully-connected layer (3), and is then concatenated with the transformed static data (4). This is then used to output a prediction score (5).

predict discharge. The number of previous days of inpatient data that should be used in the model is a hyperparameter[6] that requires tuning.

The input of the model contains both static and dynamic types of data. In order to account for this, a model that is capable of accepting both static and dynamic input was built. Specifically, the static data is transformed using a fully-connected layer, while the dynamic data is fed through an LSTM and a fully-connected layer. The outputs from the fully-connected layers are concatenated and then used to output a prediction score. This score is used to determine whether the inpatient will get discharged from the hospital within 24 hours. For a visualization of the model's structure, refer to Figure 4-9.

The classifier requires tuning several hyperparameters:

[6]A hyperparameter is parameter whose value is set before training begins.

- The number of hidden neurons in each of the fully-connected layers.

- The number of hidden neurons in the LSTM layer.

- The number of previous days of inpatient data that should be used as input into the LSTM.

To tune these hyperparameters, grid-search with 10-fold cross-validation[7] is used. In the grid-search approach, a manually specified subset of the hyperparameter space is exhaustively trained. Afterwards, the results of each training is evaluated and the hyperparameters that maximize the AUC-ROC metric are returned (see Section 4.4). To prevent overfitting, dropout[8] is used.

The discharge prediction model's optimal hyperparameters were (i) 16 hidden neurons for the fully-connected layer that takes static data as input; (ii) 32 hidden neurons for the fully-connected layer following the LSTM; (iii) 50 hidden neurons for the LSTM; and (iv) 7 previous days of inpatient data (Figure 4-10). If an inpatient was in the hospital for less than 7 days, the input is pre-padded with zeros (Figure 4-11).

4.4 Evaluating the Model

The discharge prediction model outputs a score between 0 and 1. Scores near 0 indicate that a patient will remain in the hospital for at least the next 24 hours, whereas scores near 1 indicate that a patient will get discharged within 24 hours from the hospital. It is less clear, however, what a score of 0.33 or 0.65 represents. Although it is tempting to assume that scores above 0.5 indicate that a patient will

[7] K-fold cross-validation is a technique used to assess how the results of a model will generalize to an independent data set. This technique involves dividing the dataset into K groups. For each unique group, (1) the test set is comprised of the unique group while the training set is comprised of the remaining groups; and (2) the model is fitted and evaluated using these sets. Once this process has been repeated for each of the K groups, the K results can be compared to determine how well the model will generalize.

[8] Dropout is a regularization method in which several of the layers outputs are randomly ignored. Dropout causes the training process to be noisy, which forces nodes to learn more robust features.

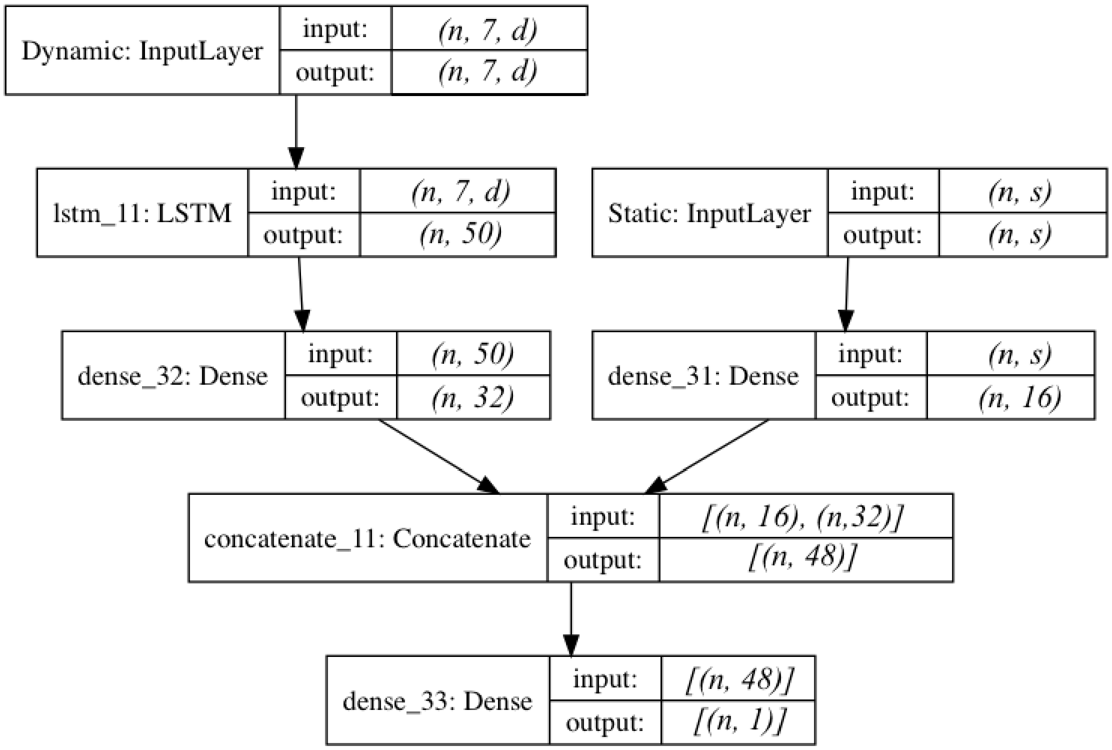

Figure 4-10: A visualization of the input shapes of the discharge prediction model. n represents the number of feature vectors, d represents the number of dynamic features, and s represents the number of static features. Notice that the following hyperparameters were used (i) 16 hidden neurons for the fully-connected layer that takes static data as input; (ii) 32 hidden neurons for the fully-connected layer following the LSTM; (iii) 50 hidden neurons for the LSTM; and (iv) 7 previous days of inpatient data.

	Facility Barrier	**Facility Resolution**
1	0	0
2	0	0
3	0	0
4	1	0
5	1	0
6	1	0
7	1	1

Figure 4-11: This figure shows a simplified example of the dynamic data that is fed into the LSTM before it has been preprocessed. Each row corresponds to a day the patient was at the hospital. Since 7 was found to be the optimal number of previous days of inpatient data that should be used as input into the LSTM model, the table above has 7 rows. In this particular example, the patient has been in the hospital overnight for 4 days, so the first three rows are assigned to zero.

get discharged within 24 hours and scores below 0.5 indicate the opposite, a decision threshold of 0.5 may not be optimal for this problem.

In order to determine a decision threshold for the classification algorithm, the discharge model's prediction scores for out-of-sample feature vectors were found. Using these prediction scores, the metrics below were calculated for 100 distinct thresholds that are evenly spaced (i.e., $0.01, 0.02, \ldots, 0.98, 0.99$):

Metrics

1. True Positives (TP)

   The number of times the model correctly predicts that a patient will get discharged within 24 hours.

2. False Positives (FP)

   The number of times the model predicts that a patient will get discharged within 24 hours, but the patient ends up staying in the hospital.

3. True Negatives (TN)

   The number of times the model correctly predicts that a patient will *not* get discharged within 24 hours.

4. False Negatives (FN)

   The number of times the model predicts that a patient will *not* get discharged within 24 hours, but the patient ends up leaving the hospital.

5. Precision

   Precision is defined mathematically as $\frac{TP}{TP+FP}$. This metric represents how often the model is correct when it predicts that a patient will get discharged within 24 hours.

6. Recall

   Recall is defined mathematically as $\frac{TP}{TP+FN}$. This metric represents how often the model is able to identify patient discharges (i.e., the fraction of discharges captured). This metric is also referred to as sensitivity or the true positive rate.

The threshold that maximizes both precision and recall was used in order to evaluate the model's results. Notice, however, that there is a tradeoff between these two metrics:

- When increasing the threshold, the number of false positives decreases but the number of false negatives increases. As a result, precision $\left(\frac{TP}{TP+FP}\right)$ increases while recall $\left(\frac{TP}{TP+FN}\right)$ decreases.
- When decreasing the threshold, the number of false positives increases but the number of false negatives decreases. As a result, precision $\left(\frac{TP}{TP+FP}\right)$ decreases while recall $\left(\frac{TP}{TP+FN}\right)$ increases.

Since the aim is for *both* the precision and recall values of the model to be as large as possible, the threshold that minimizes the absolute difference between precision and recall is selected. For example, in Figure 4-12 the threshold of 0.28 is chosen. Note that the threshold that maximizes the sum of the precision and recall metrics is not used because this may result in a large precision value but a small recall value, or vice versa.

There are two additional metrics that are used to evaluate the model:

1. AUC-ROC

   A Receiver Operating Characteristic (ROC) graph is created by plotting the true positive rate $\left(\frac{TP}{TP+FN}\right)$ against the false positive rate $\left(\frac{FP}{FP+TN}\right)$ at various threshold settings. Since the aim is to have a small number of false negatives and a large number of true negatives, this corresponds to targeting a large true positive rate and a small false positive rate.

