Leveraging Clinical Data to Optimize Oxygen Delivery to the Preterm Infant

by

Mattie F. Wasiak

B.S. Computer Science, Massachusetts Institute of Technology, 2020

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

May 2020

© Massachusetts Institute of Technology 2020. All rights reserved.

Author .....
Department of Electrical Engineering and Computer Science
May 18, 2020

Certified by .....
Thomas Heldt
Associate Professor of Electrical and Biomedical Engineering
Thesis Supervisor

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

Leveraging Clinical Data to Optimize Oxygen Delivery to the Preterm Infant

by

Mattie F. Wasiak

Submitted to the Department of Electrical Engineering and Computer Science
on May 18, 2020, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

Determining an optimal target oxygen range has been a recurring debate in neonatal intensive care units (NICUs) as researchers and clinicians have taken note of the morbidity and mortality associated with high and low arterial oxygen saturations (SpO_2), respectively. In this thesis, we validated the real time data streams from the 48-bed Beth Israel Deaconness Medical Center NICU and leveraged these massive amounts of physiological data streams to analyze and identify clinical, demographic, physiological, and workflow factors that may affect the in-target oxygenation of premature infants.

We studied a cohort of 865 preterm infants (gestational age $<$ 37 weeks) that were in the NICU for more than 24 hours between January 1, 2018 and September 12, 2019; 412 of these patients were recorded on supplemental oxygen at some point during their stay. We had SpO_2 data coverage of 90.3% at the NICU level. We determined that patients spent 90.0% of their collective stay within their target SpO_2 range; 4.9% of their stay below their target SpO_2 range; and 5.1% of their stay above their target SpO_2 range. For patients on supplemental oxygen, we determined that these patients spent 24.0% of their time above their target SpO_2 range while on supplemental oxygen.

We investigated various factors' relationships with SpO_2; these factors included gestational age, birth weight, gender, supplemental oxygen, postmenstrual age, time of day, and bed space assignment. Overall, we determined that younger patients with lower birth weights spend the least time within their target SpO_2 range. Male patients also spent a greater % time above their target SpO_2 range while on supplemental oxygen than female patients.

Thesis Supervisor: Thomas Heldt
Title: Associate Professor of Electrical and Biomedical Engineering

Acknowledgments

I would like to thank my research supervisor, Professor Thomas Heldt, and our clinical collaborator, Doctor Wendy Timpson, for all of their support and guidance that made this research possible. I have learned so much over the past two years and am extremely grateful that I got the opportunity to contribute to this incredible work.

I would also like to thank James Lynch for all his help with the massive task of organizing and wrangling with the data provided by Data Warehouse Connect.

I would also like to thank my parents for always believing in me and being there for me. I really appreciate everything you've done for me and I wouldn't have been able to accomplish all of this without you.

Lastly, I would like to thank all of the amazing friends I met during my time at MIT; the people are truly what makes MIT special, and I will forever treasure all the great memories I made at this school.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

A significant part of keeping preterm neonates healthy in neonatal intensive care units (NICUs) is monitoring and titrating their oxygen saturation (SpO_2) in order to prevent mortality associated with low oxygen saturation (hypoxemia) and morbidity associated with high oxygen saturation (hyperoxemia). Keeping infants' oxygen saturations in a pre-defined target range for a high fraction of time is important to prevent life threatening complications for these tiniest of patients. Previous work has found that target oxygen ranges of 85% to 89% can lead to risk of death while target oxygen range of 91% to 95% can lead to heightened risk of retinopathy of prematurity (a potentially blinding eye disorder) [2].

Despite these known complications of sustained oxygenation outside target ranges, NICU staff can struggle to maintain infants' oxygenation saturations within the set target range due to continuously evolving pathophysiology, the intricacies of various respiratory support modalities and possibly complex workflow factors.

Our primary goal is to appropriately leverage clinical, demographic and workflow information as well as physiological monitoring data streams to identify factors that may place preterm infants at risk for hypoxemia and hyperoxemia. Due to the vulnerability and unpredictability of preterm infants to fluctuations in oxygen levels, it is important to identify these factors and adjust clinical management of care for each neonate accordingly to limit the amount of time spent outside the target oxygen range.

We leveraged an infrastructure called Data Warehouse Connect (DWC) from the NICU at Beth Israel Deaconess Medical Center (BIDMC) which provides massive amounts of clinical and physiological patient data from the Philips bedside monitors of the 48-bed NICU. After initially validating the data integrity against the clinical medical records, we processed the cleaned data to provide data analysis insights to the clinical staff that enhance the understanding of individual factors that influence each patient's oxygenation and determine key factors associated with maintaining appropriate SpO_2. Examples of analyses include how the percentage of time spent outside of the target SpO_2 range fluctuates based on various factors such as gestational age, bed assignment, supplemental oxygen, and time of day.

Uncovering the primary factors that drive out-of-target oxygenation can provide insights to clinical staff regarding which variables should be focused on when working to minimize the amount of time preterm infants spend outside their target SpO_2 range. The massive amounts of accessible SpO_2 data was leveraged by both researchers and clinical staff to determine which patients are most likely to experience hypoxemic and hyperoxemic episodes and associated morbidites; the data uncovered characteristics found to be more indicative of spending more time outside their target oxygen range. Understanding the causes of these episodes can help structure clinical workflow and proactively focus on patients deemed at higher risk.

Chapter 2

Background

Through this section, we summarize the history of oxygenation in the NICU and the importance of maintaining a target oxygen saturation range to prevent low arterial oxygen saturation levels (hypoxemia) and high arterial oxygen saturation levels (hyperoxemia). We also review literature that has studied the ideal target oxygen saturation and factors that are associated with patient's oxygenation. This section also provides background on the NICU involved in our study as well as our data source.

2.1 History of Oxygenation in the NICU

Over the past century, there has been debate surrounding the optimal target oxygen saturation range for preterm infants, and researchers have discussed the morbidity and mortality associated with high and low oxygen saturations respectively. Since the 1940s, delivery of supplemental oxygen has been an important and commonly used therapy in neonatal nurseries in order to increase oxygenation of neonates' tissues while preventing oxygen toxicity and stress [33]. When oxygen therapy was first introduced, it was widespread practice to provide unrestricted oxygen to neonates to prevent hypoxemia (low arterial oxygen saturation). However, in 1951, Kate Campbell and Mary Crosse suggested an association between this usage of unrestricted oxygen and severe retinopathy of prematurity (ROP), a severe eye disease caused by

abnormal development of retinal blood vessels [15]. After this discovery, NICUs began restricting supplemental oxygen in order to avoid hyperoxemia (high arterial oxygen saturation), even though it was estimated in small clinical trials that there were 16 additional deaths for every prevented blindness, indicating risk of mortality due to hypoxia [15, 22]. To avoid both ROP and increased mortality, over recent years, researchers have been working to find the optimal target oxygen range for neonates and address potential trade-offs in mortality and morbidity through randomized controlled clinical trials.

2.2 Target Oxygen Saturation Range

Recent studies have confirmed Campbell and Crosse's hypothesis and encouraged defining an ideal upper oxygen saturation limit in order to reduce the incidence of ROP. Since 2003, there have been five major trials connected through the Neonatal Oxygenation Prospective Meta-analysis (NeOProM) Collaboration exploring the outcomes of low oxygen saturation versus high oxygen saturation in extremely preterm infants: SUPPORT, COT, BOOST, BOOST II (UK) and BOOST II (Australia) [19]. These were randomized, double-blind, multicenter trials that studied a total of approximately 5000 infants born before 28 weeks gestational age and enrolled within 24 hours of birth [19]. The trials aimed to compare low target oxygen saturation ranges (85 to 89%) to high target oxygen saturation ranges (91 to 95%) by comparing morbidity and mortality rates [19]. Researchers discovered that mortality and the incidence of necrotizing enterocolitis were significantly higher in the low target group than the high target group; however, there was a significantly higher risk of ROP in the high target group compared to the low target group [19]. These results led to conclusions that clinicians should be cautious when treating neonates with the lower oxygen range due to increased mortality rates and led to hospitals accordingly increasing their target oxygen saturation range limits over the last 10 years despite the potential risk of ROP [32]. The conclusions of these trials reveal the ongoing debate of what the optimal target oxygen saturation range for preterm infants should

be and demonstrate potential consequences of a neonate's oxygen levels being below or above certain values. A key question for standard clinical practice is the degree to which a particular neonate can be reliably kept within the target oxygenation range.

2.3 Maintaining the Target Oxygen Saturation Range

We can see that it is vital to maintain a neonate's oxygenation in a target range for as long as possible to prevent mortality and morbidity; however, this task can be very difficult for neonatal clinical staff due to continuously evolving pathophysiology, various respiratory modalities, and complex workflow factors. In order to maintain oxygen saturation (SpO_2) within the target SpO_2 range, clinicians can titrate oxygen manually and also leverage pulse oximetry for continuous monitoring of SpO_2 in a non-invasive manner to ensure they are preventing the hypoxemic (defined by one study as SpO_2 below 80% for more than 10 seconds) and hyperoxemic (SpO_2 above 95% for more than 10 seconds) episodes associated with the aforementioned mortality and morbidity [31].

Very few studies have worked to determine how often clinical staff are able to maintain their hospital's target oxygen ranges in patients and found that patients were within their target range anywhere from 16 to 79% of the monitored time [14, 31]. A systematic review of sixteen of these studies included 574 infants and 2935 nurses and reached the overall conclusion that compliance in targeting SpO_2 was low and that it was especially difficult to maintain the SpO_2 below the upper limit when patients were on supplemental oxygen leading to many hyperoxemic events, which indicates that there is room for improvement [31]. The studies that were reviewed can be broken down into six studies analyzing percent time spent within the intended SpO_2 target range and potential factors affecting this compliance, six studies analyzing the impact of automated FiO_2 adjustment on in-target oxygenation, two studies analyzing nurses' awareness and compliance of SpO_2 limits and the associated alarms, and two studies analyzing how often SpO_2 alarm limits are correctly set.

We referenced the subset of seven of the reviewed sixteen studies that focused

on analyzing percent time spent within the intended SpO_2 range to gain additional insight of the difficulty of maintaining the target SpO_2 range [3, 8, 10, 11, 14, 18, 30, 31]. The results for these studies can be observed in Table 2.1; the potential factors investigated affecting compliance will be explored more in Section 2.4.

Table 2.1: Breakdown of studies analyzing percent time spent within the target SpO_2 range.

Study Title	**Authors**	**Study Cohort**	**SpO_2 Sampling Frequency**	**Target SpO_2 Range**	**Target SpO_2 Range Compliance**	**Investigated factors potentially affecting compliance**
Achieved Versus Intended Pulse Oximeter Saturation in Infants Born Less Than 28 Weeks' Gestation: The AVIOx Study	Hagadorn JI, Furey AM, et al.	84 extremely preterm (<28 weeks) infants on supplemental oxygen from 14 centers across 3 countries (mean GA: 26 weeks)	72 hours weekly for the first four weeks of a patient's life	Lower limit: 83% - 92% Higher limit: 92% - 98%	Within: 48% Below: 16% Above: 36%	Width of target SpO_2 range Chronological age Postmenstrual age Gestational age Birth weight Mechanical ventilator status Pulmonary acuity score SNAP-II score Time of day

Study Title	**Authors**	**Study Cohort**	**SpO_2 Sampling Frequency**	**Target SpO_2 Range**	**Target SpO_2 Range Compliance**	**Investigated factors potentially affecting compliance**
Pulse oximetry in very low birth weight infants: can oxygen saturation be maintained in the desired range?	Laptook AR, Salhab W, et al.	72 very low birth weight (500-1250 grams) preterm infants (< 37 weeks) on continuous supplemental oxygen Group 1: 23 infants (mean GA: 27 weeks) Group 2: 49 infants (mean GA: 26 weeks)	24 hours twice a month for 6 months	Group 1: 90% - 95% Group 2: 88%-94%	Group 1: Within: 57.7% Below: 26.9% Above: 15.4% Group 2: Within: 59.4% Below: 26.6% Above: 14.0%	Target SpO_2 range limits

Study Title	**Authors**	**Study Cohort**	**SpO_2 Sampling Frequency**	**Target SpO_2 Range**	**Target SpO_2 Range Compliance**	**Investigated factors potentially affecting compliance**
Nurse:patient ratio and achievement of oxygen saturation goals in premature infants	Sink DW, Hope SAE et al.	14 preterm (<29 weeks) infants on continuous supplemental oxygen (mean GA: 26.6 weeks)	Multiple monitoring periods (median of 60 periods per infant) of up to 6 hours	85%-92%	Within: 16.1% Below: 3.1% Above: 79.3%	Nurse:patient ratio Illness severity Postmenstrual age Respiratory support mode Inspired oxygen (FiO_2)

Study Title	**Authors**	**Study Cohort**	**SpO_2 Sampling Frequency**	**Target SpO_2 Range**	**Target SpO_2 Range Compliance**	**Investigated factors potentially affecting compliance**
An observational study to quantify manual adjustments of the inspired oxygen fraction in extremely low birth weight infants	Van der Eijk AC, Dankelman J et al.	12 extremely low birth weight (<1000 grams) preterm infants on supplemental oxygen (median GA: 26 2/7 weeks)	3 days continuously from the first 2 weeks of life	88%-94%	Within: 54% Below: 16% Above: 30%	FiO_2 adjustments

Study Title	**Authors**	**Study Cohort**	**SpO_2 Sampling Frequency**	**Target SpO_2 Range**	**Target SpO_2 Range Compliance**	**Investigated factors potentially affecting compliance**
Oxygen saturation targeting in preterm infants on continuous positive airway pressure	Lim K, Wheeler KI et al.	45 preterm (<37 weeks) infants in 2 NICUs on continuous positive airway pressure (CPAP) and supplemental oxygen (median GA: 30 weeks)	Multiple 24 hour periods	88%-92%	Within: 31% Below: 9% Above: 59%	CPAP

Study Title	**Authors**	**Study Cohort**	**SpO_2 Sampling Frequency**	**Target SpO_2 Range**	**Target SpO_2 Range Compliance**	**Investigated factors potentially affecting compliance**
Sojourn in excessively high oxygen saturation ranges in individual, very low-birthweight neonates	Arawiran J, Curry J, et al.	71 preterm infants (<31 weeks) Before education intervention: 41 infants (mean GA: 25 weeks) After education intervention: 30 infants (mean GA: 25 weeks)	Multiple 12 hours shifts	85%-92%	Before education intervention: Within: 44.5% Below: 18.6% Above: 36.9% After education intervention: Within: 40.4% Below: 17.7% Above: 41.9%	Mode of ventilation Nursing shifts Education

Study Title	**Authors**	**Study Cohort**	**SpO_2 Sampling Frequency**	**Target SpO_2 Range**	**Target SpO_2 Range Compliance**	**Investigated factors potentially affecting compliance**
Nurses' compliance with alarm limits for pulse oximetry: qualitative study	Armbruster J, Schmidt B et al.	20 extremely preterm (<28 weeks) infants on supplemental oxygen	First 3 days of life	88%-92%	Within: 68%-79%	Education Response time to alarms Patient:staff ratio Root cause analysis at the bedside Amount of priority given to control of oxygen therapy

Additionally, the six studies analyzing the impact of automated adjustment of the fraction of inspired oxygen (FiO_2) on the time within the SpO_2 target range did record the percent time spent within the target SpO_2 range in manual FiO_2 adjustment mode and automatic FiO_2 adjustment mode; however, these studies focused on attempting to improve the percent time spent within the target SpO_2 range through algorithms rather than determining root causes of patients leaving their target SpO_2 range. For manual FiO_2 adjustments, we acknowledged there may be some bias in time spent within the target SpO_2 range as the clinical staff was aware of the study and additionally, some studies designated a nurse to continuously monitor the patient which, while ideal, is unrealistic in practice. The results for these studies are detailed in Table 2.2 [1, 12, 23, 24, 25, 26, 31].

Table 2.2: Breakdown of studies analyzing percent time spent within the target SpO_2 range with manual FiO_2 adjustments and automatic FiO_2 adjustments.