   The area under the ROC graph (AUC-ROC) represents how well model is able to distinguish between patients who will get discharged within 24 hours and patients who will remain at the hospital for at least 24 more hours. An AUC-ROC

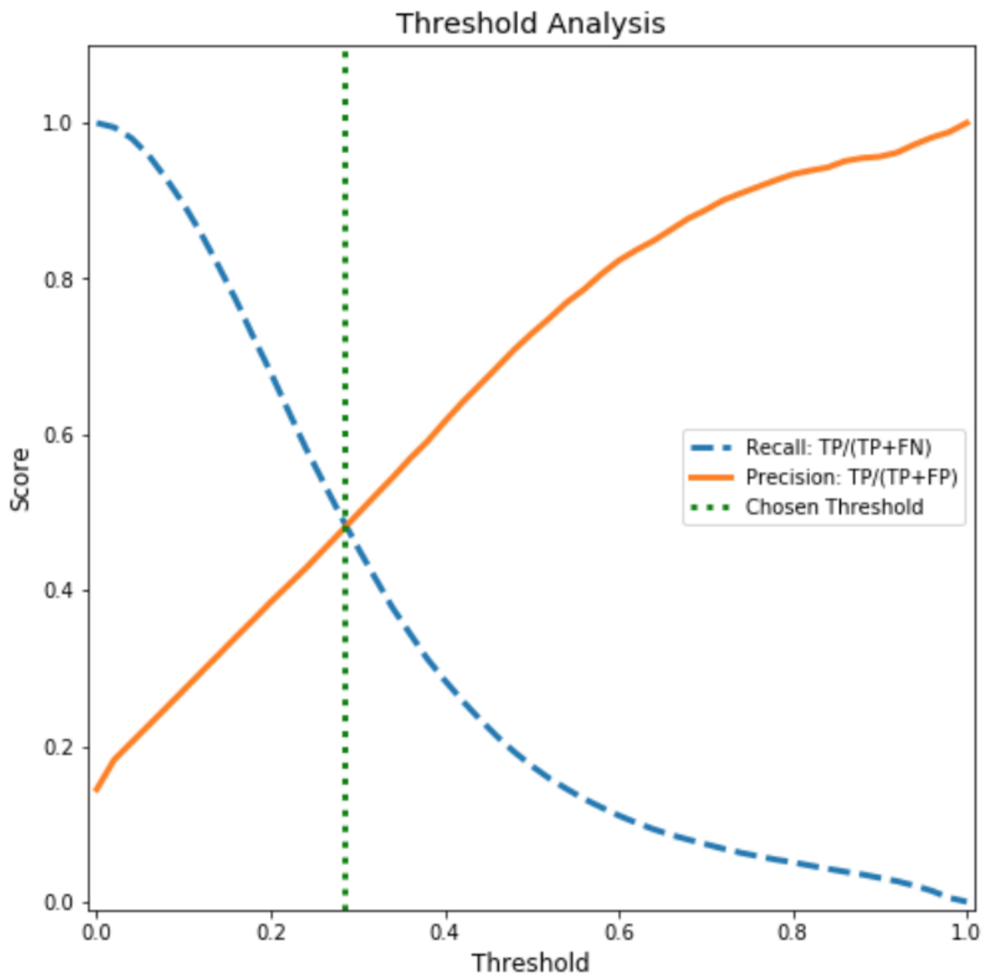

Figure 4-12: An example of a plot that illustrates the precision and recall values for various thresholds. A threshold of 0.28 is chosen since it minimizes the absolute difference between precision and recall.

of 1 represents a perfect classification model, while an AUC-ROC of 0.5 represents a classifier that no has class separation capacity. 10-fold cross-validation is used to estimate the model's AUC-ROC.

2. Precision (48 hours)

   The model is trained and tested on the discharge decisions made the by clinical teams at MGH. As a result, even though a patient may be clinically ready for discharge, the patient may end up staying at the hospital due to a physician decision or a post-hospital care constraint. For this reason, the precision of discharge within 48 hours is also evaluated. This represents how often a patient is predicted to be discharged within 24 hours by the model, but is actually discharged within 48 hours.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 5

Model Performance

This chapter presents the medicine discharge prediction model's results, and compares them to the results of two distinct discharge prediction models: (1) an *enhanced* version of Zanger's surgical model that uses additional patient data, the trigger-resolution framework, and an LSTM-based architecture; and (2) a modified version of the medicine model (*medicine-minus*) that does not make use of the techniques mentioned above.

The medicine model's results that are presented in this chapter were achieved using the data of MGH medicine inpatients who were hospitalized sometime between May 2016 and September 2018. This population included 60,993 inpatients that stayed in the hospital overnight for a total of 421,343 days. Section 5.1 evaluates the model's AUC-ROC, and Section 5.2 analyzes the model's precision and recall metrics.

5.1 ROC Curve

The medicine model's AUC-ROC was calculated using 10-fold cross-validation (Figure 5-1). The model achieved an average AUC-ROC of 0.830 (SD 0.002), which is lower than Zanger's enhanced model's average AUC-ROC of 0.912 (SD 0.006). This is expected since medicine patients usually have treatment plans that are less predictable than those of surgical inpatients, and are generally more challenging to discharge.

The AUC-ROC of the medicine-minus model was 0.05 lower than the AUC-ROC

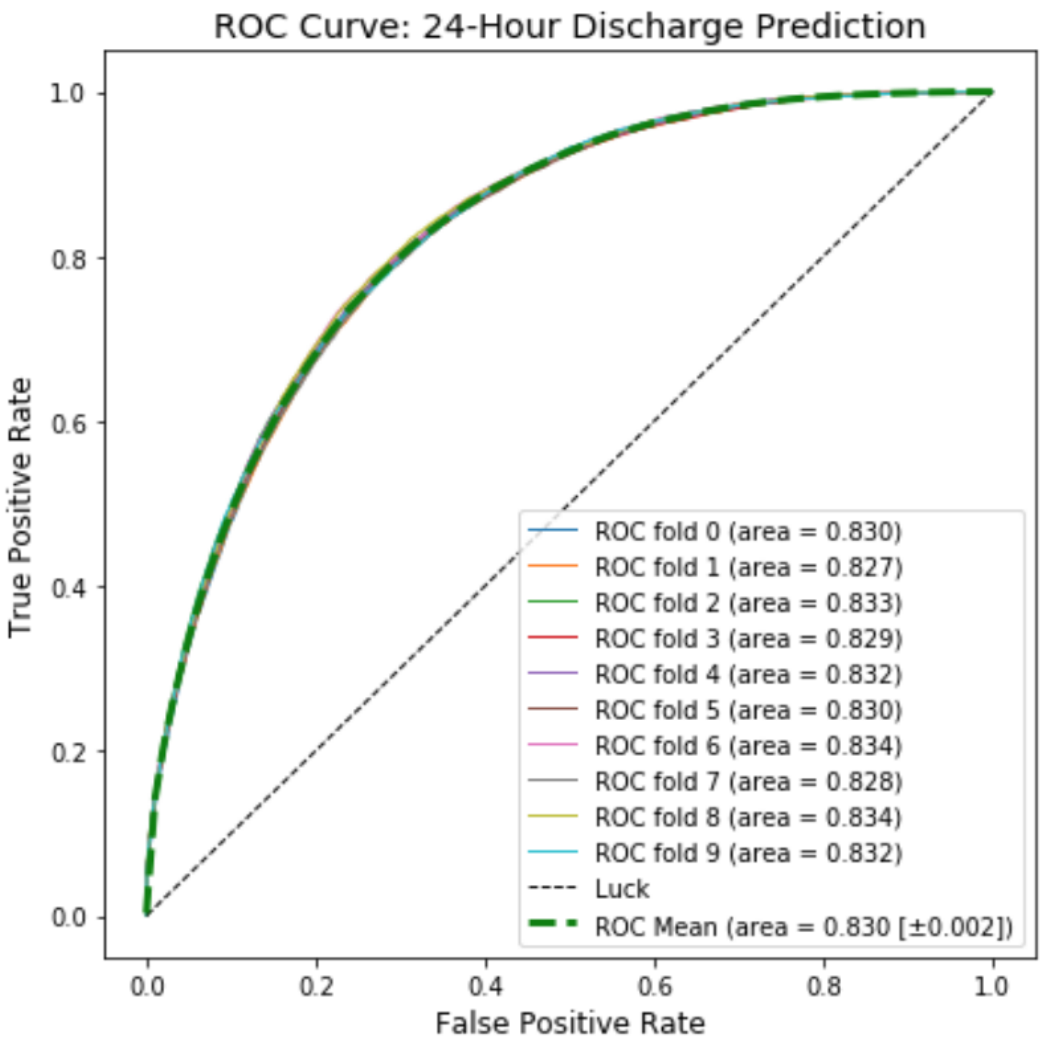

Figure 5-1: The ROC Curve for the 24-hour medicine discharge prediction model. The ROC was computed using 10-fold cross validation and inpatient data from May 2016 until September 2018.

Patient Service	Precision (24 hours)	Precision (48 hours)	Recall
Medicine	43.8%	64.8%	38.2%
Medicine (-)	35.9%	56.2%	33.5%
Cardiology	58.3%	74.3%	63.5%
Cardiology (-)	53.8%	70.0%	57.1%
Oncology	44.9%	65.6%	32.5%
Oncology (-)	34.7%	53.0%	27.3%
Neurology	45.4%	67.5%	48.5%
Neurology (-)	37.7%	58.4%	40.5%
All Medicine Services	47.6%	67.4%	43.8%
All Medicine Services (-)	40.2%	59.3%	38.3%
Surgery	62.1%	83.1%	62.5%

Table 5.1: The precision (24 hours, 48 hours) and recall metrics for the medicine model, the medicine-minus model (marked with '-'), and Zanger's enhanced surgical model.

of the medicine model (Figure 5-2). This suggests that adding additional patient data that was not in Zanger's model, using the trigger-resolution framework, and training the model using an LSTM-based architecture improved the discharge prediction algorithm.

5.2 Precision and Recall

The medicine discharge prediction model achieved 47.6% precision (24 hours), 67.4% precision (48 hours), and 43.8% recall using a threshold of 0.31. The decision threshold of 0.31 was chosen since it minimized the absolute difference between precision and recall (Figure 5-3). The precision and recall metrics were also evaluated for each patient service (medicine, cardiology, oncology, neurology) that cared for medicine inpatients for both the medicine and medicine-minus model. The results are compared with Zanger's enhanced surgical model's in Table 5.1.

Note that the oncology patient service has the lowest recall metric $\left(\frac{TP}{TP+FN}\right)$ of the four patient services. This is probably because oncology patients typically leave

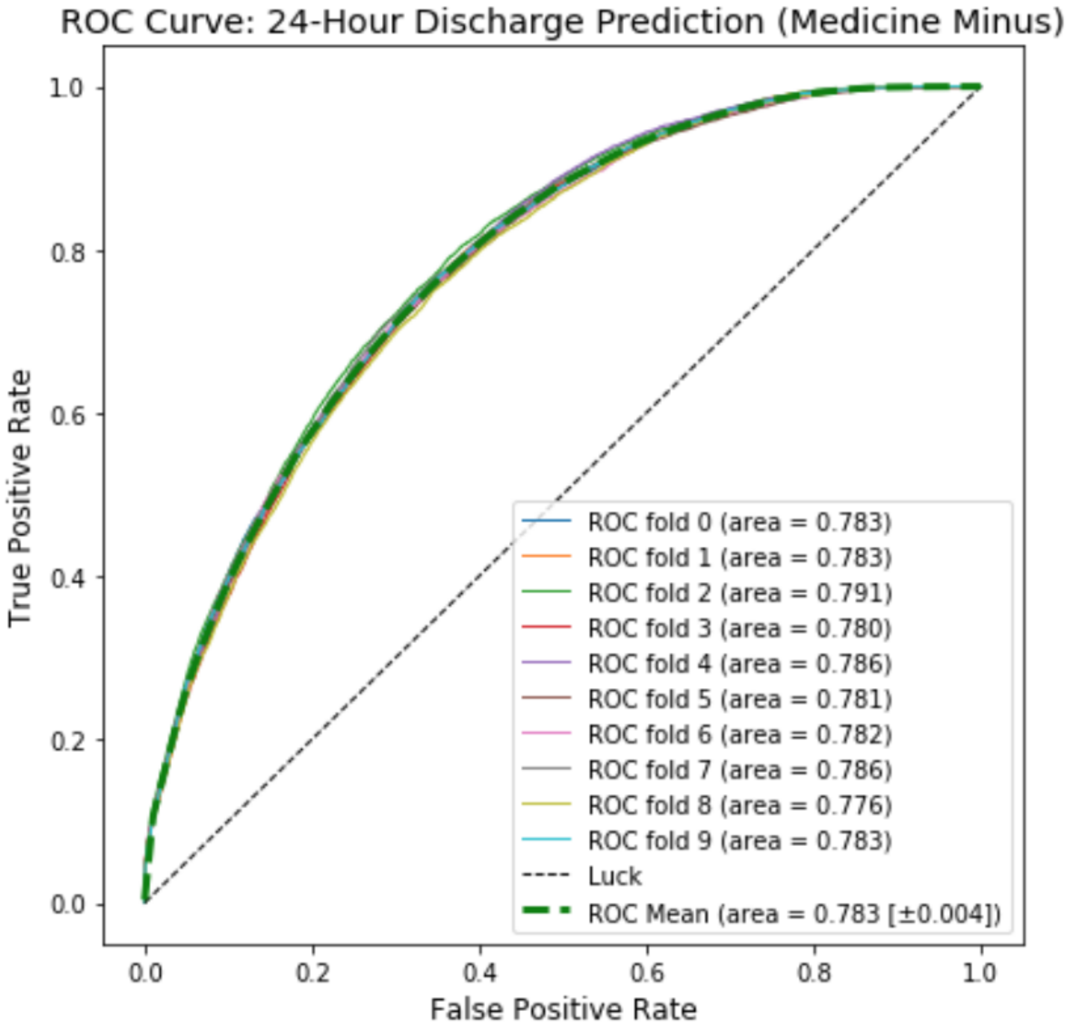

Figure 5-2: The ROC Curve for the 24-hour medicine-minus discharge prediction model. The ROC was computed using 10-fold cross validation and inpatient data from May 2016 until September 2018.

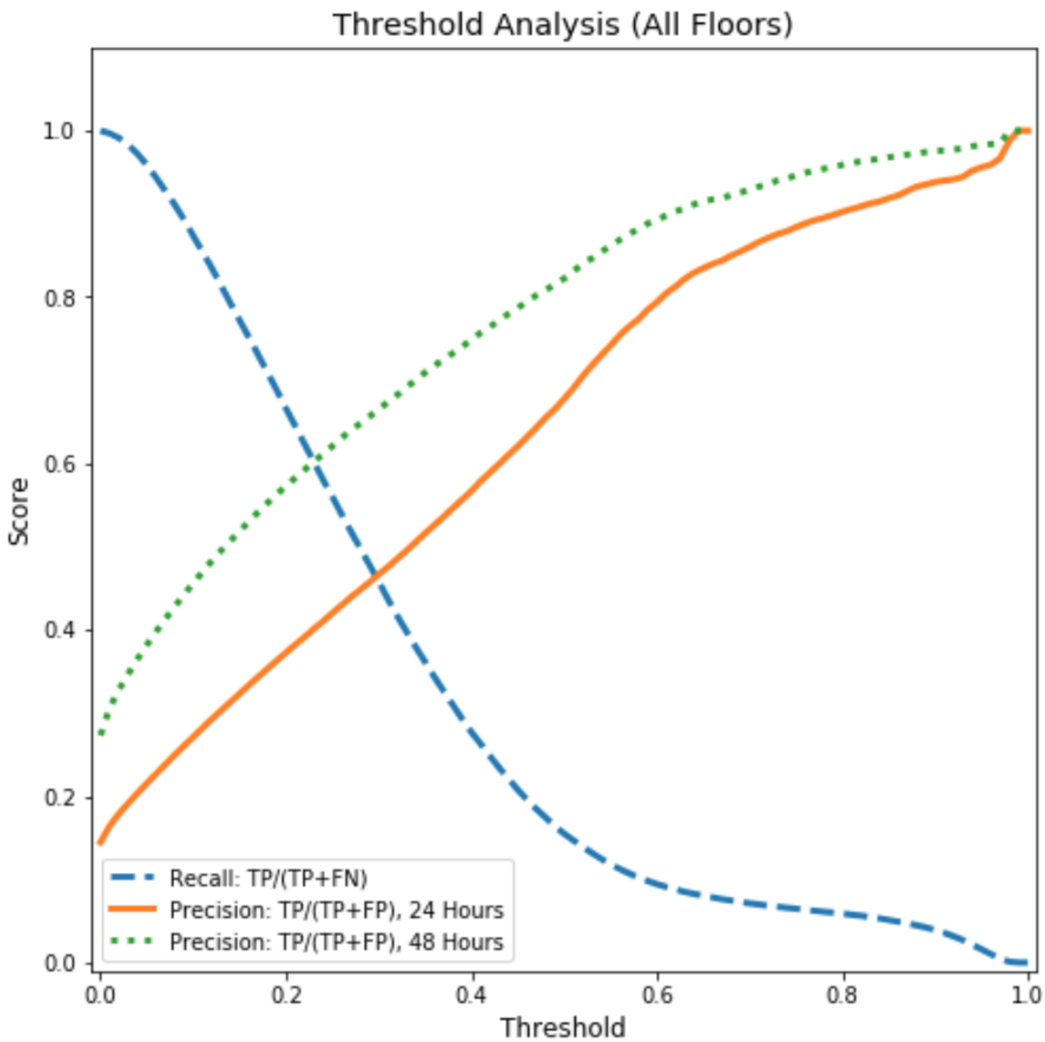

Figure 5-3: The recall and precision values for various thresholds in the medicine model. The threshold that minimizes the absolute difference between recall and precision is 0.31.