Study Title	**Authors**	**Study Cohort**	**SpO_2 Sampling Frequency**	**Target SpO_2 Range**	**Target SpO_2 Range Compliance (Manual FiO_2 Adjustments)**	**Target SpO_2 Range Compliance (Automated FiO_2 Adjustments)**
Automated adjustment of inspired oxygen in preterm infants with frequent fluctuation in oxygenation: a pilot clinical trial	Claure N, D'Ugard C, and Bancalari E	16 preterm infants on supplemental oxygen (mean GA: 24.9 weeks)	Two 4 hours periods (one on manual FiO_2 adjustment by clinical personnel and one on automatic adjustment)	88%-95%	Within: 42% Below: 27% Above: 31%	Within: 58% Below: 33% Above: 9%

Study Title	**Authors**	**Study Cohort**	**SpO_2 Sampling Frequency**	**Target SpO_2 Range**	**Target SpO_2 Range Compliance (Manual FiO_2 Adjustments)**	**Target SpO_2 Range Compliance (Automated FiO_2 Adjustments)**
Closed-loop controlled inspired oxygen concentration for mechanically ventilated very low birth weight infants with frequent episodes of hypoxemia.	Claure N, Gerhardt T et al.	14 very low birth weight (500-1250 grams) preterm infants on supplemental oxygen (mean GA: 25 weeks)	Two 2 hours periods (one on manual FiO_2 adjustment by clinical personnel and one on automatic adjustment)	88% -96%	Within: 66.3% Below: 18.7% Above: 14.9%	Within: 74.9% Below: 16.5% Above: 9.9%

Study Title	**Authors**	**Study Cohort**	**SpO_2 Sampling Frequency**	**Target SpO_2 Range**	**Target SpO_2 Range Compliance (Manual FiO_2 Adjustments)**	**Target SpO_2 Range Compliance (Automated FiO_2 Adjustments)**
Multicenter crossover study of automated control of inspired oxygen in ventilated preterm infants.	Claure N, Bancalari E et al.	32 preterm infants on supplemental oxygen (median GA: 25 weeks)	Two 24 hour periods (one on manual FiO_2 adjustment by clinical personnel and one on automatic adjustment)	87%-93%	Within: 32% Below: 23% Above: 37%	Within: 40% Below: 32% Above: 21%

Study Title	**Authors**	**Study Cohort**	**SpO_2 Sampling Frequency**	**Target SpO_2 Range**	**Target SpO_2 Range Compliance (Manual FiO_2 Adjustments)**	**Target SpO_2 Range Compliance (Automated FiO_2 Adjustments)**
Automatic control of the inspired oxygen fraction in preterm infants: a randomized crossover trial	Urschitz MS, Horn W et al.	Validation trial: 12 preterm (<34 weeks) infants on supplemental oxygen (median GA: 24.5 weeks) Efficacy trial: 12 preterm infants on supplemental oxygen (median GA: 25.5 weeks)	5 periods of 90 minutes each over 1 day	87%-96%	Validation trial: Within: Routine manual control: 79.7% Optimal manual control: 85.8% Efficacy trial: Within: Routine manual control: 81.7% Optimal manual control: 91%	Validation trial: Within: Open-loop control: 85.1% Efficacy trial: Within: Closed-loop control: 90.5%

Study Title	**Authors**	**Study Cohort**	**SpO_2 Sampling Frequency**	**Target SpO_2 Range**	**Target SpO_2 Range Compliance (Manual FiO_2 Adjustments)**	**Target SpO_2 Range Compliance (Automated FiO_2 Adjustments)**
A randomised controlled trial of an automated oxygen delivery algorithm for preterm neonates receiving supplemental oxygen without mechanical ventilation.	Zapata J, Gomez JJ et al.	20 extremely low birth weight (<1000 grams) preterm (<30 weeks) infants on supplemental oxygen	12 hours	85% - 93%	Within: 33.7% Below: 11.5% Above: 54.8%	Within: 58% Below: 14% Above: 26.5%

Study Title	**Authors**	**Study Cohort**	**SpO_2 Sampling Frequency**	**Target SpO_2 Range**	**Target SpO_2 Range Compliance (Manual FiO_2 Adjustments)**	**Target SpO_2 Range Compliance (Automated FiO_2 Adjustments)**
Closed-loop automatic oxygen control (CLAC) in preterm infants: a randomized controlled trial.	Hallenberger A, Poets CF et al.	34 preterm infants from 4 centers on supplemental oxygen (median GA: 26.4 weeks)	Two 24 hours periods (one on manual FiO_2 adjustment by clinical personnel and one on automatic adjustment)	Lower limit: 80%-90% Higher limit: 92%-95%	Within: 61.4% Below: 15% Above: 16%	Within: 71.2% Below: 9.1% Above: 15.9%

Overall, we can see that patients are within their target range around 40%-60% of the time; we aimed to compare our study cohort's target SpO_2 range compliance with these studies. These studies also demonstrate that patients seem to spend more time above their target range rather than below their target range; however, this is likely due to focusing on patients that are on supplemental oxygen.

When analyzing these studies, we note that all of the studies had comparatively small cohort sizes of less than 100 patients and were mainly looking at extremely preterm (<28 weeks) infants for a limited period of time usually only over the first few weeks of life; we aimed to expand upon these studies by analyzing a larger cohort size of all preterm (<37 weeks) infants with continuous monitoring. Additionally, there was a focus on monitoring patients only when they are on supplemental oxygen and focusing on time spent above the target SpO_2 range and associated hyperoxemic events; however, the occurrence of hypoxemic events is another valuable research question that should also be explored.

These studies also noted many gaps in knowledge that future work should explore such as how to best titrate oxygen therapy in extremely preterm infants, what the optimal target SpO_2 range should be, what center and clinical characteristics may determine successful maintenance of target SpO_2 ranges, what degree of compliance of the target SpO_2 ranges is feasible, what the short and long effects of varying compliance with the target SpO_2 range are, and what the distribution of SpO_2 value are for infants on supplemental oxygen [3, 11, 14]. Our work addresses clinical characteristics that may determine successful maintenance of target SpO_2 ranges and the degree of compliance of the target SpO_2 ranges observed in our study cohort.

2.4 Factors Associated With Maintaining the Target SpO_2 Range

As noted in Table 2.1, the studies and reviews also examined what factors may be root causes affecting this low compliance to the target SpO_2 range which can help

keep clinicians aware of how to maintain the target range.

One way that clinicians can maintain the target oxygen range is by adjusting the FiO_2 to their patients; however, it is important to deliver appropriate amounts of oxygen without creating oxygen toxicity and hyperoxemia. When a patient is on supplemental oxygen (i.e. $FiO_2 > 21\%$), studies hypothesize that there may not be proper guidelines for titrating the FiO_2 (such as when and how a patient should be weaned off of FiO_2), leading to the risk of patients being exposed to the aforementioned oxygen toxicity [34]. Additionally, clinicians may not be aware of the adverse effects of hyperoxemia on premature infants and therefore may not understand the importance of proper FiO_2 titration [34].

When nurses were interviewed in one study, they confirmed the importance of clinical awareness by stating that education is one of the factors that may affect maintaining the target SpO_2 range along with response time to alarms [11]. Other factors suggested to affect maintaining the target SpO_2 range were lack of root cause analysis at the bedside, and the amount of priority given to control of oxygen therapy [11]. Another study found that tackling some of these issues through incorporating both nurse education and an algorithm for FiO_2 titration led to less time above the target oxygen saturation range and, consequently, less instances of ROP [34]. However, other studies have found that educating nurses has not changed or even increased the amount of time that patients spend outside their target SpO_2 range [31]. On the other hand, multiple studies demonstrated improved time in the target SpO_2 range with decreased hyperoxemic events with the use of automated FiO_2 control compared to manual adjustment, as seen in Table 2.2; however, time below the SpO_2 limit did sometimes increase with the use of automated FiO_2 control [12, 23, 24, 26, 25]. Providing proper guidelines and education may or may not be sufficient in increasing awareness of the clinical staff to prevent time outside the target SpO_2 range.

Other factors that may relate to maintaining the target SpO_2 range includes nurse to patient ratio, gestational age, postmenstrual age, respiratory support mode, and setting of alarm limits [8, 14, 31]. Studies have found that if a nurse has a greater

patient load, there is an increase in the time the SpO_2 is above the target upper limit [8, 31]. Another study found that lower gestational age was associated with improved time spent in the target range [14]. Similarly, patients at lower postmenstrual ages had higher target range compliance with lower hyperoxemic time; however, patients had more hypoxemic time at lower postmenstrual age [8]. These findings could be because it becomes more difficult to maintain oxygen saturations within the target range as infants mature, require less invasive respiratory support, and may develop other complications such as lung disease [11]. Regarding respiratory support mode, a study demonstrated that nasal cannulae was associated with a higher amount of time above the target SpO_2 range while high-frequency ventilation was associated with the least amount of time [10].

Exploring SpO_2 alarms, studies have found that the alarm limits may not always be set correctly; for example, upper alarm limits may be set correctly only approximately 25% of the time and are otherwise set too high [4, 20, 31]. One reason for incorrect alarm limit setting is because clinicians want to decrease the number of unnecessary alarms to prevent alarm fatigue. Alarm fatigue occurs due to having a large number of alarms going off in a busy NICU; many alarms do not require intervention due to being brief or false, so increasing the upper alarm limit is an attempt to mitigate the alarm burden [15]. However, this incorrect alarm setting may lead to decreased compliance since clinical staff would not be alerted to when a patient is over the target SpO_2 range and would miss treating the corresponding hyperoxemic episodes. To support this hypothesis, another study found that having a policy to set the alarm limits to be close to the target SpO_2 range limits led to improved target range compliance [14]. This study also determined that higher target range upper limits and wider alarm limits were also associated with improved target range compliance [14].

Addressing all of these aforementioned factors in the NICU is important for clinical staff to understand what to focus on to best improve compliance to the target SpO_2 saturation range. In our study, we explored these factors and examined how our findings compare to previous studies.

2.5 Intermittent Hypoxemia

Most of the aforementioned non-compliance is associated with the upper SpO_2 limit leading to studies focusing on hyperoxemia; however, it is also important to note the causes and consequences of intermittent hypoxemia (IH). There is not currently a universal definition of IH and some proposed definitions used in multiple studies include: $SpO_2 < 80\%$ for 10 seconds or longer [6], $SpO_2 < 80\%$ for between 10 seconds and 3 minutes [16, 28], a drop in SpO_2 of more than 3% from a baseline [9], and a drop in SpO_2 of more than 10% from a baseline for more than 2 seconds [21].

IH can be caused by respiratory instability which is common in preterm infants due to immature respiratory control resulting in increased metabolic oxygen consumption, respiratory pauses, and poor respiratory function [17]. A study has shown that late preterm infants specifically are more at risk for IH than term infants [21]. These IH episodes tend to progressively increase over the first month of a preterm infant's life until they plateau and then begin to decline around Week 7 [28]. Researchers have found that episodes usually take place in clusters less than 20 minutes in total duration; about half of the events occur less than 1 minute apart [17].

2.6 Outcomes and Prevention of Intermittent Hypoxemia

There are many adverse short and long term outcomes associated with the occurrence of intermittent hypoxemia episodes. Examples of morbidities include ROP, wheezing, bronchopulmonary dysplasia, and neurodevelopmental impairment [17]. A study found that hypoxemic episodes that last more than 1 minute in extremely preterm infants were associated with an increased risk of late death or disability such as motor impairment, cognitive or language delay, severe hearing loss, or bilateral blindness [6]. Through these outcomes, we can see the importance of understanding potential factors contributing to intermittent hypoxemia in order to prevent it. Potential ways of preventing IH episodes include similar aforementioned strategies for preventing hy-

peroxemic episodes such as optimizing the baseline SpO_2 and target SpO_2 range and automating control of inspired oxygen as well as xanthine therapy [28].

2.7 Neonatal Intensive Care Unit at Beth Israel Deaconess Medical Center

Our study focuses on the NICU at BIDMC which is comprised of two nurseries: Reisman and Stoneman. Our focus is on the Reisman NICU which is structured such that there are 24 rooms that can each hold up to three beds in a room. Typically, each room is occupied by two bed spaces (one by the window and one by the door); however, a third bed space may be introduced as an extension bedspace when necessary. Additionally, there are virtual bed spaces recorded in the electronic medical record which a patient may be temporarily recorded in when the clinical staff is aware that the patient will soon be admitted; these virtual bed spaces do not represent a physical location in the NICU. The physical rooms are arranged on the outside margin of a large U-shaped corridor which is labeled North and South where there are staff work areas and six central monitoring systems throughout the inner margin of the U [29]. This layout can be seen in Figure 2-1.

Figure 2-1: Layout of Reisman NICU at BIDMC.

2.8 Data Warehouse Connect

Data Warehouse Connect (DWC) is a bedside data recording and archiving infrastructure by Philips Healthcare that is implemented in the BIDMC intensive care units. This infrastructure contains real time data regarding patient alarms, trends, and physiological waveforms which flow from up to 73 NICU beds in the Reisman NICU to database servers at MIT every four hours starting from December 12, 2017 to present; the Stoneman unit acquired DWC licenses in October 2019 and was not included in our analysis. There is an abundance of data, including waveforms (such as the electrocardiogram, arterial blood pressure, and airway pressure) in their native formats of up to 500 samples per second; trend data (such as heart rate and oxygen saturation) in the format of 1 sample per second and accurate to the tenth of a percent; and patient alarm data detailing alarm severity, alarm start and end time and if the alarm has been silenced.

This resulting data is organized into multiple tables in a SQL database that contain information about the patient's stay (i.e. admission time, discharge time, bed assignment), information about alarms that went off during the patient's stay, information about the patient's SpO_2, and some degree of information regarding the care providers' interactions with the monitor. Key tables for our purposes are the patient table, the patient string attribute table, the alert table, the numeric value table, the patient mapping table and the numeric mapping table; the structure of these tables can be seen in Table A.1.

The patient table documents patient information and contains columns recording a patient's unique identifier, time stamps for a patient's stay, the bed spaces a patient is recorded in (room number and bed number), the patient's gender, height and weight, the patient's resuscitation status, the patient's admission state, and the patient's clinical unit.

The patient string attribute table documents additional information that is related to each patient. For each unique patient identifier, there may be associated string attributes (i.e. FirstName, LastName, LifetimeId) and their corresponding values.

The alert table documents all of the monitoring alarm information and creates a new row any time a patient monitoring alarm updates. This table contains columns recording the associated patient ID, time stamps for alerts, the source and label of the alerts, the alert's severity, if the alert is silenced, and the announce, onset, and end time for the alert. The patient ID recorded in this table is actually an alternative mapping ID that is recorded in the patient mapping table. The patient mapping table documents a mapping between this alternative mapping ID and the ID used in the patient and patient string attribute table.

The numeric value table documents numeric values (such as the second-by-second SpO_2 values) for a given patient and contains columns recording the associated alternative patient ID, the time stamps for the values, the ID of the type of numeric recorded, the value of the numeric, and if the trend is uploaded. The numeric table is used in conjunction with the numeric value table as this table documents all of the different physiologic trend data that are collected in DWC. This table contains

columns mapping the ID of the type of numeric recorded to the label of the numeric, the upper and lower limits of the numeric, and the unit of measure for the numeric.

From December 12, 2017 to December 31, 2019, in the Reisman unit, there are 5771 unique patient IDs recorded in the patient table and 2648 unique MRNs recorded in the patient string attribute table. Each patient ID has associated rows recorded in DWC for an average of 5 days, 5 hours, 55 minutes and 57 seconds and 2143 alerts recorded on average. Each patient ID has on average 1 day, 22 hours, 6 minutes and 57 seconds of SpO_2 data leading to a total of approximately 11,225 patient days of SpO_2 data. We note that we should not assume from this that we only have approximately 25% SpO_2 data coverage and this demonstrates the importance of data integrity checks that are performed later. As discussed in Section 4.1, there are a large number of patient IDs recorded in the patient table and not all of these are necessarily "valid" patients; therefore, these IDs would not have any associated SpO_2 data. The true SpO_2 data coverage will be discussed more in depth in Section 5.1.

2.9 Open Questions

The aforementioned previous studies inspired the current work as we aimed to continue to find specific factors that may be influencing a neonate's oxygen levels and bring these to the attention of clinical staff to determine if it is possible to improve the workflow and treatment for patients to keep them in their target oxygen range. Through our research, we have access to massive amounts of real time physiological and clinical data streams that we have taken advantage of and analyzed retrospectively to gain insights into factors that may be affecting the oxygen saturation of premature infants. Many studies have emphasized educating clinical staff on the importance of maintaining patients' oxygen saturation within the target oxygen range. We are able to encourage this education by closing the loop between researchers and clinicians and present insights and analyses to the clinicians that the researchers uncover from the data.

Additionally, we have found that previous studies tend to involve small study co-

horts of a subset of the NICU population (< 100 patients) such as extremely preterm (< 28 weeks) or extremely low birth weight (< 1000 grams) patients and have a focus on either hypoxemia or hyperoxemia. These studies also did not have access to continuous SpO_2 recordings and other studies have also stated that manual documentation of hyperoxemic and hypoxemic episodes may result in significant underreporting [15]. Due to the massive nature of our dataset, we are able to have a more comprehensive study cohort that includes all preterm infants in the NICU for more than 24 hours over the course of 21 months; we also have access to more granular data to leverage to obtain more accurate reports of both potential hypoxemic and hyperoxemic episodes. Having access to such a large database allows us to perform comprehensive analyses regarding oxygenation of preterm infants which have not been explored before.

Chapter 3

Study Cohort

Through this chapter, we define our study cohort and provide a description of the cohort in terms of commonly used clinical indicators.

3.1 Study Cohort Definition

Our study cohort consists of all preterm neonates (gestational age < 37 weeks) in neonatal critical care in the Reisman unit for more than 24 hours from January 1, 2018 to September 12, 2019.

3.2 Cohort in Beth Israel Deaconess Medical Center Clinical Logs

In addition to the data streams received through DWC, we also receive clinical logs from the BIDMC NICU electronic medical record system. These logs are snapshots of the NICU census taken around 22:00:00 daily starting from January 1, 2018 through to September 12, 2019 from 105 bed spaces comprising the beds in both the Reisman and Stoneman units as well as virtual bed spaces. Each row lists the date and time of the snapshot as well as a patient's medical record number, the bed the patient is in at that time (the unit, room number and bed number), the patient's first and last

Table 3.1: Clinical characteristics of study cohort and overall NICU population. Data are medians and interquartile ranges.

	Study Cohort	**NICU Population in Clinical Logs**
Number of Patients	865	1633
Gestational Age [weeks]	34.1 (32 - 35.3)	35.6 (33.4 - 38.7)
Birth Weight [grams]	2050 (1580 - 2430)	2540 (1890 - 3280)
Length of Stay [days]	14.7 (6.1 - 34.7)	6.1 (1.2 - 18.4)

name, the patient's admission and discharge date, the patient's date of birth, and the patient's gestational age and birth weight.

These clinical logs contain 1633 unique medical record numbers (MRNs) with 865 MRNs fitting our inclusion criteria of being preterm and in the NICU for more than 24 hours. Table 3.1 summarizes the median and interquartile range of the demographic information of our study cohort as well as the overall NICU population recorded in these clinical logs; note that if the patient is in the NICU for less than 24 hours, they will not be recorded in the clinical logs unless their stay overlaps with the time the data snapshot is taken. For our study cohort, all of our patients are in the NICU for more than 24 hours, so the BIDMC clinical logs are sufficient.

To further explore the characteristics of our study cohort, we produced distributions of gestational age, birth weight, and length of stay which can be seen in Figure 3-1. It should be noted that of these 865 patients, there are 122 very low birth weight (less than 1500 grams and greater than 1000 grams) patients, 62 extremely low birth weight (less than 1000 grams and greater than 500 grams) patients and 5 micropreemies (less than 500 grams). There are 140 very preterm infants (less than 32 weeks and greater than 28 weeks) and 64 extremely preterm infants (less than 28 weeks). Additionally, we evaluated various trends in the data to check if they aligned with our expectations.

We found a positive relationship between gestational age and birth weight which aligns with the literature as well as our intuition as if a patient is born more premature,

they tend to be less developed and typically weigh less. This relationship can be seen in Figure 3-2.

We found an inverse relationship between the patient's birth weight and their length of stay; this intuitively makes sense because if a patient is born at a lower birth weight, the patient will most likely need more time in the hospital to get healthy before discharge. This relationship is seen in Figure 3-3.

We also found an inverse relationship between the patient's gestational age and their length of stay; this intuitively makes sense because if a patient is born more premature, the patient will similarly need more time in the hospital to get healthy before discharge. This relationship is seen in Figure 3-4. We explored some of the outliers (patients with GA < 27 weeks and length of stay < 50 days) and found that their short length of stay was due to redirected care, the patient being transferred or neonatal demise in labor and delivery.

Through analyzing these relationships, we can determine that the trends found confirm our expectations.

Through these clinical logs, we also have access to supplemental oxygen usage through hourly recordings of patient's fraction of inspired oxygen (FiO_2). Table 3.2 summarizes the median and interquartile range of the demographic information of the 412 patients (out of the 865 patients in our study cohort) recorded on supplemental oxygen at some point of time while in the Reisman unit.

3.3 Summary

Our study cohort of 865 patients is larger than study cohorts in previous literature as discussed in Chapter 2 and we can also take subsets of this study cohort to get comparable numbers of extremely preterm or extremely low birth weight patients. Leveraging this large study cohort along with the continuous SpO_2 monitoring offered through DWC allows us a great opportunity to provide insights regarding factors that affect a patient's time within their target SpO_2 range.

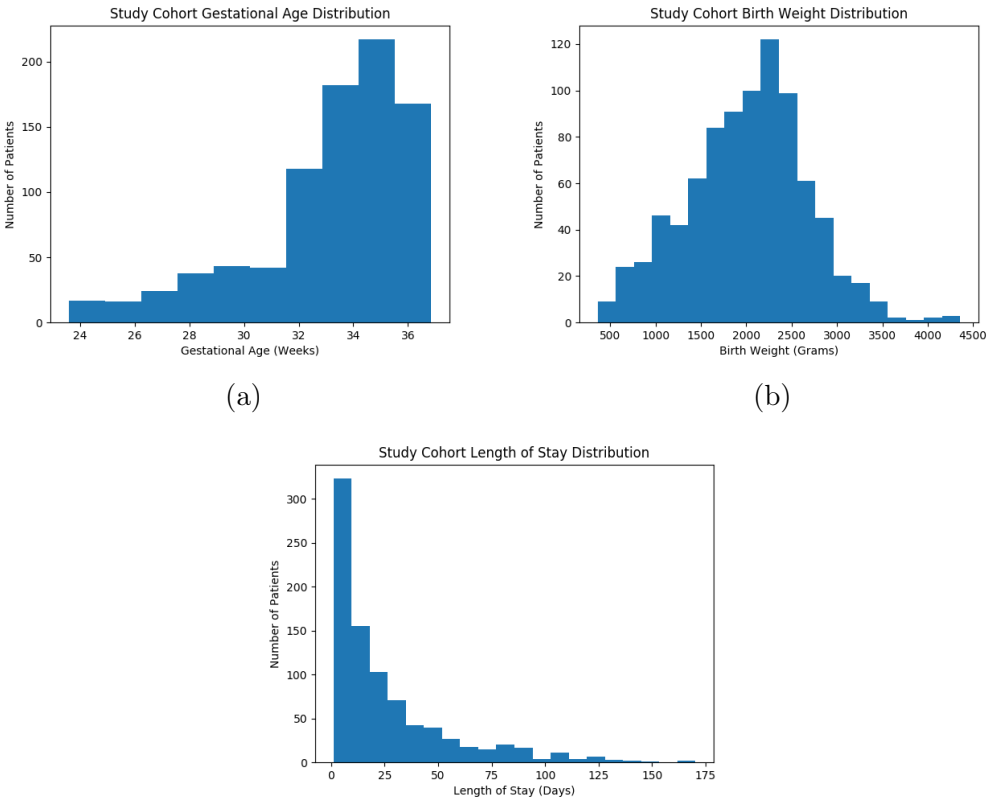

Figure 3-1: Distribution of (a) gestational age, (b) birth weight, and (c) length of stay of study cohort.

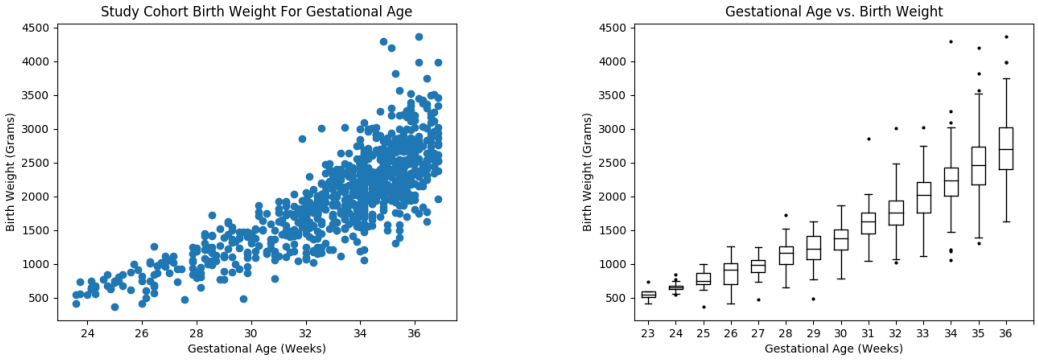

(a) (b)

Figure 3-2: Scatter and box plot illustrating the relationship between gestational age and birth weight.

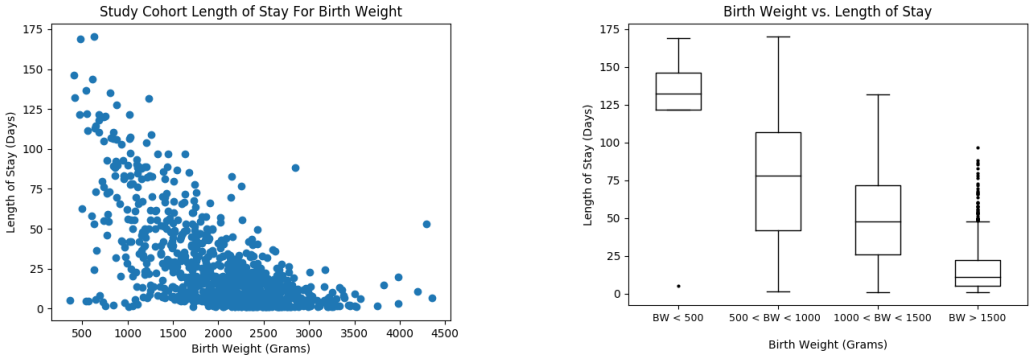

(a) (b)

Figure 3-3: Scatter and box plot illustrating the relationship between birth weight and length of stay.

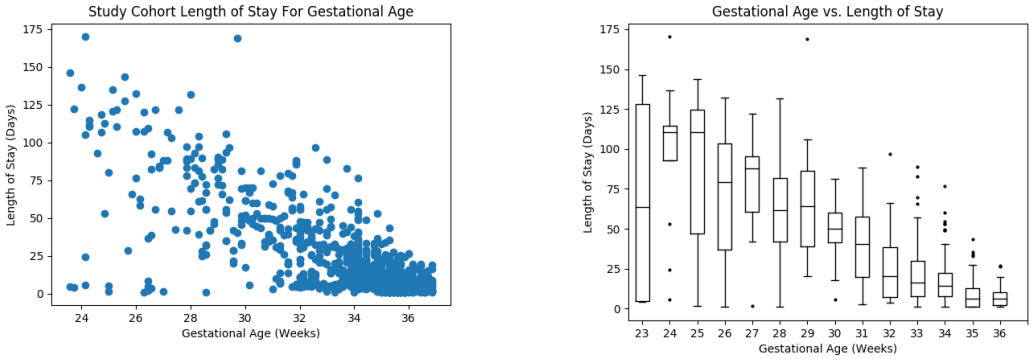

(a) (b)

Figure 3-4: Scatter and box plot illustrating the relationship between gestational age and length of stay.

Table 3.2: Clinical characteristics of study cohort patients recorded on supplemental oxygen. Data are medians and interquartile ranges.

	Study Cohort Patients on Supplemental Oxygen	**Study Cohort Patients not on Supplemental Oxygen**
Number of Patients	406	459
Gestational Age [weeks]	33.1 (29.7 - 34.9)	34.3 (33.2 - 35.4)
Birth Weight [grams]	1915 (1250 - 2420)	2150 (1753 - 2452)
Length of Stay [days]	21.1 (8.0 - 53.1)	11.3 (4.8 - 24.1)
Time on Supplemental Oxygen (Hours)	6.5 (1.1 - 82.5)	-
Number of FiO_2 adjustments per day	2 (1 - 4)	-
Time after Birth Patient First Receives Supplemental Oxygen [hours]	16 (10.5 - 23.1)	-

Chapter 4

Data Integrity

As DWC is a recently introduced infrastructure at BIDMC and the associated data elements have not been examined in depth yet, the first major step focused on exploring the data and determining its cleanliness, completeness, and correctness. We can accomplish this through two different ways: exploring the composition of each of the tables provided by DWC and matching the data recorded by DWC to the aforementioned BIDMC clinical logs.

Through DWC, we have key information to identify patients in the NICU, such as the patient ID, medical record number (MRN), first name, last name, and bed space assignment. We can explore each of these features to determine the validity of the recorded data.

4.1 Patient IDs in DWC

For our purposes, we used DWC data spanning from January 1, 2018 to September 11, 2019. DWC recorded 4867 patient IDs during this timeframe in the NICU. However, we note that, upon further inspection, not all of these patient IDs necessarily correspond to a patient; in some cases, the patient ID may never be recorded as admitted to the hospital or may be recorded in DWC for a minimal amount of time.

For each patient ID, we observed records of a patient ID's admit state (admit state = 0 means not admitted, admit state = 1 means admitted, admit state = 2 means

Table 4.1: Example of a non-admitted patient ID recorded at the same time that an actual patient is admitted or discharged from the patient table. In this example, this is the only occurrence of Patient ID B and we can see that the patient is never admitted to the hospital and is recorded in DWC for only 1 second; during this second, a patient ID that was previously admitted is discharged. Both of these patient IDs are associated with the same bed spaces. Based on these observations, we can determine that Patient ID A refers to a patient while Patient ID B is an artifact which can be disregarded.

Patient ID	**Time Stamp**	**Admit State**	**Bed Label**
A	2019-07-30 04:49:47	1 (Admitted)	Bed 1
B	2019-08-01 13:13:43	0 (Not Admitted)	Bed 1
A	2019-08-01 13:13:43	2 (Discharged)	Bed 1
B	2019-08-01 13:13:44	0 (Not Admitted)	Bed 1

discharged). We found that 2678 of the 4867 patients IDs are never admitted to the NICU and only have an admit state of 0 recorded. Additionally, of these 2678 patient IDs, 2665 have records spanning less than 24 hours indicating that the majority of these would not be included in our study cohort anyway.

We looked into these non-admitted patient IDs to determine what information may be associated with them. These non-admitted patient IDs are introduced to the patient table for a very minimal amount of time (2519 of the 2678 are recorded for less than 5 seconds); we also observe that these non-admitted patient IDs often are recorded at the same time that an actual patient is admitted or discharged from the patient table, so we treat these patient IDs as artifacts (see Table 4.1). Based on this, we can conclude that the non-admitted patient IDs very likely do not contain important information for the purpose of our study and we are not losing patient information by disregarding them. Therefore, we concluded that there are 2189 valid patient IDs admitted to the NICU between January 1, 2018 and September 11, 2019; 1696 of these patient IDs are recorded in DWC for more than 24 hours which is comparable to the 1633 MRNs recorded in the clinical logs.

4.2 Patient String Attributes: Medical Record Numbers, First Names, and Last Names

In the patient_string_attribute table, we have access to a patient's medical record number (MRN) which is entered by staff through a patient's monitor under the field name "LifetimeId". We also have information about a patient's first name and last name which are entered in a similar way. We found that 2173 patient IDs of the 4867 patient IDs have either a MRN, first name and/or last name recorded, and all of these patient IDs are recorded as admitted to the hospital. We also note that none of the 2683 patient IDs which were never admitted to the monitor have an MRN, first name, or last name, which confirms that these are most likely artifacts that can be disregarded.

4.2.1 First Names and Last Names

Looking at the records of first names and last names specifically, we found that a subset of the 2173 patient IDs have a first name or last name. There are 1357 patient IDs with a recorded first name; 535 of these first names are unique. The two most popular first names are Boy (411 occurrences) and Girl (316 occurrences) which indicates that clinical staff may not always know or be specifying the first name of the patients when they are admitted to the monitors. 1350 of these 1357 patient IDs also have a recorded MRN.