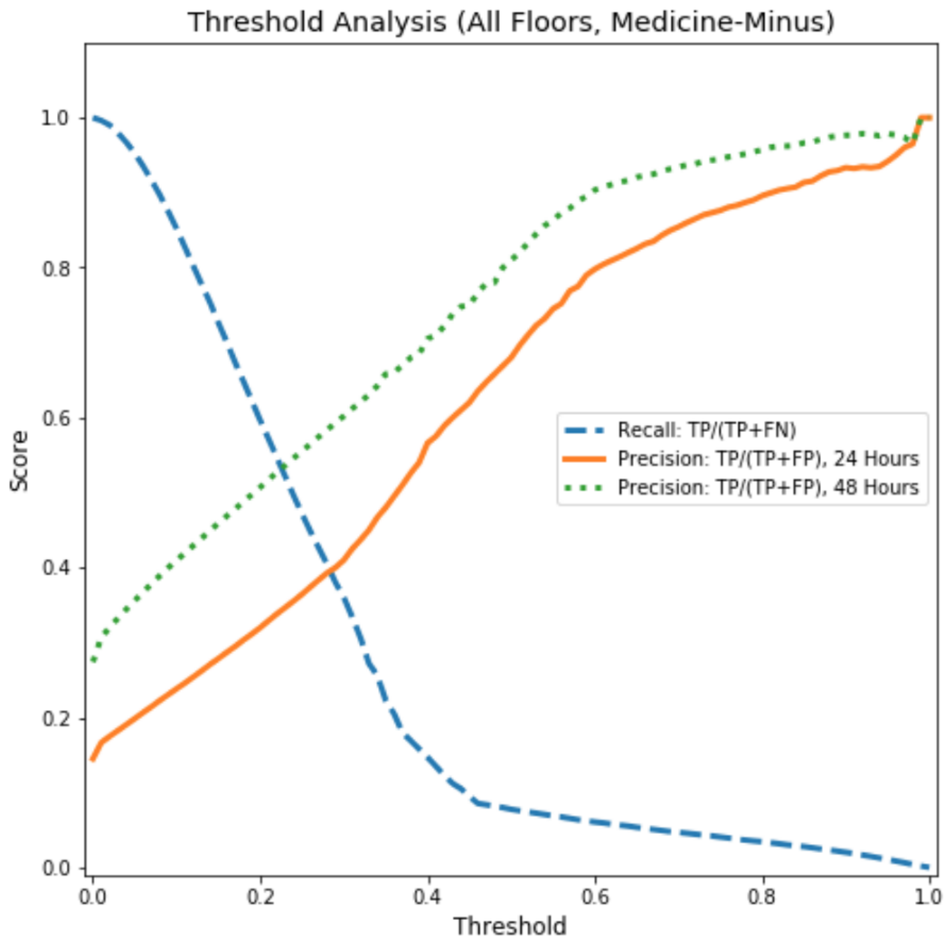

Figure 5-4: The recall and precision values for various thresholds in the medicine-minus model. The threshold that minimizes the absolute difference between recall and precision is 0.29.

the hospital sick, which makes it difficult to detect when these patients are ready for discharge. This results in a high number of false negatives and a low recall metric.

Also note that the precision and recall metrics of the medicine-minus model are lower than the medicine model's by approximately 7.5%, and 5.5% respectively (Figure 5-4, Table 5.1). This suggests that (1) adding patient data that was not in Zanger's model; (2) generating features with a the trigger-resolution framework; and (3) training the discharge prediction model using an LSTM-based architecture improve the medicine discharge prediction algorithm.

In addition, the medicine model's precision (24 hours and 48 hours) is approximately 15% less than Zanger's enhanced surgical model's precision (24 hours and 48 hours), and the medicine model's recall is roughly 20% less than Zanger's enhanced surgical model's recall (Table 5.1). This is most likely because the medicine population is much more clinically diverse than the surgical population, and it is therefore more difficult to identify medicine inpatient discharges correctly.

Notice however that the medicine model's precision (24 hours) and recall results for the cardiology service are only 5% lower than Zanger's enhanced surgical model's. This is most likely because cardiology inpatients have treatment plans that are more standardized and similar to surgical inpatients, and medical conditions that are easier to diagnose.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 6

Recommendations and Conclusions

This section provides recommendations for improving the bed-management process at MGH using the results of the discharge prediction model (Section 6.1). Areas for improvement and opportunities for further study are also discussed (Section 6.2).

6.1 Recommendations

Including the discharge prediction tool as the part of the bed-management process at MGH has the potential to impact clinicians, the Admitting Department, and the hospital's practices.

1. Care teams

   Currently clinical teams at the hospital do not have a data-driven way to reliably identify medicine patients who will get discharged from the hospital. This results in delays in discharge process, since patients who are approaching readiness for discharge are often only identified late in the day. Providing the discharge prediction tool to clinicians would allow for earlier identification of patients who are ready for discharge and for better prioritization methods to discharge patients.

2. Admitting Department

   Bed managers generally become aware of prospective discharges by receiving self-reports from medicine floors. This type of communication is manual and subjective, and prevents bed managers and the hospital from proactively monitoring the discharge process. Using the discharge prediction tool would provide bed managers and the Capacity Coordination Center with an approximation of the number of available beds and a methodical way to identify medicine inpatient discharges.

3. <u>Hospital Practices</u>

   The discharge prediction tool has the capability not only to provide a centralized approach to manage discharges, but also to pinpoint the bottlenecks that cause delays in patient discharges. An in-depth analysis of the barriers-to-discharge could highlight opportunities to reduce length-of-stay throughout the hospital and to ameliorate overcrowding challenges such as Code Help and Code Disaster in the Emergency Department at MGH.

6.2 Areas for Improvement and Opportunities for Further Study

In order to identify areas of improvement for the model, false positives and false negatives of the medicine discharge prediction model were analyzed. The areas of improvement that were identified from this analysis are listed below.

- <u>Adding Clinical Barriers</u>

  The oncology patient service had the lowest recall metric $\left(\frac{TP}{TP+FN}\right)$ of the four medicine patient services. In order to better detect when these patients are ready for discharge, additional clinical features for oncology patients should be added to the medicine model (e.g. monitoring chemotherapy treatments, detecting future chemotherapy treatments).

  In addition, the comments written after microbiology tests and imaging results are not currently captured by the model. Processing these notes and including them as features in the model may enhance the discharge prediction tool.

- Changing the Method of Predicting Discharge Based on Score

  Using a decision threshold may not be the best method to determine whether or a not a patient will get discharged within 24 hours. Instead, checking whether a patient's prediction score has increased by a certain amount within 24 hours or using a combination of this suggested method with a decision threshold may give better results.

- Adding Milestones and Barriers for Long Length-of-Stay Patients

  If a medicine patient has stayed in the hospital for longer than approximately a week, it becomes extremely challenging to determine when the patient will leave the hospital. These patients usually have exceptional medical or psychosocial conditions that make it very difficult to predict discharge. Identifying and adding clinical milestones and barriers-to-discharge that are particularly relevant for long length-of-stay patients may improve the discharge prediction model.

The medicine inpatient discharge prediction tool developed in this thesis also provides several opportunities for further study.

1. Identifying Significant Barriers-to-Discharge

   A trigger-resolution framework was used in order to process the raw patient data into features. This framework can be used to identify crucial clinical milestones that must happen before discharge and significant barriers-to-discharge that must be resolved prior to discharging a medicine inpatient. This would allow clinicians to focus on reaching those milestones and addressing those barriers in order to advance the discharge process.

2. Forecasting Other Hospital Events

   The discharge prediction framework can be refined to forecast additional hospital events, such as (1) patient readmissions; (2) in-hospital mortality; (3) length-of-stay; (4) capacity crises in the Emergency Department; and (5) patient readiness for discharge within 48 hours. This can be done by changing the model's objective function.