On the other hand, there are 2168 patient IDs with a recorded last name; 1979 of these last names are unique indicating that last names are recorded more often than the first names. The most frequently reoccurring last names seem to be associated with car seat tests; 29 last names are Carseat, 13 are Carseat test, and 11 are Car seat. 2165 of these 2168 patient IDs also have a recorded MRN. This follows a recommendation by the American Academy of Pediatrics (AAP) that preterm infants should be observed in a car safety seat before discharge to monitor for potential apnea, bradycardia or oxygen desaturation [27].

Figure 4-1 illustrates the composition of the MRNs, first names, and last names recorded.

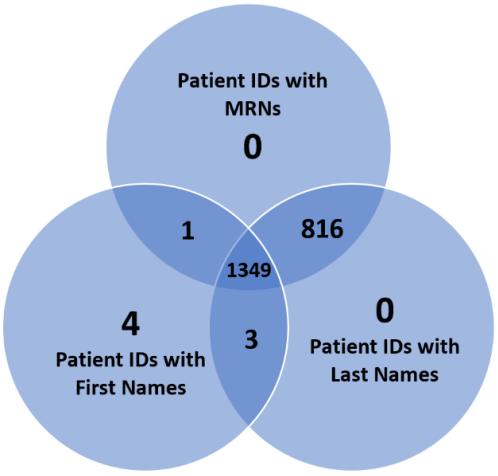

Figure 4-1: Venn diagram illustrating the composition of the MRNs, first names, and last names recorded in the patient_string_attribute table in DWC.

Based on these results, we can conclude that MRNs and last names are the most commonly inputted identifiers and we can leverage this when we later match DWC patient IDs to patients in the clinical logs.

There are 16 patients that do not have an MRN, first name or last name, but are recorded as admitted to the NICU in the patient table; however, only 5 of these are recorded for more than 24 hours and would potentially be included in our study cohort. This is a minimal number, but we acknowledge that these patient IDs may refer to actual patients that the clinical staff did not manually input the identifying information for in the patient monitor; this will be accounted for when we match DWC patient IDs to patient MRNs in Section 4.4.

4.2.2 Medical Record Numbers

As aforementioned, through DWC, we have access to a patient's medical record number (MRN) for a subset of the 4867 patient IDs; there are MRNs for 2166 unique patient IDs from January 1, 2018 to September 11, 2019. There are 2298 total recorded

MRNs indicating that some patient IDs have multiple MRNs associated with it; 2219 of these recorded MRNs are unique indicating that some MRNs have multiple patient IDs associated with it. Figure 4-2 illustrates these potential issues.

We discovered that 120 patient IDs have multiple MRNs assigned. Diving into these patient IDs, we observed that some patient IDs have multiple MRNs associated with them because the first recorded MRN is a "dummy" MRN; this may be because the clinical staff was unable to document a correct MRN for a patient until a later time. Some examples of dummy MRNs are multiple 0s in a row (which occurs 33 times), 1234567 (6 times), and 007 (5 times).

We also discovered that 39 MRNs have multiple patient IDs associated with them. The aforementioned dummy IDs may also be one of the reasons for this occurrence.

We acknowledge that there are some valid MRNs that have multiple patient IDs associated with them and some patient IDs that have multiple valid MRNs associated with it. These two issues will be addressed in more depth when we match patient IDs to MRNs in Section 4.4.

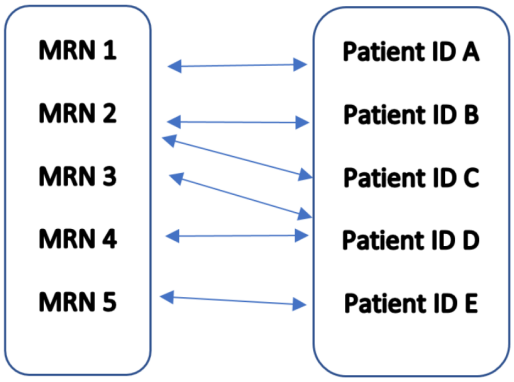

Figure 4-2: Chart illustrating the potential issue of having an MRN mapped to multiple patient IDs and a patient ID mapped to multiple MRNs. In this example, MRN 2 is assigned to two patient IDs and Patient ID D is assigned to two MRNs.

4.3 Bed Spaces

We also have access to bed spaces that a patient ID is associated with. We can analyze the bed spaces recorded for each patient ID in DWC and reconstruct a patient's NICU

stay in an interpretable way to assess the cleanliness and completeness of the data. Bed spaces will also be able to be leveraged when we match patient IDs to MRNs as we can observe what bed space they are recorded in which maps to the physical location in the NICU.

We explored a sample patient's stay in DWC by extracting unique entries in the patient table as data points; each data point represents the unique time stamp and bed space recorded in DWC. Figure 4-3 illustrates a sample patient's stay in DWC; it illustrates each data point that we have for the patient and with some assumptions about the time between each data point, we can transform these data points to the chart in Figure 4-4.

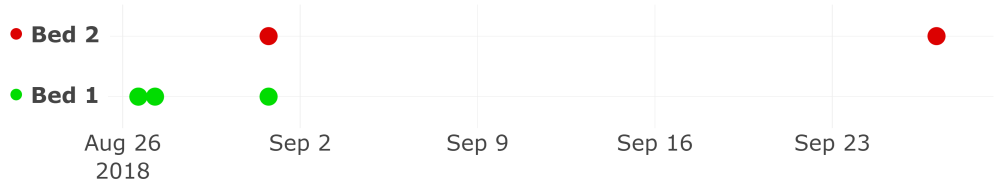

Figure 4-3: Chart illustrating a sample patient's stay in DWC; each data point represents a time stamp that a patient was recorded in a bed space. Through this, we can see how a patient is recorded in both Bed 1 and Bed 2 over the course of 3 days.

Through analyzing charts like Figure 4-4 for NICU patients in our study cohort, we discovered two inconsistencies in the data which should be addressed. One inconsistency is unreliable time between bed spaces. We want to assume a straightforward progression of the patient through the bed spaces where a patient is moved from one bed space to another. However, there are some times when the first time a patient is seen in one bed space occurs before the last time a patient is seen in the original bed space; we consider this an "overlap" of the bed spaces.

An example and description of this overlap is detailed in Figure 4-5. If we interpreted a patient's stay in DWC without accounting for potential overlaps, we would end up recording unnecessary movement between bed spaces. This occurs in approximately 14% of the admitted patient IDs (305 patients IDs / 2189 admitted patient IDs); the median "overlap" time is 13 seconds.

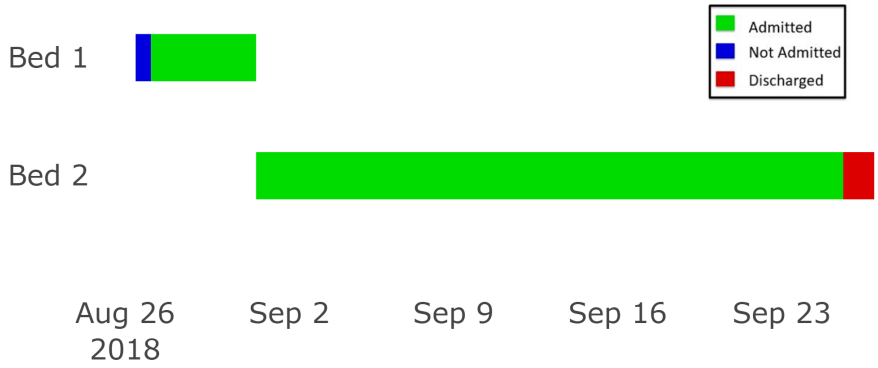

Figure 4-4: Chart reconstructing the patient stay in Figure 4-3 by assuming that the first time stamp recorded in the bed space is the time that the patient is admitted to that bed space and similarly, the last time stamp recorded in the bed spaces is the time that the patient is discharged from that bed space and either is discharged from the hospital or moved to another bed space. Through this, we can see how a patient moves from Bed 1 to Bed 2 over the course of 3 days.

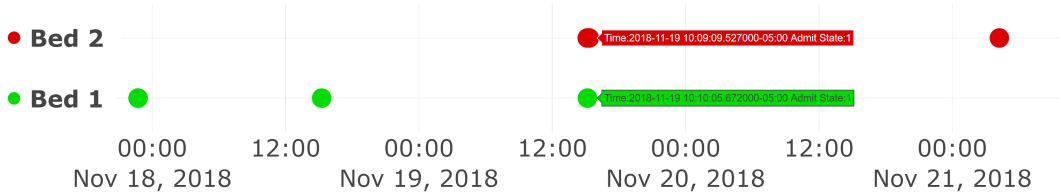

Figure 4-5: Chart illustrating an example of an "overlap" of the bed spaces. The patient is observed to be in Bed 2 at 10:09:09 which is approximately one minute before they were last observed to be in Bed 1 at 10:10:05. If we interpreted this without accounting for this inconsistency, we would believe that the patient moves from the Bed 1 to Bed 2 to Bed 1 back to Bed 2.

We decided to generally combat this inconsistency by still assuming the straightforward progression and swapping the timestamps of the “overlap” when the overlap is found to be less than 5 minutes. For example, in Figure 4-5, the last point in the Bed 1 would now occur at 10:09:09 and the first point in Bed 2 would now occur at 10:10:05 leading to the linear progression and allowing us to understand where a patient is in the NICU. We chose this arbitrary threshold of 5 minutes as it is unlikely that a patient is monitored in one bed space for less than 5 minutes before returning to their original bed space.

We should note that this is not a major discrepancy and most likely occurs due to a short delay in discharging the patient from one bed space before admitting them to another; this delay only accounts for less than 5 minutes of the data which is insignificant.

The other discovered inconsistency is unaccounted time in a bed. This occurs when there is only one data point in a bed space so it may be unclear when the patient was admitted or discharged from that bed space.

In the examples from Figure 4-3 and Figure 4-5, the patient has data recorded at least two points for each bed space: the start of their stay in a bed space as well as at the end of their stay in the same bed space; this is the ideal situation as we can easily assume that the patient was in the single bed space from the first point to the last point. In Figure 4-3, the patient is recorded in Bed 1 three times making it clear what the admit and discharge point is for this bed; immediately after the end point in Bed 1, we observe the patient in Bed 2. Similarly, there are two data points of the patient in Bed 2 so we can assume the patient was in the Bed 2 during this whole duration. However, there are some times when a patient ID only has one data point recorded in a given bed space which leads to uncertainties regarding data coverage and the amount of time spent in this bed space; we consider this a "gap" in time. This occurs in approximately 13% of the admitted patient IDs (287 patient IDs).

An example and description of this gap is detailed in Figure 4-6. If we interpreted a patient's stay in DWC without accounting for potential gaps, we would have periods of time where we are unsure where a patient is in the NICU.

We decided to combat this discrepancy by assuming that when a bed space discharge time is missing, the patient ID is in that bed space for the whole duration until they are recorded in the next bed at the next sequential time stamp (or when the bed space admit time is missing, we assume they are in the bed starting from the last sequential time stamp for the patient ID). In Figure 4-6, this would mean that the patient was in Bed 2 for the whole duration until the patient is in Bed 3. If there is only one record for the patient ID, we will assume that this is their admit time and the discharge time will be the next sequential time stamp of the next patient ID

recorded in the same bed space; this will lead to an overestimate of the patient ID's discharge time, but attempts to ensure that any alarm or SpO_2 data associated with this ID is captured.

Both of these inconsistencies demonstrate that there are some discrepancies in the data; however, we are able to account for these issues by making assumptions regarding how the amount of time a patient spends in a bed space so that we can cleanly interpret and use the data.

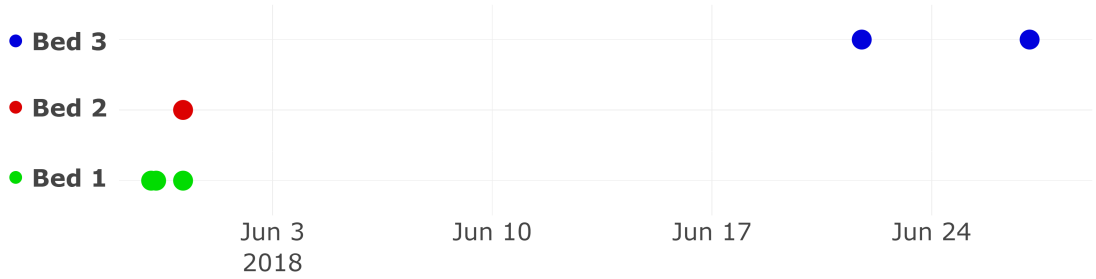

Figure 4-6: Chart illustrating an example of a "gap" in time. The patient is observed to be in Bed 2 at only one data point leading us to be unsure how long the patient remains in the Bed 2 for or if they were moved to another unrecorded bed before moving to the Bed 3.

4.4 Matching DWC Patient IDs to Medical Record Numbers

Another way of validating the data in DWC is by matching the data recorded by DWC to other data sources including the aforementioned BIDMC clinical logs. While we have observed that some Patient IDs already have mappings to MRNs through the patient_string_attribute table, it is still necessary to map the remaining Patient IDs to MRNs so that we can confirm that our study cohort can be identified in the DWC data. We developed an algorithm to match DWC patient IDs to the MRNs by cross referencing the BIDMC clinical logs to match patient's last names, MRNs and bed spaces.

Based on the aforementioned explorations and analyses, we developed an algorithm to map DWC patient IDs to MRNs so that we can match DWC patient IDs to relevant data in the clinical logs such as gestational age, birth weight, and usage of supplemental oxygen; this mapping process also validates the data in DWC as it demonstrates the correctness and completeness of the data when compared with the clinical logs. We note that this algorithm is generalizable and scalable so when the clinical logs are updated leading to a new study cohort, we can rerun the code to retrieve an updated patient mapping.

The following provides a brief overview of the steps to generate this patient mapping.

> Patient Mapping Algorithm:
>
> Input: Clinical log, DWC patient table and DWC patient_string_attribute table
>
> Output: A mapping of MRNs to DWC patient IDs.
>
> **Step 1.** Get study cohort from clinical logs: list of MRNs fitting our inclusion criteria of having a gestational age of < 37 weeks and being in the Reisman NICU for > 24 hours.
>
> **Step 2.** Ensure each patient in our study cohort has visited at least one of the bed spaces recorded in DWC.
>
> **Step 3.** Match Patient IDs to MRNs that are recorded in the patient_string_attribute table in DWC.
>
> **Step 4.** Determine what patients still need to be matched.
>
> **Step 5.** Match remaining MRNs to Patient IDs that visited the same bed spaces in the same time span and also may have the same last name in DWC as recorded in the clinical logs.
>
> **Step 6.** Return the mapping.

We will now describe the execution of each step in the above algorithm and the associated results.

This algorithm took the clinical logs that contained data from January 1, 2018 to September 11, 2019, and the patient and patient_string_attribute tables in DWC

as inputs; the clinical logs were used to accurately map the patient IDs to MRNs based on information stored in the patient and patient_string_attribute tables. The algorithm outputted a mapping that mapped a MRN to a patient ID.

Step 1. The first step is getting the study cohort from the BIDMC clinical log file; this consists of 865 MRNs that fit our inclusion criteria of having a gestational age of <37 weeks and a length of hospital stay of >24 hours. The patient mapping algorithm will attempt to map each of these MRNs to a DWC patient ID.

Step 2. The clinical logs contain patient data for Reisman, Stoneman, and virtual bed spaces; however, DWC only has data from the Reisman unit up until September 11, 2019. Therefore, our study cohort excludes patients that are exclusively recorded in Stoneman and virtual bed spaces as we will not have any DWC data for these patients. This occurs for three patients: two of these patients are only recorded in virtual bed spaces and one patient is recorded in Stoneman for approximately one week. Therefore, we are only able to match at most 862 MRNs to DWC patient IDs.

Step 3. Next, we determined how many of our study cohort MRNs we can directly match in the patient_string_attribute table so that we can match the associated patient IDs. We found that there are 847 (patient_ID, MRN) mappings found through this direct search; this consists of 842 unique patient IDs and 817 MRNs.

We noticed that there are 5 patient IDs that have 2 MRNs mapped to them; 26 MRNs that have 2 patient IDs mapped to them; and 2 MRNs which have 3 patient IDs mapped to them. This introduces the concern of not having a strict one-to-one mapping. In the case when a patient ID is mapped to multiple MRNs, we understand the importance of knowing the admit and discharge times of an MRN which will be addressed in Section 4.5. This allows us to be aware of when a patient ID corresponds to which MRN. In the case when an MRN is mapped to multiple patient IDs, we understand that we will have to aggregate the statistics across all the patient IDs in DWC data; for example, when retrieving all the bed spaces a patient MRN visited in DWC, we will have to aggregate all the bed spaces across all the mapped patient IDs.

We note that the MRNs are recorded as Lifetime IDs in DWC and are inputted by clinical staff which leads to inconsistent formats of the Lifetime IDs as well as occasional typos. Alternative formats of the Lifetime IDs include having the numbers separated by spaces (### ## ##) or the numbers separated by dashes (###-##-##). To account for this, we performed a regex analysis to return all the patient IDs that have a Lifetime ID that contains all of the digits of the desired MRN; this also accounted for potential typos where the clinical staff may have inputted an extra digit.

Through this process, we discovered one Lifetime ID that consisted of two concatenated MRNs which we do not want to incorrectly match to only one MRN. To account for this and other potential similar cases, we discarded patient IDs that have LifetimeIDs comprising of more than 10 characters. This ensures that we only matched MRNs to LifetimeIDs that are composed of the same digits, but may have extra spaces, dashes or minor typos.

After incorporating the result of this regex analysis, we identified 41 additional (patient_ID, MRN) mappings leading to a total of 871 mappings found; this is comprised of 866 unique patient IDs and 832 MRNs. There are now 5 patient IDs which have 2 MRNs mapped to them; 31 MRNs which have 2 patient IDs mapped to them; and 4 MRNs which have 3 patient IDs mapped to them. There were 29 MRNs that were recorded with spaces in DWC and there were 6 MRNs with extra digits recorded (typos). This leaves 30 MRNs with no patient ID mapping.

Step 4. At this point, we wanted to ensure that the mappings that we have currently found are accurate and determine if any MRNs needed to be matched to additional patient_IDs. To do this, we compared the last names, bed spaces, and admit times that are stored in both the clinical logs and DWC. First, we verified that for each (patient_ID, MRN) mapping, the last name associated with the MRN in the clinical logs matched the last name associated with the patient ID in DWC. To account for formatting differences, we removed dashes, spaces, commas, numbers, and #s from the last names. Additionally, we noticed that last names in DWC may

contain additional text including "girl" or "boy"; alternatively, last names in DWC may be an abbreviated version of the last name in the clinical logs. Therefore, we considered the last names a match if the last name in the clinical logs was a substring of the last name in DWC or vice versa. Another consideration is the possibility for typos as the clinical staff inputted the last name into the monitor; we introduced an additional condition to determine if two last names match that used an edit distance calculation to quantify the similarity between two strings as the minimum number of operations (adds, deletes, or changes) needed to change one string into the other. If the edit distance between two last names was less than 3, then we considered this a match as we reasoned that this difference would be caused by a typo. Using these conditions, we determined that all 871 (patient_ID, MRN) mappings have matched last names.

Additionally, we want to check that for each (patient_ID, MRN) mapping, the bed spaces that the MRN visits in the clinical logs matches the bed spaces the patient_ID visits in DWC. We retrieved all the bed spaces visited in DWC and all the Reisman bed spaces visited in the clinical logs. Since the clinical logs are made up of daily snapshots, the clinical logs may not have records of all the bed spaces a patient visited; the logs may be missing bed spaces that the patient is in for less than 24 hours. Therefore, if the bed spaces recorded in the clinical logs are a subset of the bed spaces recorded in DWC, then we considered this a match. Based on this, we found that 20 MRNs had a patient recorded in more bed spaces in the clinical logs than in DWC with the associated patient ID. Since we knew that it was possible for an MRN to be mapped to multiple patient IDs, we assumed that these MRNs for which we don't have a complete bed space timeline may have additional patient IDs for which they should be mapped; these additional patient IDs should include the bed spaces that are currently missing from the DWC timeline. Therefore, we concluded that we may still need to map these 20 MRNs to other patient IDs.

Lastly, we wanted to check that for each (patient_ID, MRN) mapping, the admit date recorded for the MRN in the clinical logs matches the admit date for the patient_ID in DWC. As we have seen before, we know that all the patient IDs with

MRNs recorded are admitted to DWC (i.e. have an admit_state = 1 in the patient table). Therefore, we considered a patient ID's admit time in DWC to be the first time stamp when the patient ID's admit_state is 1. If the time between the admit time in the clinical logs is less than 6 hours before the patient ID's admit time in DWC, then we consider this a matched admission time. Otherwise, we figured that there is a discrepancy and the MRN may need to be mapped to additional patient_IDs to account for the missing time. Based on this, we found that 9 (patient_ID, MRN) mappings had a patient recorded in the clinical logs more than 6 hours before being recorded in DWC. Therefore, we concluded that we may still need to map these 9 MRNs to other patient IDs; 4 of these MRNs also were discovered to have an incomplete bedspace timeline as detailed above. Additionally, as aforementioned, we have 30 MRNs that don't have a patient ID mapping yet, leading to 55 total MRNs which still need to be mapped to patient IDs.

Step 5. Next, we matched the 55 remaining MRNs to patient IDs based on if the patient IDs visited the same bed spaces and may also have the same last name in DWC. This process can be broken down into multiple steps.

First, we generated the bed space timeline based on the clinical logs; this is composed of the first time and last time (an interval) that the MRN is recorded in each bed space.

Second, we iterated through the time intervals recorded for each bed space and query the patient table; for each time interval in the format (bed space, first recorded time, last recorded time), we recorded all the patient IDs that are recorded in that bed space from the MRN's first recorded time minus 2 days to the MRN's last recorded time plus 2 days. We included this 2 day cushion due to discrepancies that may occur due to the clinical logs being composed of only a daily snapshot or potential delayed monitor input. The result of this step is a list of patient IDs and which bed space from the clinical log bed space timeline they visited.

Third, we retrieve the patient IDs which are recorded in all of the bed spaces. For example, if the MRN visited 3 bed spaces in the clinical logs, then potential patient

ID matches are those which also visited all of those 3 bed spaces; we do not want to consider patient IDs which only visited 2 of those bed spaces. If there was only one patient ID which visited all the bed spaces during the correct time frame, then we mapped the MRN to this patient ID after verifying that the last name in the clinical logs matched the last name in DWC (using the same aforementioned considerations from Step 4). 25 of the 55 MRNs had only one patient ID which visited all the bed spaces during the correct time frame; the other 30 MRNs had multiple potential patient ID mappings. Out of these 25, 20 were found to have matched last names, 4 were found to not have a last name recorded in DWC, and 1 was found to have a mismatched last name. For the 4 patient IDs with no last name recorded in DWC, we still considered this a match. We discarded the (patient_ID, MRN) potential match which had a mismatched last name leading to this MRN with no patient_ID match. Overall, this led to 24 additional (patient_ID, MRN) mappings.

Otherwise, for the remaining 30 MRNs, 25 had patient IDs that could be potential mappings; 5 of these MRNs had multiple patient IDs that occurred across all the desired bed spaces. However, the majority (25 out of the 30) had no patient IDs that occurred across all the bed spaces; in this case, our list of potential patient IDs consisted of all the patient IDs that visited at least one of the desired bed spaces.

5 of the 30 MRNs had no patient IDs that occurred at any of the desired bed spaces during the desired time; 4 of these MRNs were already matched to a patient ID through Step 3, but were found to be incomplete; 1 MRN was unable to find any patient ID match. Therefore, for 4 MRNs, we are unable to find a complete timeline and may be missing data; this will be explored more when we compare length of stays in DWC to the clinical logs in Section 4.5.

For the 25 MRNs with multiple potential patient ID mappings, to determine which patient ID should be mapped to the MRN, we checked if the last name in DWC matched the last name in the clinical logs; if the last names matched, this was considered a (patient_ID, MRN) mapping. We were able to match 14 MRNs using this process leading to 14 additional (patient_ID, MRN) mappings. The 11 MRNs that were not able to be matched were already matched to a patient_ID in Step 3,

indicating that we were unable to find a complete timeline and may be missing data.

Through Step 5, we were able to generate 38 additional (patient_ID, MRN) mappings; we also noted that 9 of the 25 MRNs that were initially thought to have incomplete data in Step 4 were able to be matched to additional patient IDs. We ran a verification check similar to Step 4 to confirm that the number of bed spaces and the admit date from the clinical logs match DWC for the MRNs that had mappings generated by this step. We found that 3 MRNs were still recorded in more bed spaces in the clinical logs than DWC. We also found that 6 MRNs were still admitted to the clinical logs more than six hours before being admitted to DWC; 2 of these also were discovered to have an incomplete bedspace timeline. This indicated that there were 7 additional incomplete timelines along with the 15 incomplete timelines of MRNs that were already unable to be mapped during Step 5. This demonstrated that 22 MRNs (2.5% of the study cohort) still had incomplete timelines.

Step 6. At this point, we had found 909 total (patient_ID, MRN) mappings and outputted these mappings; this consisted of 860 unique MRNs and 902 unique patient IDs. There are now 7 patient IDs which have 2 MRNs mapped to them; 39 MRNs which have 2 patient IDs mapped to them; 5 MRNs which have 3 patient IDs mapped to them.

We found that 860 out of the 862 MRNs in our study cohort were able to be matched; 2 MRNs were unable to be matched to a patient ID. Both of these MRNs are recorded in the Reisman unit in the clinical logs for less than 3 days, which indicates that we are missing minimal data in DWC. Additionally, we saw that the (patient_ID, MRN) mappings were able to be verified by cross-referencing the last names recorded in DWC and the clinical logs. This indicated the validity of the DWC data and demonstrated that patients are recorded in DWC and can be tracked by MRN.

A flow chart of this mapping algorithm can be seen in Figure 4-7.

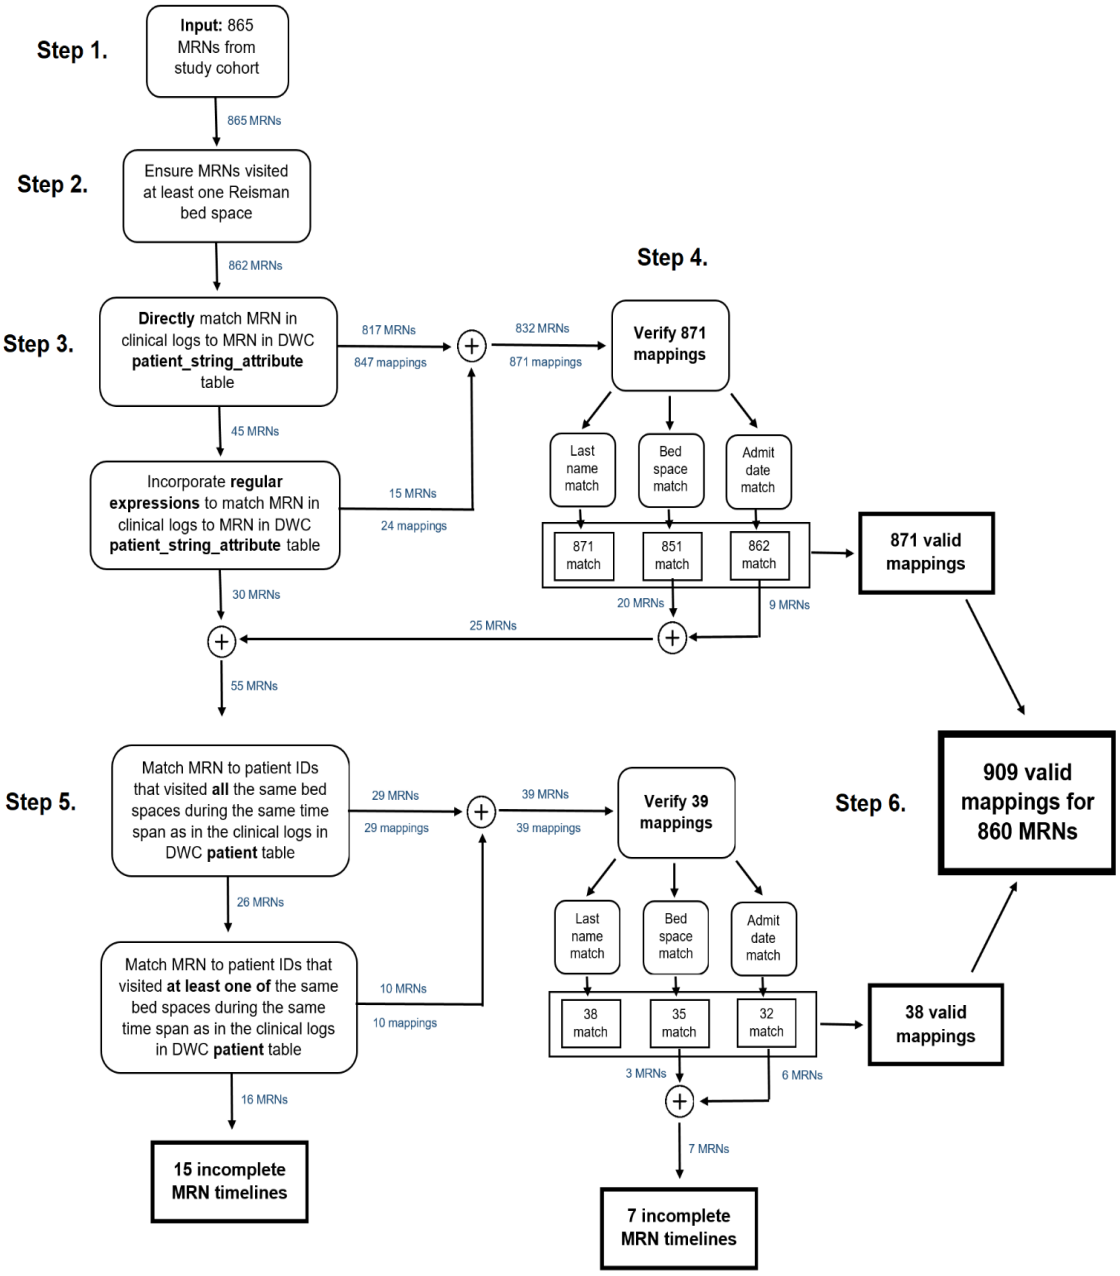

Figure 4-7: Flow chart describing the patient mapping algorithm.

4.5 Determining Expected Admit and Discharge Times

Through the patient mapping algorithm, we had determined all the (patient_ID, MRN) mappings, so next, we wanted to determine the expected admit and discharge time for each mapping so that we would know at what time the DWC data associated

with a patient ID should be corresponding to an MRN; this is especially important when a patient ID is mapped to multiple MRNs. This also allowed us to calculate the data coverage in DWC as we could compare the admit and discharge times in the clinical logs to the admit and discharge times in DWC. We also determined the extent of the incomplete data which may occur due to not recording the patient in each bed space or not recording the patient's correct admission date.

The admit and discharge times were calculated using the clinical logs. If the patient started their NICU stay in a bed in the Reisman unit, then the admit time is the admission time recorded in the clinical logs. Otherwise, the admit time is the timestamp associated with the first time the clinical logs records the patient in the Reisman unit. Similarly, if the patient ends their NICU stay in a bed in the Reisman unit, then the discharge time is the discharge time recorded in the clinical logs. Otherwise, the discharge time is the timestamp associated with the last timestamp that a patient is in the Reisman unit.

We also observed that there were 20 MRNs that had multiple admission and discharge dates in the clinical logs; this indicated that the patient may have been discharged from the NICU before returning at a later date. In this case, we created multiple rows for this MRN corresponding to each admission. For example, for a patient that is admitted twice to the NICU, their mapping may look like the following: (patient_ID_1, MRN_1, first_admission_date, first_discharge_date), (patient_ID_1, MRN_1, second_admission_date, second_discharge_date).

We also noticed that there are 24 MRNs that leave the Reisman unit and return to the Reisman unit at a later date. We handle this similarly to how we handle multiple admission and discharge dates and create multiple rows in the mapping corresponding to each stay in the Reisman unit; the admit and discharge times will be associated with the first and last time a patient is recorded in the Reisman unit.

These considerations allow us to map each (patient_ID, MRN) mapping to the associated clinical log admission and discharge time; this is useful as when we query SpO_2 data, for each MRN, we will retrieve all the data that occurs between this clinical log admission and discharge time.

4.6 Comparing Admission and Discharge Times

After determining the admission and discharge time of each MRN and associated patient ID, we determined the difference between the admission and discharge time from the clinical logs and the admission and discharge time found in DWC to see how complete the DWC data is.