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix A

Overview of the Prediction Tool's Framework

This section presents an overview of the steps necessary to train the medicine inpatient discharge prediction tool.

Steps to Train the Model

Training the discharge prediction tool involves a five-step process (Figure A-1):

1. Filter Patients

   A list of medicine patients that meet the following conditions is used to train the discharge prediction tool:

   - The patient was admitted to MGH after May 1, 2016. From this date and onwards, clinical data from EPIC became more reliable.
   - The patient was discharged before a date that can be specified by the user.
   - The patient was on a surgical, medicine, or mixed floor at least once during the hospital encounter.
   - The patient was taken care of by the medicine, cardiology, oncology or neurology patient service while on a surgical, mixed, or medicine floor.
   - The patient is not an expired patient.

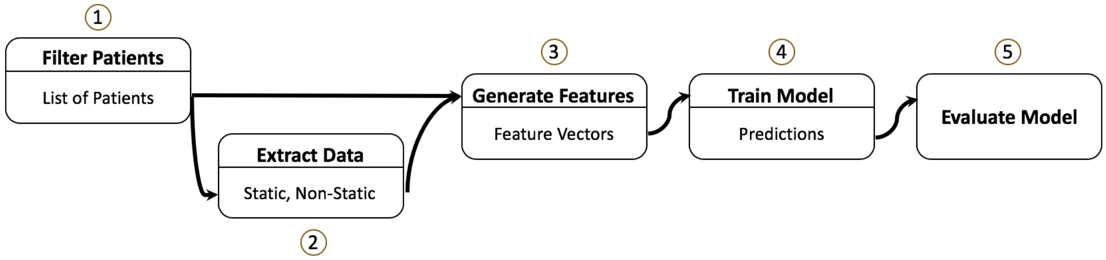

Figure A-1: This diagram illustrates the steps necessary to train the discharge prediction model.

2. Extract Data

   The static data (admission source, demographic information, diagnosis, location history, and patient service) and dynamic data (4Next referrals and acceptances, calendrical information, flowsheets, labs, LDAs, medications, orders, and notes) of the patients found in Step 1 are extracted from EPIC.

3. Generate Features

   Feature vectors are generated using the list of patients found in Step 1 as well as the patient data extracted in Step 2. This involves processing the raw data, including using the trigger-resolution framework or a sentiment analysis classifier.

4. Train Model

   The feature vectors generated in Step 3 along with a multi-input, fully-connected layer and LSTM neural-network based classifier are used to train the discharge prediction tool.

5. Evaluate Model

   The model's performance is evaluated using the predictions from Step 4. In particular, the AUC-ROC, precision (24 hours and 48 hours), and recall are calculated.

Appendix B

Data Sources Overview

This appendix presents the Partners Enterprise Data Warehouse (EDW) tables used by the medicine inpatient discharge prediction model. EDW contains approximately 26,000 data objects from multiple sources, including EPIC, financial systems, and a variety of external sources. Currently, EDW has nearly 1,350 enabled users with roughly 7.4 million database queries since its inception in 2014 [45].

EDW tables were used in this thesis to extract patient data recorded in MGH's electronic medical record system, EPIC. The following information is provided for each EDW table that was used in the model: (1) the field the table was joined on; (2) the columns that were extracted; and (3) the conditions that were set.

4Next Referrals and Acceptances

<u>EDW Table</u>: [Epic].[Patient].[Identity_MGH]

- Join On: PatientID (list of relevant medicine patient IDs)
- Columns Extracted: PatientIdentityID (MRN)
- Conditions: IdentityTypeID = 67 (MGH)

<u>EDW Table</u>: [Clinical].[FourNext].[TransitionOfCareReferralSummary_MGH]

- Join on:
  [Epic].[Patient].[Identity_MGH].[PatientIdentityID] =

[Clinical].[FourNext].[TransitionOfCareReferralSummary_MGH].[PatientMRN]

- Columns Extracted: CaseID, VisitNBR, ProviderID, ProviderNM, ProviderTypeCD, AcuteDischargeDTS, AcuteReferralDTS, BedAcceptDTS, AcceptDTS, PlacedDTS, BedOfferDTS, ProviderRetractDTS, ProviderRetractReasonTXT, DeniedDTS, DeniedReasonTXT, CancelAcceptDTS, CancelAcceptReasonTXT, CancelPlacementDTS, CancelPlacementReasonTXT, NoPlacementRequiredDTS
- Conditions: AcuteReferralDTS >= "05-01-2016" and is not Null

Admission Source

EDW Table: [Epic].[Encounter].[PatientEncounterHospital_MGH]

- Join On: PatientEncounterID (list of relevant medicine CSNs)
- Columns Extracted: AdmitSourceCD, AdmitSourceDSC
- Conditions: AdmitSourceCD is not Null

Demographic Information

EDW Table: [Epic].[Patient].[Patient_MGH]

- Join On: PatientID (list of relevant medicine patient IDs)
- Columns Extracted: BirthDTS, CityNM, MaritalStatusCD, MaritalStatusDSC, SexCD, SexDSC, StateCD, StateDSC, ZipCD

Diagnosis

EDW Table: [Epic].[Clinical].[HospitalProblem_MGH]

- Join On: PatientEncounterID (list of relevant medicine CSNs)
- Columns Extracted: PrincipalProblemFLG

EDW Table: [Epic].[Patient].[ProblemList_MGH]

- Join On:
  [Epic].[Clinical].[HospitalProblem_MGH].[ProblemListID] =
  [Epic].[Patient].[ProblemList_MGH].[ProblemListID]
- Columns Extracted: DiagnosisDTS, DiagnosisID

EDW Table: [Epic].[Reference].[ICDDiagnosis]

- Join On:
  [Epic].[Patient].[ProblemList_MGH].[DiagnosisID] =
  [Epic].[Reference].[ICDDiagnosis].[DiagnosisID]
- Columns Extracted: DiagnosisNM, ParentDiagnosisID

Flowsheet Records

EDW Table: [Epic].[Clinical].[FlowsheetRecordLink_MGH]

- Join On: PatientID (list of relevant medicine patient IDs)

EDW Table: [Epic].[Clinical].[FlowsheetMeasure_MGH]

- Join On:
  [Epic].[Clinical].[FlowsheetRecordLink_MGH].[FlowsheetDataID] =
  [Epic].[Clinical].[FlowsheetMeasure_MGH].[FlowsheetDataID]
- Columns Extracted: MeasureTXT, RecordedDTS, FlowsheetMeasureID
- Conditions:
  - RecordedDTS >= "05-01-2016" and is not Null
  - MeasureTXT is not Null
  - FlowsheetMeasureID is in the following list:
    * Activity Score: 305140, 305560

∗ Ambulation Assistance: 15924, 318180, 1600100065
∗ Assistive Device: 305620
∗ Blood: 3040990000
∗ Blood Pressure: 5, 1492, 5142, 8634
∗ Braden Scale Score: 305180
∗ Cardiac Rhythm: 301320
∗ Case Management: 11218, 304080500, 304080505, 304080507, 304080525, 304080783
∗ Cognition: 301880
∗ Emesis Occurrence: 304350
∗ First Surface to Surface Device Used: 304480144, 3040400254
∗ Friction and Shear Score: 305170, 7090680
∗ Level of Assistance: 7060230, 7060350
∗ Level of Consciousness: 301860
∗ Mobility Score: 305150, 7090640
∗ Morse Fall Risk Score: 305110
∗ Nutrition Score: 305160, 7090690
∗ O2 Device: 28920, 301030, 7086000, 107301030, 3041900111
∗ O2 Flow Rate: 15707, 250026, 7086010
∗ Orientation Level: 301870, 7083010
∗ OT Discharge Recommendation: 3040402126
∗ P.O.: 51
∗ Pain Score: 18, 20421, 3040104280
∗ Passing Flatus: 11049
∗ Physical Therapy: 3040402125, 3040401237
∗ Pulse: 8, 112892, 7085920, 1040129597, 1180102006, 3043200102
∗ Resp: 9, 7085930, 1180102005

    * Sensory Perceptions: 305120
    * Sit to Stand Level of Assistance: 318130
    * Speech-Language Pathology: 304100427, 304100528, 304100626, 304100629, 304250215, 304250322, 304250326
    * Soft restraint wrist: 300030, 300031, 300122, 300123
    * Speech: 301890
    * SpO2: 10, 115, 8572, 112894, 1601001, 30419010, 1040129596, 1180102007, 1180000005
    * Stand to Sit Device or Assistance: 304480143, 3040400527, 3040400529
    * Stool Occurrence or Amount: 340280, 304340
    * Supine to Sit Level of Assistance: 318110
    * Temp: 6, 891
    * Urine: 61, 304330