4.6.1 Admission Time

The admission time of a patient ID in DWC was calculated to be the first timestamp associated with an admit_state of 1. For MRNs with multiple patient IDs, the DWC admission time was the earliest admission time of all the associated patient IDs. If a MRN is associated with a patient ID that is mapped to multiple MRNs, the DWC admission time was the time stamp that the desired MRN is recorded in the patient_string_attribute table; if the desired MRN is not recorded in the table, then the DWC admission time was the first timestamp that the patient ID had an admit_state of 1.

We compared the DWC admission time to the expected admit time and found that 48 MRNs had more time recorded in DWC while 812 MRNs had more time recorded in the clinical logs. Initially, this was concerning as it may indicate that we are missing patient data in DWC; however, we realized that these findings made sense as patients are unlikely to be admitted to a monitor and recorded in DWC immediately upon being admitted to the NICU. Therefore, a slight delay in DWC admission is expected.

For the 48 MRNs with more time recorded in DWC, we found the median and interquartile range for the excess time in DWC was 1.9 (0.13 - 23.97) hours; 11 MRNs had more than 1 day of excess data. For these 11 MRNs, we discovered two potential reasons for an early admission to DWC. One finding is that the first recorded bed space for 7 of these MRNs in the clnical logs are virtual bedspaces; when this occurs, the first timestamp that the MRN is recorded in a Reisman bed space is considered the expected admission time which may be a later estimate than what actually occurred

since the clinical logs only take daily snapshots. The other finding is that the other 4 MRNs have a date of birth that is up to 12 days before their expected admission date, so it is possible that the patient was admitted to a monitor and recorded in DWC after they were born but not recorded as admitted in the clinical logs. A distribution of the excess time in DWC can be seen in Figure 4-8.

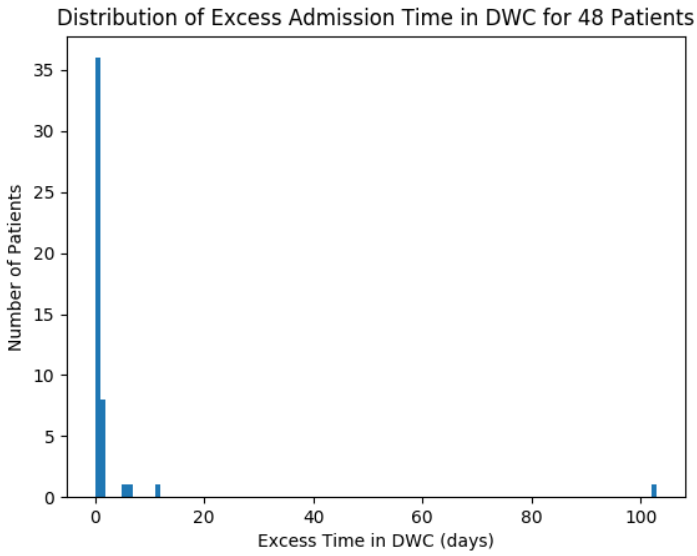

Figure 4-8: Distribution of excess admission time recorded in DWC.

For the 812 MRNs with more time recorded in the clinical logs, we found the median and interquartile range for the amount of time missing from DWC was 0.28 (0.15 - 0.64) hours; 7 MRNs were missing more than 1 day of data. For these 7 MRNs, we discovered two potential reasons for a late admission to DWC. One finding is that the expected admission date was a few days earlier than the timestamp of the first snapshot taken in the clinical logs for 1 of the MRNs; upon investigation, we determined that this patient was triaged on their admission date and then was not admitted to a NICU bed space until 3 days later indicating that it is correct that there is no data in DWC until their NICU admission. The other finding was that for 3 of the MRNs, in DWC, they do have rows recorded before the expected admission date; however, the admit_state equals 0. If we account for these rows, then we would

have excess time in DWC indicating that DWC may still be capturing relevant data such as SpO2 even if the patient didn't have an admit_state = 1 yet. For the 3 MRNs which have unexplained late admissions, we noted that 1 MRN was missing approximately 2 days of data, but had a NICU stay of approximately 85 days which indicated that we are missing a minimal amount of data. The other 2 MRNs are unaccounted for approximately 75 days leading to less desirable coverage; however, it was encouraging that only 2 MRNs (<1% of our study cohort) have a significant percentage of data missing from the start of a patient's stay. A distribution of the time missing from DWC can be seen in Figure 4-9.

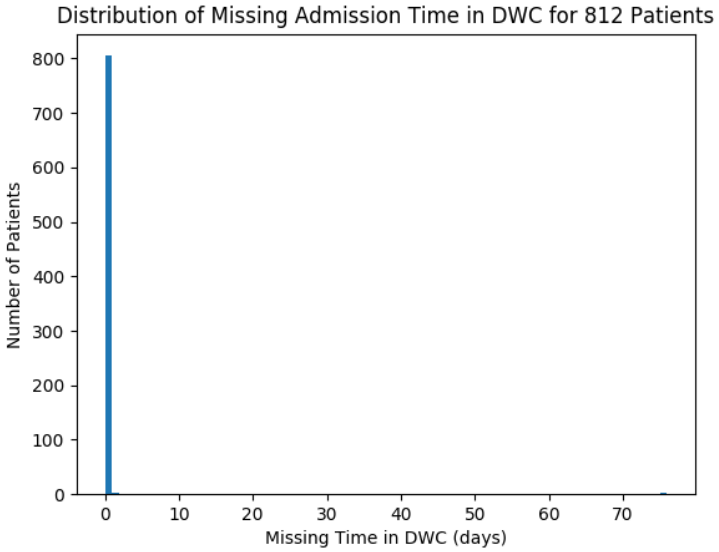

Figure 4-9: Distribution of missing admission time from DWC.

Through this analysis, the admission time from the clinical logs is consistent with the DWC admission time as we found the majority of patients have the DWC admission time within 1 day of the expected admission time. We were able to explain excess and missing data seen in DWC.

4.6.2 Discharge Time

Next, we calculated the discharge time of a patient ID in DWC. A discharge is indicated in the patient table when the admit_state = 2; however, we found that 148

patient IDs either never had an admit_state equal to 2 or there were more rows recorded after the admit_state is last set equal to 2. This indicates that it is not always possible to determine the exact discharge time based on admit_state. If the last time that a patient ID's admit_state is equal to 2 is the last record of the patient ID in DWC, then we considered this the DWC discharge time. Otherwise, the discharge time of a MRN was the time stamp that the next patient ID was recorded in the last bed space the current patient ID was recorded in.

For MRNs with multiple patient IDs, the DWC discharge time was the maximum discharge time of all the associated patient IDs. If a MRN is associated with a patient ID which is mapped to multiple MRNs, the DWC discharge time was the time stamp that the next MRN was recorded in the patient_string_attribute table; if the desired MRN is not recorded in the table or the desired MRN is the final recorded MRN in the patient_string_attribute table, then the DWC discharge time is either the last timestamp that the patient ID had an admit_state of 2 or the time stamp that the next patient ID was recorded in the last bed space the current patient ID was recorded in.

We compared the DWC discharge time to the expected discharge time and found that 693 MRNs had more time recorded in DWC while 167 MRNs had more time recorded in the clinical logs. We believed these findings made sense as a patient may have more time in DWC if the patient was not discharged from the monitor immediately upon leaving the NICU; additionally, as we have seen that the discharge time of the patient may not be the most reliable, we may be overestimating the discharge time when we set the DWC discharge time to be the time stamp that the next patient ID was recorded. Alternatively, a patient may have more time recorded in the clinical logs if the patient was discharged from the monitor prematurely. Therefore, slight discrepancies between the DWC discharge time and the expected discharge time is expected.

For the 693 MRNs with more time recorded in DWC, we found the median and interquartile range for the excess time in DWC was 16.20 (0.25 - 37.51) hours; 257 MRNs had more than 1 day of excess data. Excess data may occur due to reasons

explained above such as not discharging the patient from the monitor immediately upon leaving the NICU; discharging a patient from the monitor is not as urgent and it is not concerning to have excess data in DWC. We also are not concerned about excess data in DWC as when we query for SpO_2 we will be only taking values recorded between the admission and discharge time according to the clinical logs. A distribution of the excess time in DWC can be seen in Figure 4-10.

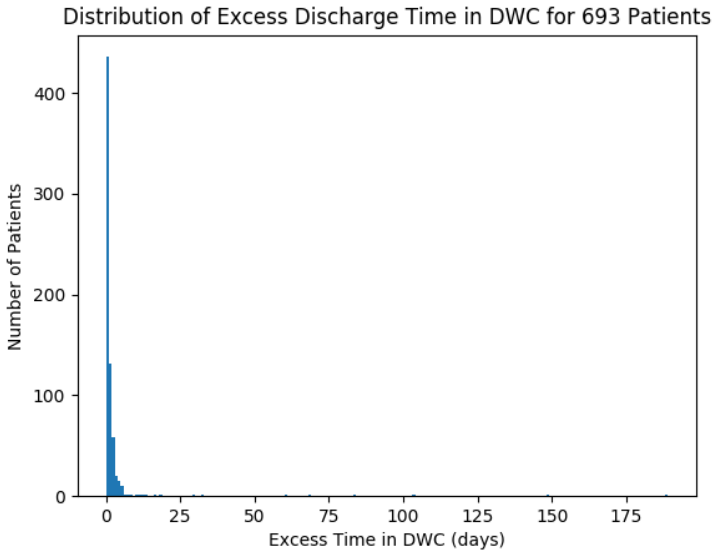

Figure 4-10: Distribution of excess discharge time from DWC.

For the 167 MRNs with more time recorded in the clinical logs, we found the median and interquartile range for the amount of time missing from DWC was 0.64 (0.25 - 1.59) hours; 4 MRNs were missing more than 1 day of data. One reason for this discrepancy is that 2 MRNs had a discharge date that were at most 4 days after the last timestamp recorded in the clinical logs; this indicates that these patients may not have been in the Reisman unit at the end of their stay and therefore, would not have any data in DWC. We noted that one MRN was missing approximately 1.5 days of data, but had a NICU stay of approximately 78 days which indicated that we are missing a minimal amount of data. The other MRN was also missing approximately 1.5 days of data, but had a Reisman NICU stay of approximately 5 days leading

to less desirable coverage; however, similar to analyzing the admission times, it was encouraging that only 1 MRN (<1% of our study cohort) had a significant percentage of data missing from the start of a patient's stay. A distribution of the missing time in DWC can be seen in Figure 4-11.

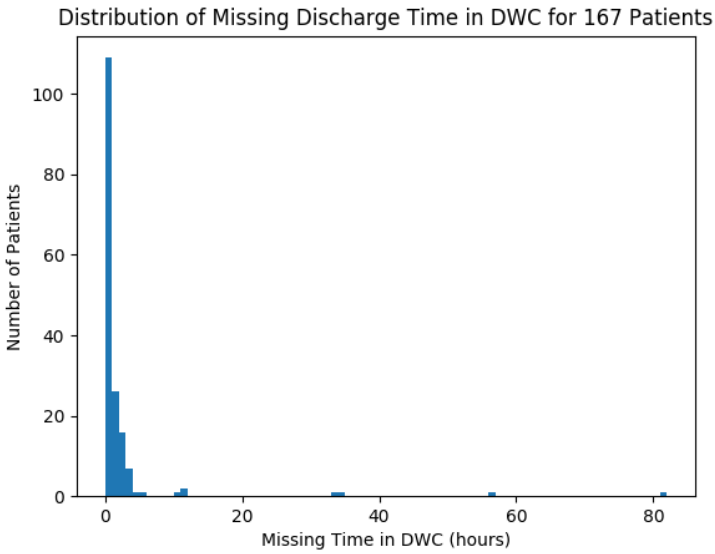

Figure 4-11: Distribution of missing discharge time from DWC.

Through this analysis, we found the discharge time from the clinical logs is relatively consistent with the DWC discharge time; however, we noted that we may have excess data in DWC due to how we are calculating DWC discharge times.

4.6.3 Length of Stay

We combined our analysis of admission and discharge times to determine the difference in the expected Reisman NICU length of stay and the DWC length of stay in order to estimate the amount of data coverage we have in DWC. The expected Reisman NICU length of stay is equal to the expected discharge time minus the expected admit time; similarly, the DWC length of stay is equal to the DWC discharge time minus the expected admit time.

We found that 660 MRNs had more time recorded in DWC while 200 MRNs had more time recorded in the clinical logs.

For the 660 MRNs with more time recorded in DWC, we found the median and interquartile range for the excess time in DWC was 20.21 (4.65 - 46.55) hours; 287 MRNs had more than 1 day of excess data (due to having more than 1 day of admission or discharge time recorded in DWC). A distribution of the excess time in DWC can be seen in Figure 4-12.

For the 200 MRNs with more time recorded in the clinical logs, we found the median and interquartile range for the amount of time missing from DWC was 1.06 (0.46 - 2.30) hours; 6 MRNs were missing more than 1 day of data (due to missing more than 1 day of admission or discharge time). A distribution of the excess time in DWC can be seen in Figure 4-13.

This analysis ensures the validity of DWC data as we can see that for the majority of MRNs, DWC is not missing a significant amount of time compared to the clinical logs; however, DWC does have significantly more data corresponding to each MRN which is likely due to the way we are calculating DWC discharge times.

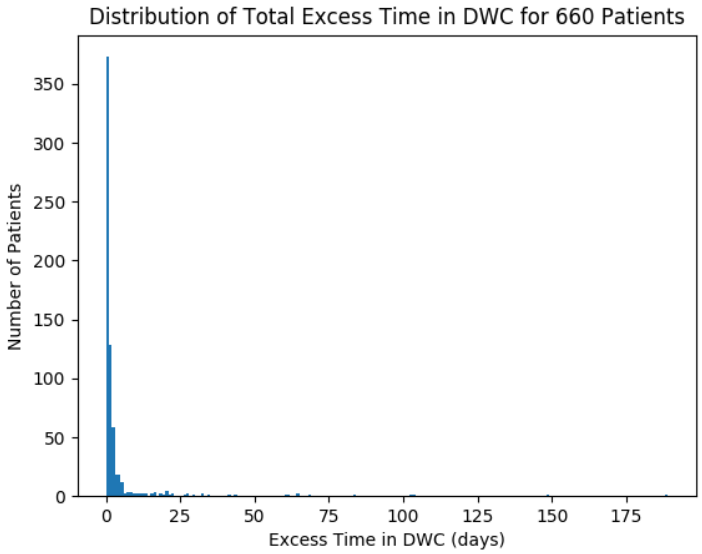

Figure 4-12: Distribution of excess time from DWC.

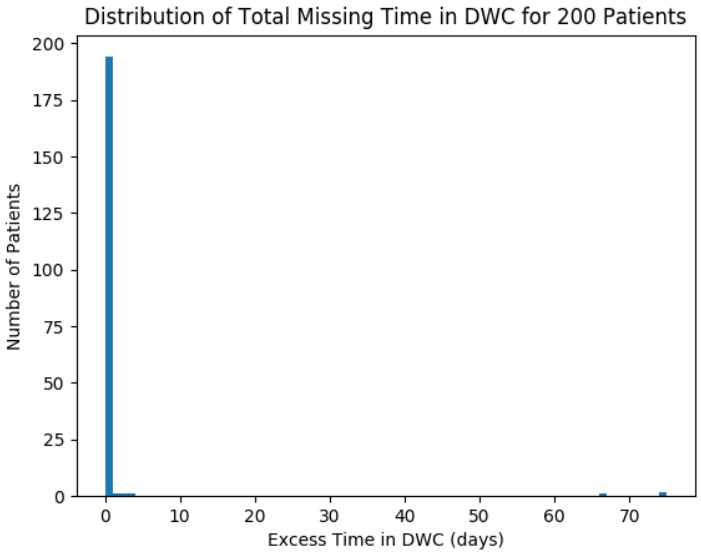

Figure 4-13: Distribution of missing time from DWC.

4.7 Unmatched Medical Record Numbers

As we have found that the vast majority of the MRNs in our study cohort were able to be mapped to patient IDs, we also wanted to explore the other MRNs that are recorded in DWC and the clinical logs. As aforementioned, we found 2219 unique MRNs recorded in DWC; we also have 1633 unique MRNs in the clinical logs. We can see that there are more MRNs in DWC and we also noted that the developed patient mapping algorithm only maps the 860 MRNs in our study cohort. To continue to confirm the validity of the DWC data, we first attempted to verify the remaining 773 MRNs in the clinical logs can be mapped to patient IDs in DWC; next, we explored the remaining MRNs that are recorded in DWC, but aren't found in the clinical logs.

4.7.1 Additional MRNs in Clinical Logs

We determined that 644 of the 773 remaining MRNs (83.3%) in the clinical logs had at least one patient ID mapped to it. The median and interquartile range of the missing MRNs' gestational age is 39.4 (38.4 - 40.2) weeks. The median and interquartile range of the missing MRNs' length of stay is 3.43 (3.03 - 4.42) hours. This indicates that we are mainly missing patients that are not born premature and who stay in the hospital

for a minimal amount of time. These patients are likely not recorded in DWC as the clinical staff did not input the patient information into the monitor as the patient was only temporarily staying in the NICU before being discharged.

However, we noted that there were 12 preterm (<37 weeks) patients that were unable to be mapped to a patient ID; the minimum gestational age missing a mapped patient ID was 24.6 weeks. 7 of these were recorded only in virtual bedspaces (including the patient with a gestational age of 24.6 weeks); 1 of these was only recorded in a non-Reisman bed space; 1 was not discharged by September 11, 2019. The two remaining MRNs truly could not be found in DWC; their gestational ages were 34.6 and 36.1 weeks.

Additionally, there were 3 patients with a length of stay of more than 24 hours that were unable to be mapped to a patient ID; the maximum length of stay missing a mapped patient ID was approximately 11 days. The gestational age of these patients are 38.5, 38.1, and 40.6 weeks. We manually checked the presence of these patients in DWC and found that there are potential patient ID mappings, but they are unable to be verified through the recorded MRN and last name.

Through this, we concluded that the majority of MRNs in the clinical logs are recorded in DWC. The 129 MRNs that cannot be found in DWC are mainly not preterm patients with short length of stay.

4.7.2 Additional MRNs in DWC

There are 709 MRNs in the DWC patient_string_attribute table that are not in the clinical logs. The median and interquartile range of the associated patient ID's DWC length of stay is 15.51 (4.00 - 57.03) hours; 318 of these MRNs are recorded in DWC for more than one day. This indicates that somehow there are extra MRNs recorded in DWC that are not captured by the clinical logs.

Since the clinical logs only contain daily data snapshots, we introduced another data source known as the Admit/Discharge/Transfer (ADT) logs to verify the additional MRNs in DWC. ADT is able to provide a more comprehensive view of which patient occupies which bed space in the NICU at every point in time; it contains

columns recording the MRN, timestamp, type of transaction (admit, discharge or transfer), and bed space. 548 of the 709 MRNs not found in the clinical logs are found in ADT indicating they are valid.

This indicated that 161 MRNs recorded in DWC are unable to be found in the clinical logs or ADT. The median and interquartile range of the associated patient ID's DWC length of stay is 32.21 (4.20 - 139.70) hours; 105 of these MRNs are recorded in DWC for more than one day. We focused on these 105 MRNs as these are the patients that could have been included in our study cohort. Analyzing these MRNs, we found that 55 have incorrect number of digits for a MRN; a MRN should have 7 digits. MRNs with an incorrect number of digits (i.e. 007, 000, 02000) would not be found in the clinical logs or ADT and are likely placeholder MRNs. Additionally, 5 MRNs consisted of only one digit (i.e. 0000000) and are also likely placeholder MRNs. We also found that 16 MRNs had an associated first name or last name that indicated that the corresponding patient ID was for a car seat test and would likely have a placeholder MRN. Finally, 15 MRNs had another MRN recorded in the patient_string_attribute table for the same patient ID; in this case, we assumed that the clinical staff may have first inputted a placeholder MRN and changed it to the correct MRN at a later time. Therefore, we are left with 24 potentially valid MRNs that are not found in the clinical logs or ADT.

For these 24 potentially valid MRNs, we explored the possibility that the MRN recorded in DWC had a typo and we attempted to match patients based on bed spaces similar to Step 5 in the patient mapping algorithm. We were successful in matching all the patient IDs associated with the potentially valid MRNs to the correct MRN in the clinical logs indicating that the MRNs in the patient_string_attribute table may not always be correct and are susceptible to typos.

A flow chart of this process can be seen in Figure 4-14.

Overall, we were able to account for all the MRNs recorded in DWC and indicate why some may not be found in the clinical logs or ADT. This is encouraging in demonstrating the completeness and correctness of the data in DWC. While the MRN recorded in the patient_string_attribute table may be incorrect, the associated data

in the patient table (such as bed labels and time stamps) are correct and can be used to verify that the data recorded in DWC matches patients recorded in the clinical logs or ADT.

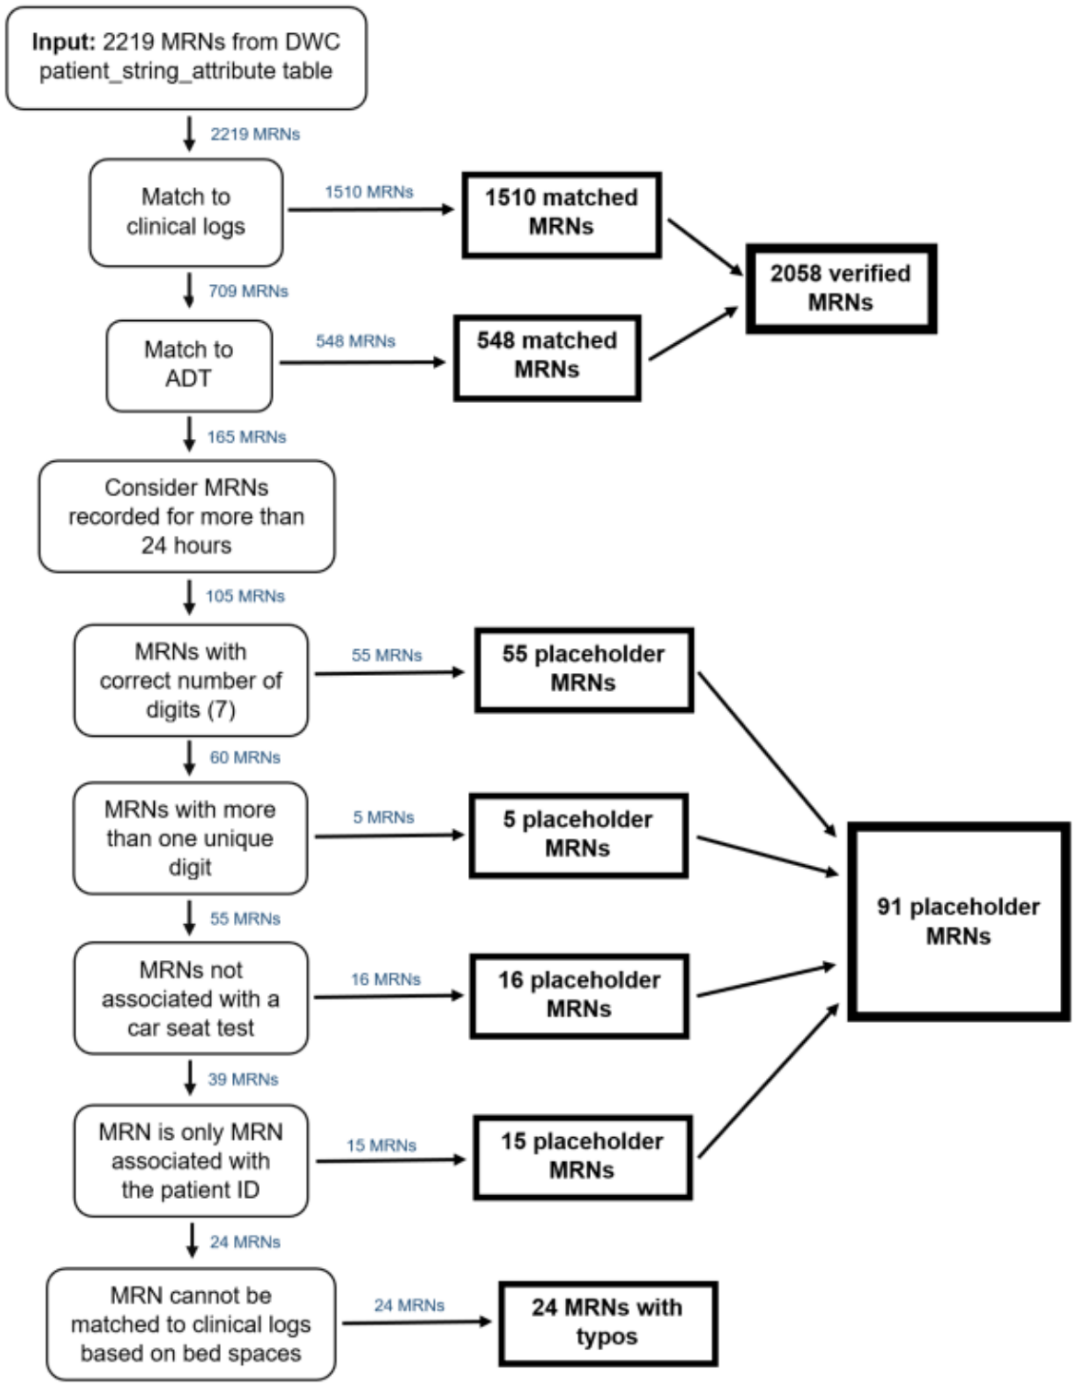

Figure 4-14: Flow chart illustrating how we verified all the MRNs in DWC.

4.8 Summary of Findings

Through multiple validity checks, we were able to confirm that the data recorded through DWC is correct, complete and clean for the most part. First, we assessed the composition of the patient and patient_string_attribute table in DWC to determine the number of Patient IDs, MRNs, first names, last names and bed spaces that we have access to. We found 2189 valid patient IDs that were admitted to the NICU between January 1, 2018 and September 11, 2019 and the majority of these had either an MRN, first name and/or last name recorded. The recorded bed spaces allowed us to recreate a timeline of a patient's stay and identify when each patient occupies various bed spaces.

Next, we matched the MRNs recorded for our study cohort in the clinical logs to the DWC patient IDs to ensure that the data in DWC represented the data in the clinical logs. By comparing last names, bed spaces visited, and admit date, we were able to confirm that 860 out of the 865 MRNs in our study cohort were able to be mapped and validated. Using this mapping, we compared admit dates, discharge dates, and length of stays and found that the majority of patients are not missing a significant amount of data in DWC when compared to the clinical logs. DWC does usually have more data recorded for patients; however, since we will be using the admission and discharge times from the clinical logs to determine what data (such as SpO_2) we will be pulling from DWC, this does not present a major issue.

Finally, we made sure that any MRNs in either the clinical logs or DWC that were not matched yet by this point were accounted for. We found that the majority of the additional MRNs in the clinical logs could be matched in DWC; the MRNs that were unable to be matched were mainly patients born at term or who stayed in the hospital for less than 24 hours. All of the additional MRNs in DWC were able to be accounted for.

Overall, these checks were able to demonstrate the coverage and validity of the data collected in DWC which allowed us to continue our analysis and begin incorporating SpO_2 data for the 860 patients found in DWC. This confirms that we can

leverage this large data set to expand on the reviewed literature discussed in Chapter 2. Additionally, since DWC provides real time data streams, this data set is constantly growing and the analyses run in this study can be continuously rerun leading to a bigger study cohort every day.

Chapter 5

Oxygen Saturation Data

Through the DWC numerics table, we have access to massive amounts of oxygen saturation (SpO_2) data that are sampled second by second for each patient in the BIDMC NICU. We leveraged this data to determine the percent of time patients are within, below, and above their target SpO_2 range and what factors may be affecting this compliance.

5.1 Data Coverage

Initially, we assessed the coverage of the SpO_2 data recorded in DWC from January 1, 2018 to September 12, 2019 for our study cohort. For each MRN, we determined how much SpO_2 data we had from the MRN's admission date to discharge date as calculated in Section 4.5. To do this, we queried the numerics table for all SpO_2 values taken between the MRN's admission and discharge time and grouped these values by second to ensure that we had only one value collected per second. The % SpO_2 data coverage was determined as

$$\frac{\text{number of seconds with } SpO_2 \text{ recorded in DWC}}{\text{number of seconds between admission and discharge date recorded in clinical logs}}$$

. Therefore, if there is no SpO_2 value associated with a second, then this second would be considered missing. This was calculated at the NICU level as well as the

patient level.

We found that there was approximately 90.5% SpO_2 NICU level data coverage. At a patient level, we had 87.9% ($\pm$15.4%) coverage indicating that from a patient's admit time to discharge time we have second-by-second SpO_2 data for 87.9% of a patient's stay on average. Figure 5-1 illustrates a breakdown of this data coverage. We also graphed the relationship between data coverage and gestational age, birth weight, and length of stay, as seen in Figure 5-2, and we can see that we have better data coverage for patients that are born more premature and stay in the NICU for longer.

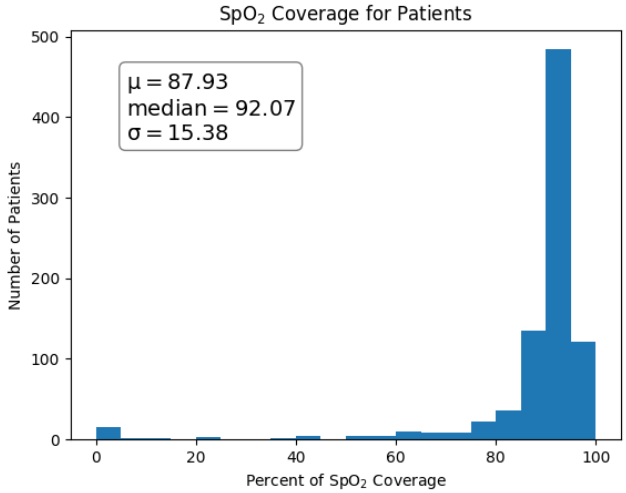

Figure 5-1: Distribution of coverage of SpO_2 data across patients.

We found that 254 patients had less than 90% data coverage, and we found that 9 patients had no SpO_2 data recorded. It is unclear why this 0% coverage occurred, but none of these 9 patients were extremely preterm or extremely low birth weight.

(a)

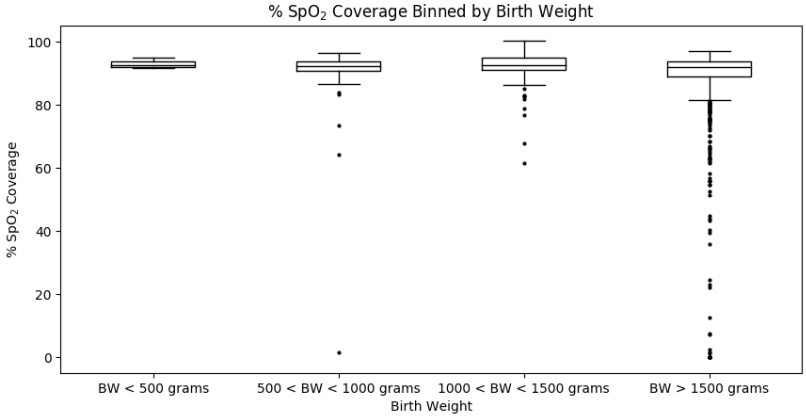

(b)

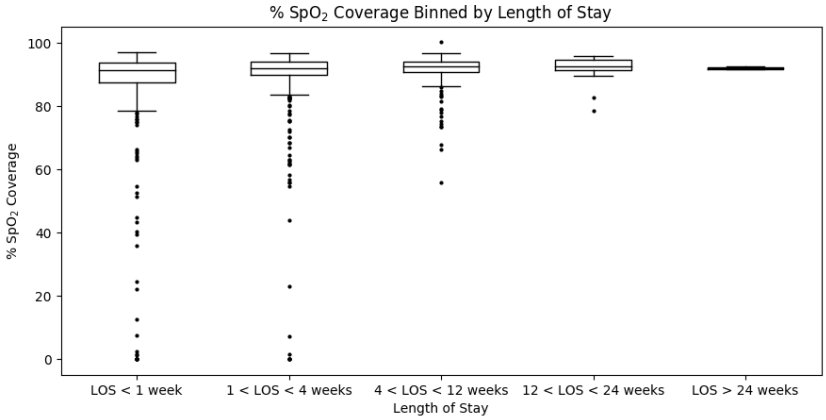

(c)

Figure 5-2: Relationship between percent SpO_2 coverage and (a) gestational age, (b) birth weight, and (c) length of stay of study cohort.