Isolation

EDW Table: [Epic].[Encounter].[HospitalIsolation_MGH]

- Join On: PatientEncounterID (list of relevant medicine CSNs)
- Columns Extracted: IsolationCD, IsolationDSC, IsolationAddedDTS, Isolation-RemovedDTS, IsolationCommentTXT
- Conditions: IsolationAddedDTS >= "05-01-2016" and is not Null

EDW Table: [Epic].[Orders].[Procedure_MGH]

- Join On: [Epic].[Encounter].[HospitalIsolation_MGH].[IsolationOrderID] = [Epic].[Orders].[Procedure_MGH].[OrderProcedureID]
- Columns Extracted: OrderProcedureID, ReasonForCancellationCD, Reason-ForCancellationDSC

Labs

EDW Table: [Epic].[Orders].[Result_MGH]

- Join On: PatientEncounterID (list of relevant medicine CSNs)
- Columns Extracted: ResultDTS, ResultValueNBR, ComponentID
- Conditions:
  - ResultDTS >= "05-01-2016" and is not Null
  - ResultValueNBR is not Null or equal to 9999999 (an error value)
  - ComponentNM is in the following list: Alkaline Phosphatase (ALP), Alanine Aminotransferase (ALT), Aspartate Aminotransferase (AST), Bilirubin, Creatinine (CRE), Glucose, Hemoglobin, International Normalization Ratio (INR), Potassium, Lactate/Lactic Acid, Sodium, N-terminal (NT) Brain Natriuretic Peptide (BNP), Troponin, White Blood Count

Location History

EDW Table: [Epic].[Encounter].[PatientADTDepartment_MGH]

- Join On: PatientEncounterID (list of relevant medicine CSNs)
- Columns Extracted: InDTS, OutDTS, ADTDepartmentID, ADTDepartmentNM

LDAs

EDW Table: [Epic].[Clinical].[LineDrainAirwayNoAddSingle_MGH]

- Join On: PatientEncounterID (list of relevant medicine CSNs)
- Columns Extracted: PlacementDTS, RemovalDTS, LineDrainDSC, SiteCD, PropertiesDisplayDSC, FlowsheetMeasureID
- Conditions:
  - PlacementDTS >= "05-01-2016" and is not Null

- FlowsheetMeasureID is in the following list:
  - Biliary Drain: 304740
  - Chest Tube: 16137, 3040102641, 3040102664
  - EFMS: 3040200491
  - Epidural Catheter: 7080610
  - GI Tube: 3040200057
  - Negative Pressure Wound Therapy: 3040200649, 304730
  - Open Surgical Wound: 3040200502
  - Penrose Drain: 304720
  - PICC: 304620
  - SuctionDrain: 304690
  - Suprapubic Catheter: 7085670
  - UrinaryCatheter: 7085610, 7085730

Medications

EDW Table: [Epic].[Orders].[Medication_MGH]

- Join On: PatientEncounterID (list of relevant medicine CSNs)
- Columns Extracted: MedicationID, MedicationDSC
- Conditions:
  - MedicationDSC is in the following list:
    - Antiemetic: Ondansetron, Zofran
    - Antipsychotic: Haldol, Haloperidol
    - Colonoscopy Prep: Golytely, Nulytely, Polyethylene Glycol
    - Diuretic: Bumetanide, Chlorothiazide, Furosemide, Torsemide
    - Heparin Drip: Heparin

    * IV Antibiotics: Ceftazidime, Ceftriaxone, Imipenem-Cilastatin, Linezolid, Piperacillin, Piperacillin-Tazobactam, Vancomycim, Zosyn
    * IV Fluid: Lactated Ringers
    * IV Narcotic : Fentanyl, Hydromorphone, Morphine
    * Nebulizers: Albuterol, Ipratropium-Albuterol, Levalbuterol
    * Nitroglycerin: Nitroglycerin
    * Phenobarbital: Belladonna-Phenobarbital, Phenobarbital
    * Steroids: Methylprednisolone, Prednisone
    * Tamsulosin: Tamsulosin
    * TPN: TPN, PN
    * Tube feeds: Glucerna, Jevity, Nepro, Osmolite
    * Warfarin: Coumadin, Jantoven, Warfarin

<u>EDW Table</u>: [Epic].[Clinical].[AdministeredMedication_MGH]

- Join On:
  [Epic].[Orders].[Medication_MGH].[OrderID] =
  [Epic].[Clinical].[AdministeredMedication_MGH].[OrderID]
- Columns Extracted: MedicationTakenDTS
- Conditions:
  - MedicationTakenDTS >= "05-01-2016" and is not Null
  - MARActionCD = 1 (Given)

Orders

The following EDW tables are used to extract consults, diet, heart procedures and imaging, imaging, laboratory orders, and microbiology data.

EDW Table: [Epic].[Orders].[Procedure_MGH]

- Join On: PatientID (list of relevant medicine patient IDs)
- Columns Extracted: OrderProcedureID, ProcedureID, ProcedureDSC, OrderDisplayNM, OrderDTS, ResultDTS, AbnormalFLG, LabStatusCD, LabStatusDSC, OrderStatusCD, OrderStatusDSC, ReasonForCancellationCD, ReasonForCancellationDSC
- Conditions: OrderDTS >= "05-01-2016" and is not Null

EDW Table: [Epic].[Orders].[ParentOrder_MGH]

- Join On:
  [Epic].[Orders].[Procedure_MGH].[OrderProcedureID] =
  [Epic].[Orders].[ParentOrder_MGH].[OrderID]
- Columns Extracted: ParentOrderID

CONSULTS

Finding when a consult order was placed: use the tables listed above under Orders.

- Additional Condition: ProcedureID is in the following list:
  - IP consult to Addiction Services: 165695
  - IP consult to Allergy: 387
  - IP consult to Anesthesiology: 389, 88576, 3040500
  - IP consult to Cardiology: 403, 165698, 165699, 165700, 165701, 165702
  - IP consult to Dentist: 433
  - IP consult to Dermatology:435
  - IP consult to Endocrinology: 437, 490666
  - IP consult to Gastroenterology: 453, 165704
  - IP consult to General Surgery: 325

- IP consult to Gerontology: 427
- IP consult to Gynecology: 409, 165713
- IP consult to Hematology: 431
- IP consult to Infectious Diseases: 327
- IP consult to Internal Medicine: 441
- IP consult to Interventional Radiology: 407, 516935
- IP consult to Nephrology: 333, 490716
- IP consult to Neurology: 335, 165710, 165711
- IP consult to Neurosurgery: 337
- IP consult to Nutrition Services: 385
- IP consult to Oncology: 357, 411, 429, 165678, 165680, 490663, 490668, 490669, 490670, 490671, 490672, 490673, 490674, 490676, 490684, 490685, 490688, 490689, 490690, 490691, 490693, 490695, 490698, 490699
- IP consult to Ophthalmology: 413, 3040000413
- IP consult to Otolaryngology: 165714, 30400165714
- IP consult to Palliative Care: 371, 399
- IP consult to Pain Management: 381, 165715
- IP consult to Physical Medicine Rehab: 351
- IP consult to Podiatry: 417,
- IP consult to Psychology: 355, 165669
- IP consult to Psychiatry: 423, 165668, 165670, 3041500, 3041505, 304000050190, 304000050192
- IP consult to Pulmonology: 425, 496746, 496747, 509831, 490682, 490683
- IP consult to Rheumatology: 405
- IP consult to Social Work: 447, 304099890, 304099891

- IP consult to Smoking Cessation Program: 7916
- IP consult to Surgery: 339, 369, 393, 397, 421, 457, 165679, 165683, 165685, 165694, 165697, 165705
- IP consult to Urology: 395

Finding the time when a consult note was written: use the Notes tables described below.

Diet

- Additional Conditions: OrderDisplayNM is not Null, OrderTypeCD = 9 (Diet)

Heart Procedures and Imaging

- Additional Conditions: OrderTypeCD is in the following list
  - Electrophysiology (EP) Lab: 27
  - Cardiac Catheterization (Cath) Lab: 58, 1007
  - Echocardiogram: 24, 28, 29

Imaging

- Additional Conditions: OrderTypeCD = 5 (Imaging)

Laboratory Orders

- Additional Conditions:
  - ProcedureNM is in the following list: Prepare Red Blood Count (RBC), Prepare Plasma, Prepare Cryoprecipitate, Prepare Platelets

Microbiology

- Additional Conditions:
  - ResultDTS is not Null
  - OrderTypeCD = 3 (Microbiology)
  - ProcedureNM is in the following list: Anaerobic, Blood, Cerebrospinal Fluid (CSF), Fluid (not CSF), Fluid Culture Smear, Fungal, Group B Streptococcal (GBS), Genital, Mycobacterial, Sputum, Stool, Throat, Tissue & Tissue Culture Smear, Urine, Wound & Wound Culture Smear

Patient Service

EDW Table: [Epic].[Encounter].[ADT_MGH]

- Join On: PatientEncounterID (list of relevant medicine CSNs)
- Columns Extracted: PatientServiceDTS, EffectiveDTS
- Conditions: EffectiveDTS >= "05-01-2016" and is not Null

Notes

Extracting Text:
EDW Table: [Epic].[Clinical].[Note_MGH]

- Join On: PatientID (list of relevant medicine patient IDs)
- Columns Extracted: NoteID, CreateDTS
- Conditions: CreateDTS >= "05-01-2016" and is not Null

EDW Table: [Epic].[Clinical].[NoteText_MGH]

- Join On:
  [Epic].[Clinical].[Note_MGH].[NoteID] =
  [Epic].[Clinical].[NoteText_MGH].[NoteID]

- Columns Extracted: NoteCSNID, ContactDateRealNBR, LineNBR, NoteTXT
- Conditions: NoteTXT is not Null

EDW Table: [Epic].[Clinical].[NoteEncounterInformation_MGH]

- Join On:
  [Epic].[Clinical].[NoteText_MGH].[NoteCSNID] =
  [Epic].[Clinical].[NoteEncounterInformation_MGH] .[CSNID]
- Columns Extracted: AuthorServiceCD, AuthorServiceDSC
- Conditions:
  - AuthorServiceDSC is in the following list: Addiction, Cardiology, Case Management, Hospice, Hospital Medicine, Medicine, Neurology, Nursing, Occupational Therapy, Oncology, Palliative Care, Physical Therapy, Psychiatry, Psychology, Smoking Cessation, Social Work

Smart Text:

EDW Table: [Epic].[Clinical].[NoteSmartTextID_MGH]

- Join On: Note CSN ID (list of relevant note identifiers)
- Columns Extracted: SmartTextID
- Conditions: SmartTextID is not Null

Smart Phrase:

EDW Table: [Epic].[Clinical].[NoteSmartPhraseID_MGH]

- Join On: Note CSN ID (list of relevant note identifiers)
- Columns Extracted: SmartPhraseID
- Conditions: SmartPhraseID is not Null

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix C

Sentiment Analysis Classifier Overview

This appendix presents an overview of the sentiment analysis classifier that was used to categorize dispo phrases. Recall that a dispo phrase is a group of words written in a physician, nursing, physical therapy, or social worker note that begins with the character sequence "Dispo:" or "Disposition:". Since this phrase usually signifies a patient's progress towards discharge, the goal of building the sentiment analysis classifier is to (1) determine whether a dispo phrase suggests that a patient will get discharged within 24 hours; and to (2) incorporate this information in the medicine inpatient discharge prediction model.

The headings below describe the extraction process of dispo phrases from unstructured notes, the formation of the sentiment analysis classifier, and the incorporation of the classifier's output into the medicine inpatient discharge prediction tool.

Extracting Dispo Phrases from Notes

In order to extract a dispo phrase from a physician, nursing, physical therapy, or social worker note, the dispo phrase's starting location and ending location needs to be determined. Locating where a dispo phrase begins is straightforward, since a dispo phrase will always start with the character sequence "Dispo:" or "Disposition:". However, determining where a dispo phrase ends is more challenging. This is because there is no singular condition that is able to detect where a dispo phrase terminates.

By reading over 500 dispo phrases, a list of conditions was developed that usually identifies the ending location in a dispo phrase. Several of the conditions that were assembled include:

- Character sequences that indicate that a new section in the text has begun, such as “Addendum:”, “Billing:”, “Code:”, “Code Status:”, “Contact Info:”, “Diet:”, “Hospital Bundle:”, “Medications & Allergies:”, “Overnight Events:”.

- Character sequences that identify the signature of the note’s author, such as “Signed:”, “MD”, or “CNP”[1]. Since an author’s signature is usually located at the end of a physician, nurse, physical therapist, or social worker note, a dispo phrase must terminate before the signature.

- Character sequences that signal the end of a paragraph, such as more than ten white-spaces in a row, “xxx”, “===”, “- - -”, “___”, and “***”.

The reliability of the developed conditions was checked by applying the conditions to 100 randomly selected physician, nursing, physical therapy, or social worker notes that each contained a dispo phrase. These processed notes were then manually reviewed in order to check whether or not the ending locations of the dispo phrases were accurately found. This process was repeated ten times. The average accuracy was 90.7% with a standard deviation of 3.9%.

After preprocessing the notes to find dispo phrases, the character sequence “Dispo:” or “Disposition:” was removed from the dispo phrases. This is because the character sequence “Dispo:” or “Disposition:” does not provide any information about whether a patient will get discharged within 24 hours.

Building a Sentiment Analysis Classifier

An embedding layer[2] and LSTM neural-network sentiment analysis classifier are used in order to determine the sentiment of dispo phrases (Figure C-1). Specifically, the

[1]CNP stands for certified nurse practitioner.
[2]An embedding layer maps a word to a dense-vector representation.

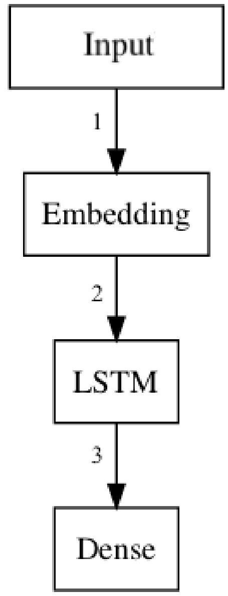

Figure C-1: A visualization of the dispo-phrase sentiment-analysis classifier's architecture. A dispo phrase is fed into an Embedding Layer (1). The resulting output is inputted into an LSTM that analyzes the sequence of words in the dispo phrase (2). Afterwards, the fully-connected layer (3) takes the output from the LSTM and transforms it into a score that indicates whether the dispo phrase contains positive sentiment (i.e., the patient will leave within 24 hours) or not.

input of the classifier is a dispo phrase, and the output of the classifier is score between 0 and 1. Scores near 0 indicate that a dispo phrase is *negative* (i.e., the patient will remain in the hospital for at least the next 24 hours), whereas scores near 1 signal that a dispo phrase is positive (i.e., a patient will get discharged within 24 hours from the hospital).

The embedding layer in the sentiment analysis classifier was used to map a dispo phrase to a word embedding (i.e., vector) that can be used as input into the model. The reason why an embedding layer was chosen to transform dispo phrases into vectors is because it is (1) computationally efficient (it can map a word from a large vocabulary into a dense vector); and because it can (2) transform words with comparable meanings to similar vector representations. An LSTM was used because it is capable of making predictions based on long sequences of words.

Training the sentiment analysis classifier required labelling a set of dispo phrases as positive or negative. Instead of doing this manually, dispo phrases written on the day-of-discharge or on the day-before-discharge were assumed to have positive sentiment[3], while the remaining dispo phrases were assumed to have negative sentiment. The sentiment analysis classifier was trained using over 200,000 dispo phrases from physician, nursing, physical therapy, and social worker notes. These notes were written for approximately 61,000 medicine inpatients from May 2016 and September 2018.

The sentiment analysis classifier's accuracy was 80.5% on an out-of-sample data set. One possible way to improve this accuracy is to manually label dispo phrases as positive or negative, instead of using the assumption outlined in the previous paragraph.

Incorporating the Sentiment Analysis Classifier's Output in the Medicine Discharge Prediction Model

The sentiment analysis classification's score of an inpatient's most-recent dispo phrase was used as a feature in the medicine inpatient discharge model. Notice that because this feature is dynamic, it is processed in the medicine model's LSTM input layer.

Several further opportunities for incorporating the classifier's output in the medicine discharge prediction model include:

- Checking whether a dispo phrase has changed at least once over the span of several days during a patient's hospitalization. If it has not changed, then the dispo phrase is most likely being copy-and-pasted and should not be used in the medicine inpatient prediction model.

- Instead of using the sentiment analysis classifier's score as a feature in the medicine discharge prediction model, the sentiment analysis classifier can be combined with the medicine discharge prediction model. More specifically, the

[3]Note that a dispo phrase on the day-of-discharge might sometimes actually be negative. For example "Dispo: more than 2-3 days" might be written in a note on the same day that a medicine inpatient was discharged from the hospital. This can occur because occasionally dispo phrases are copy-and-pasted from previous notes and are not updated.

medicine discharge prediction algorithm would have three inputs: static data, dynamic data, and text data (i.e., dispo phrases). A dense input layer would be used for the static data; an LSTM input layer would be used for the dynamic data; and an embedding input layer would be used for the text data.

THIS PAGE INTENTIONALLY LEFT BLANK

Bibliography

[1] Massachusetts General Hospital (MGH). Mass general history. https://www.massgeneral.org/museum/history/. Accessed: 2019-08-19.

[2] Massachusetts General Hospital (MGH). Overview. https://www.massgeneral.org/about/overview.aspx. Accessed: 2019-08-19.

[3] Massachusetts General Hospital (MGH). Facts & figures. https://www.massgeneral.org/news/facts/. Accessed: 2019-08-19.

[4] Partners HealthCare. About partners healthcare. https://www.partners.org/About/. Accessed: 2019-08-19.

[5] Massachusetts General Hospital (MGH). News release: Exeter health resources, wentworth-douglass hospital and massachusetts general hospital file an affiliation agreement to develop a proposed new seacoast network to ensure local access to high quality health care for the seacoast region, May 23, 2019. https://www.massgeneral.org/News/pressrelease.aspx?id=2407. Accessed: 2019-08-19.

[6] Jonas Hiltrop. Modeling neuroscience patient flow and inpatient bed management. Master's thesis, Massachusetts Institute of Technology, 2014.

[7] Sean McNichols. Reducing intraday patient wait times through just-in-time bed assignment. Master's thesis, Massachusetts Institute of Technology, 2015.

[8] Elizabeth Ugarph. Patient flow optimization in the department of medicine at mgh. Master's thesis, Massachusetts Institute of Technology, 2017.

[9] Kyan C. Safavi, et al. Development and validation of a machine learning model to aid discharge processes for inpatient surgical care. *JAMA Network Open*, 2(12):e1917221–e1917221, 2019.

[10] Jonathan Zanger. Predicting surgical inpatients' discharges at massachusetts general hospital. Master's thesis, Massachusetts Institute of Technology, 2018.

[11] Massachusetts General Hospital (MGH). About the department of neurology. https://www.massgeneral.org/neurology/about/. Accessed: 2020-01-14.

[12] Massachusetts General Hospital (MGH). Department of neurology- historical overview. https://www.massgeneral.org/neurology/about/history. Accessed: 2020-01-14.

[13] Katheleen Conti. Mgh continues to grapple with patient overcrowding. https://www.bostonglobe.com/business/2016/03/01/mgh-continues-grapple-with-patient-overcrowding/FUl4h1kIrlawLfSvroiIxO/story.html. Accessed: 2018-10-09.

[14] Brian Zimmerman. Mass general forms task force to combat ed overcrowding. https://www.beckershospitalreview.com/patient-flow/mass-general-forms-task-force-to-combat-ed-overcrowding.html. Accessed: 2018-10-09.

[15] Christopher Gavin. Here's what we know about mass. general hospital's planned $1 billion project. https://www.boston.com/news/health/2019/01/30/mass-general-hospital-project. Accessed: 2019-11-29.

[16] Thomas J. Best , et al. Managing hospital inpatient bed capacity through partitioning care into focused wings. *Manufacturing & Service Operations Management*, 7(2):157–176, February 2015.

[17] Phillip Reese. As er wait times grow, more patients leave against medical advice. https://khn.org/news/as-er-wait-times-grow-more-patients-leave-against-medical-advice/. Accessed: 2020-01-14.

[18] Liz Kowalczyk. At the er, the stay can reach 8 hours. http://archive.boston.com/yourlife/health/other/articles/2007/03/25/attheerthestaycanreach8hours/. Accessed: 2018-11-18.

[19] Liz Kowalczyk. Hospitals strive to reduce waiting lines for beds. http://archive.boston.com/yourlife/health/diseases/articles/2005/03/02/hospitalsstrivetoreducewaitinglinesforbeds/. Accessed: 2018-11-19.

[20] Scott R. Levin, et al. Real-time forecasting of pediatric intensive care unit length of stay using computerized provider orders. *Critical Care Medicine*, 40(11):3058–3064, 2012.

[21] Sean Barnes, et al. Real-time prediction of inpatient length of stay for discharge prioritization. *Journal of the American Medical Informatics Association*, 23:e2–e10, 2015.

[22] David H. Gustafson. Length of stay: prediction and explanation. *Health Services Research*, 3(1):12–34, 1968.

[23] Walter E. Pofahl Steven Steven and Ronald J. Scorpio. Predicting hospital length of stay with neural networks. *In FLAIRS conference*, pages 333–337, 1998.

[24] Vincent Liu, et al. Length of stay predictions: improvements through the use of automated laboratory and comorbidity variables. *Medical Care*, 48(8):739–744, 2010.

[25] Vandana P. Janeja Ali Azari and Alex Mohseni. Predicting hospital length of stay (phlos): A multi-tiered data mining approach. *In 2012 IEEE 12th International Conference on Data Mining Workshops*, pages 17–24, 2012.

[26] April et al. Morton. A comparison of supervised machine learning techniques for predicting short-term in-hospital length of stay among diabetic patients. *In 2014 13th International Conference on Machine Learning and Applications, IEEE*, pages 428–431, 2014.

[27] Hrayr Harutyunyan, et al. Multitask learning and benchmarking with clinical time series data. *Scientific Data*, 6(1):96–113, 2010.

[28] Shahadate Rezvy Tahmina Zebin and Thierry J. Chaussalet. A deep learning approach for length of stay prediction in clinical settings from medical records. *In 2019 IEEE Conference on Computational Intelligence in Bioinformatics and Computational Biology (CIBCB), IEEE*, pages 1–5, 2019.

[29] Arnold W. Pratt and Milos G. Pacak. Automated processing of medical english. *In International Conference on Computational Linguistics Coling*, pages 1–23, 1969.

[30] Yiming Yang, et al. An example-based mapping method for text categorization and retrieval. *ACM Transactions on Information Systems (TOIS)*, 12(3):252–277, 1994.

[31] Serguei V. Pakhomov, et al. Prospective recruitment of patients with congestive heart failure using an ad-hoc binary classifier. *Jama psychiatry*, 38(2):145–153, 2005.

[32] Guergana K. Savova, et al. Mayo clinical text analysis and knowledge extraction system (ctakes): architecture, component evaluation and applications. *Journal of the American Medical Informatics Association*, 17(5):507–513, 2010.

[33] Thomas H. McCoy, et al. Improving prediction of suicide and accidental death after discharge from general hospitals with natural language processing. *Jama psychiatry*, 73(10):1064–1071, 2016.

[34] Mitali Chakravarty , Anirban Hom Choudhuri and Rajeev Uppal. Influence of admission source on the outcome of patients in an intensive care unit. *Indian journal of critical care medicine: peer-reviewed, official publication of Indian Society of Critical Care Medicine*, 21(4):213, 2017.

[35] James M. Watters, et al. Impaired recovery of strength in older patients after major abdominal surgery. *Annals of surgery*, 218(3):380, 1993.

[36] Mark D. Neuman and Rachel M. Werner. Marital status and postoperative functional recovery. *JAMA surgery*, 151(2):194–196, 2016.

[37] Carl van Walraven and Chaim M. Bell. Impact of socio-economic status on hospital length of stay following injury: a multicenter cohort study. *BMC health services research*, 15(1):285, 2015.

[38] Partners HealthCare. 4next. https://4next.partners.org/cleartrust/ct_logon.asp?CTAuthMode=BASIC&language=en&CT_ORIG_URL=https://4next.partners.org/LoginSuccess.aspx. Accessed: 2019-11-24.

[39] Carl van Walraven and Chaim M. Bell. Risk of death or readmission among people discharged from hospital on fridays. *Cmaj*, 166(13):1672–1673, 2002.

[40] Infection Prevention and You. What are transmission-based precautions? http://professionals.site.apic.org/what-are-transmission-precautions/. Accessed: 2019-12-09.

[41] Kim Tran, et al. The effect of hospital isolation precautions on patient outcomes and cost of care: a multi-site, retrospective, propensity score-matched cohort study. *Journal of general internal medicine*, 32(3):262–268, 2017.

[42] VNA Care. Transforming care to improve lives. https://vnacare.org/. Accessed: 2019-11-24.

[43] Scikit Learn. Simpleimputer. https://scikit-learn.org/stable/modules/generated/sklearn.impute.SimpleImputer.html. Accessed: 2019-11-24.

[44] Scikit Learn. Maxabsscaler. https://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.MaxAbsScaler.html. Accessed: 2019-11-24.

[45] Partners HealthCare. Data underpins every activity. https://populationhealth.partners.org/data-analytics/. Accessed: 2019-08-19.