5.2 Data Validity

We also assessed the validity of the SpO_2 data recorded in DWC; when the SpO_2 drops too low, it may be unclear if it is still a valid measurement as either a sensor may have came off or there may be some introduced noise. An example of a sample day of SpO_2 data can be seen in Figure 5-3; through this, we can see many fluctuations in SpO_2, and we aimed to confirm that this data was usable before we began our analysis.

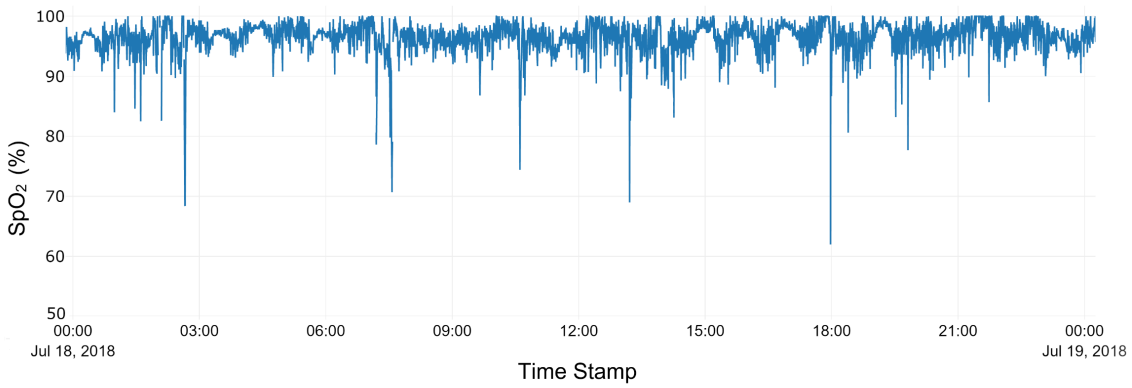

Figure 5-3: Example SpO_2 record for a patient in DWC over the course of a day.

To determine the validity of the SpO_2 data, we considered two different approaches. The first approach was analyzing the heart rate and pulse rate measurements and confirming that these measurements were within 10 beats of each other when averaged across 3 seconds. The heart rate is derived from the ECG and the pulse rate is derived from the PPG. When the signal quality of both of these measurements is high, the two values should agree. The second approach was considering if there was any missing time from the beginning of an episode when the SpO_2 dipped below the target SpO_2 range.

5.2.1 Matching Heart Rate and Pulse Rate Measurements

In the NICU, we observed that heart rate (HR) and pulse rate (PR) measures were usually less than 10 beats per minute (BPM) of each other when the SpO_2 was clean;

we attempted to incorporate this into a quality check that defined clean SpO_2 data as occurring when HR and PR were within 10 BPM averaging across 3 seconds so that we disregard potential noise.

First, we investigated the coverage of the HR and PR measurements to determine if DWC collected enough data for our purposes. The % HR data coverage was determined to be equal to

$$\frac{\text{number of seconds with HR recorded in DWC}}{\text{number of seconds between admission and discharge date recorded in clinical logs}}$$

Similarly, the % PR data coverage was determined to be equal to

$$\frac{\text{number of seconds with PR recorded in DWC}}{\text{number of seconds between admission and discharge date recorded in clinical logs}}$$

We found that, at the NICU level, there was approximately 90.6% HR data coverage and approximately 90.5% PR data coverage which is similar to our SpO_2 data coverage. We also calculated how often we have both HR and PR values recorded on a second-by-second basis and determined that this occurred 88.3% of the time at the NICU level. Distributions of this coverage can be seen in Figure 5-4.

Next, we investigated how often the HR and PR measurements were within 10 BPM of each other averaging across 3 seconds which we hypothesized would imply clean SpO_2 data. This occurred 77.7% of the time and indicated that approximately 85.9% of the SpO_2 data we had collected would be considered valid. However, since our original NICU level SpO_2 coverage found in Section 5.1 was 90.5%, removing SpO_2 data which did not have matching HR or PR resulted in 77.6% SpO_2 valid coverage which was not ideal.

After evaluating these results, we decided to examine a sample timeline and analyze the SpO_2, HR, and PR values to determine if this approach made sense. We found several examples from a patient where HR and PR either deviated or were consistent, and explored if this had any indication of the validity of the SpO_2 data;

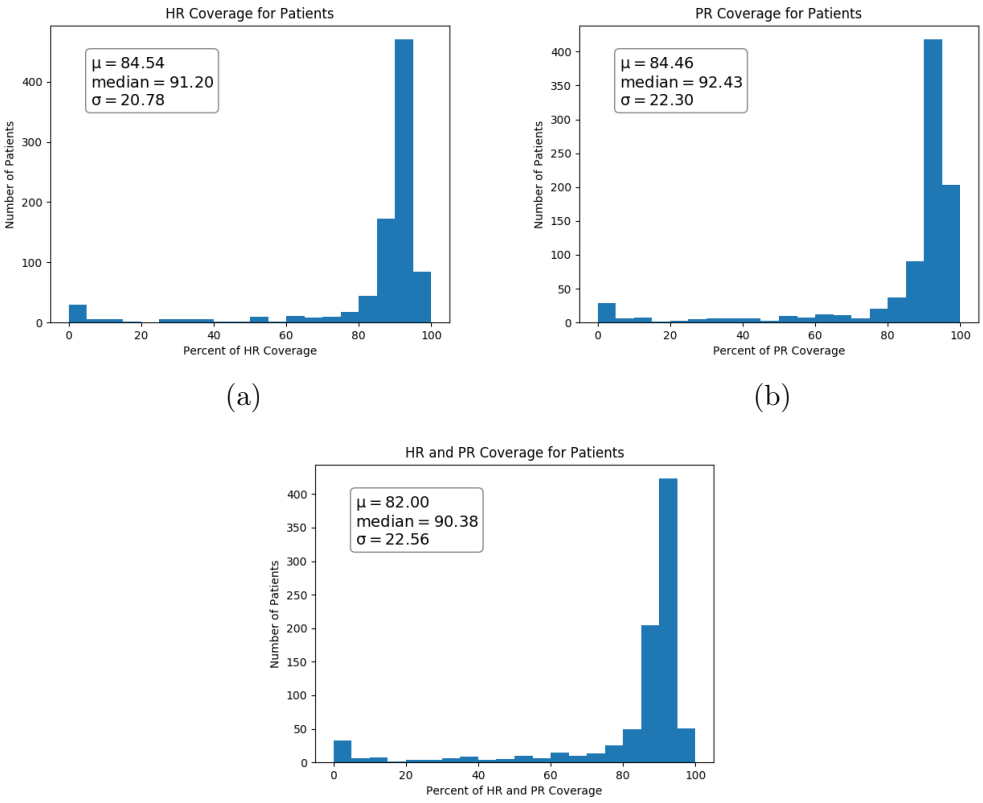

Figure 5-4: Distribution of % coverage of (a) HR, (b) PR, and (c) both HR and PR.

the patient we analyzed had a gestational age of 29 weeks and a length of stay of 90 days.

In Figure 5-5, we observed that HR and PR were within 10 BPM for the majority of this sample even when there were dips in SpO_2. However, around 20:30, as illustrated by the red box, we observed that HR significantly dropped lower than PR; with our current approach, this would invalidate the SpO_2 data. However, the SpO_2 at this time looks valid and no noise seemed to be introduced. This began to indicate that this approach may not be the most appropriate quality check.

Figure 5-6 provides another example of HR and PR significantly deviating (this time with PR significantly dropping lower than HR) even though the SpO_2 signal detailed a seemingly clean hypoxemic episode. Through these two examples, we began to question if we should pursue this quality check.

In Figure 5-7, we observed a few examples where HR and PR deviate by more

than 10 BPM. When this occurred, the SpO_2 seemed to remain clean and valid except for around 9:25, as illustrated by the red box. In this case, there were two consecutive potential hypoxemic episodes, but the second episode was missing some SpO_2 data and may be invalid. Figure 5-8 illustrates another example of missing significant amounts of SpO_2 data during a potential hypoxemic episode. These examples suggested that an alternative quality check could be disregarding hypoxemic episodes that are missing some of their start time as this may be a better indicator of noise introduced in the SpO_2 data.

Additionally, through this exploration, we observed that in this patient, the HR lagged a few seconds behind the PR; this can be seen in Figure 5-5. We shifted the HR forward multiple different values from 0 to 20 seconds and calculated the correlation coefficient between PR and HR for a sample hour at each of these different shifts to determine where the PR and HR best match. For this patient, we determined the optimal shift was approximately 8 seconds as this led to a peak correlation coefficient.

We sampled several hours from the beginning, middle, and end for this single patient's stay and determined that the optimal shift was consistent at 8 seconds; however, when we sampled 5 other patients, we found that the optimal shift changed to various values ranging from 0 seconds to 13 seconds. Initially, we thought that this could be due to different settings for each bedside monitor which would lead to patients in different bed spaces having different optimal shifts; however, we discovered that even patients that were recorded in the same bed space had different optimal shifts.

Through this analysis, we decided that the HR and PR measurements had an inconsistent shift across patients which made it difficult to compare; we also decided that this approach could lead to us throwing out too much potentially valid SpO_2 data. Therefore, we decided to explore an alternative approach to determine if the SpO_2 data is valid that focused on time periods with missing SpO_2 data.

While we did not use this approach in the end to validate SpO_2 data, this analysis was important to demonstrate the coverage of HR and PR in DWC as well as the relationship between HR and PR.

Figure 5-5: Sample SpO_2, HR, and PR values over 30 minutes. The red box illustrates an example when HR and PR deviate by more than 10 BPM.

Figure 5-6: Sample SpO_2, HR, and PR values over 20 minutes. The red box illustrates an example when HR and PR deviate by more than 10 BPM.

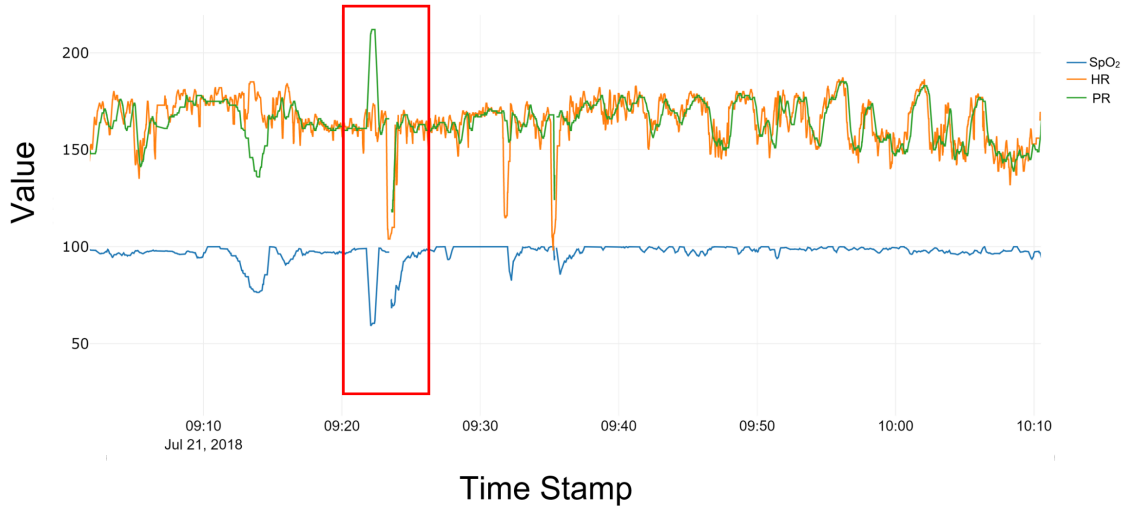

Figure 5-7: Sample SpO_2, HR, and PR values over 60 minutes. The red box illustrates an example when HR and PR deviate by more than 10 BPM.

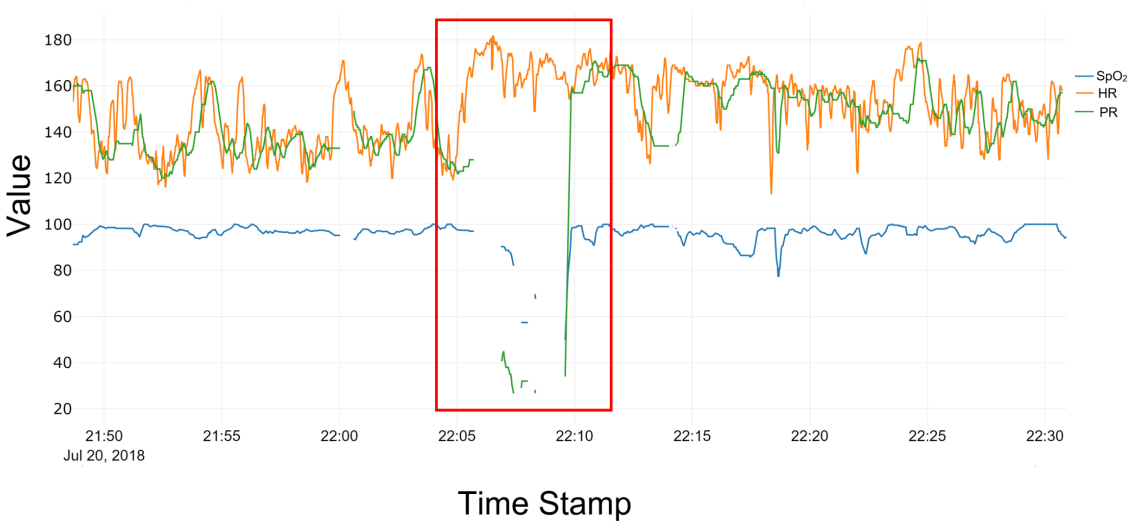

Figure 5-8: Sample SpO_2, HR, and PR values over 40 minutes. The red box illustrates an example when we are missing significant amounts of data.

5.2.2 Missing Time

As we saw in Section 5.2.1, a more telling sign that the SpO_2 values may be invalid is missing data. If a patient's SpO_2 dips below their target SpO_2 range, but they are missing time from the beginning of this episodes, then we considered that these episodes may be invalid. Examples of this were seen in Figure 5-7 and 5-8.

We determined that if a patient was missing more than 2 seconds of the start time of the episode, then we would discard this episode and not count it towards the total time spent below the target SpO_2 range. We chose this arbitrary threshold of 2 seconds as this ensured that we have coverage for the beginning of the hypoxemic episode; this threshold can be adjusted in the future. Through this, we discovered that we will be disregarding approximately 52 minutes of data for each patient on average; this accounts for approximately 0.1% of an average patient's length of stay and indicates that 3.7% of our study cohort's total time spent below their target SpO_2 range would be considered invalid. In terms of total NICU level coverage, disregarding this potentially invalid SpO_2 leads to an overall SpO_2 coverage of 90.3%. Therefore, this approach enables us to continue to have high SpO_2 coverage while ensuring the cleanliness of our data and we decided to use this approach throughout our analysis.

5.3 Target SpO_2 Range Compliance

After exploring the SpO_2 data coverage and validity of DWC, we explored the target SpO_2 range compliance and the percent of time that patients were recorded within, below and above their target SpO_2 range. At BIDMC, the clinicians determine a patient's target SpO_2 range based on the patient's post menstrual age (PMA) and whether the patient is on supplemental oxygen or not as displayed in Figure B-1. We used these ranges to determine how long patients in our study cohort spent in their target SpO_2 range based on their PMA (calculated based on the gestational age recorded in the clinical logs and the time stamp that the SpO_2 value was recorded) and the supplemental oxygen (determined based on the FiO_2 information in the clinical logs).

As we had 851 patients with SpO_2 data recorded in DWC, we analyzed these patients and broke down the time that they spent within, below and above their target SpO_2 range. The % time spent within the target SpO_2 range was defined as

$$\frac{\text{number of seconds with valid } SpO_2 \text{ recorded in DWC within the target range}}{\text{number of seconds with valid } SpO_2 \text{ recorded in DWC}}$$

The % time spent below the target SpO_2 range was defined as

$$\frac{\text{number of seconds with valid } SpO_2 \text{ recorded in DWC below the target range}}{\text{number of seconds with valid } SpO_2 \text{ recorded in DWC}}$$

Similarly, the total % time spent above the target SpO_2 range was defined as

$$\frac{\text{number of seconds with valid } SpO_2 \text{ recorded in DWC above the target range}}{\text{number of seconds with valid } SpO_2 \text{ recorded in DWC}}$$

We calculated that, at the NICU level, our study cohort spent a total of 90.0% of their collective stays within their target SpO_2 range; at the patient level, the median and interquartile range for an individual patient's % time within their target SpO_2 range was 98.3 (95.6 - 99.3)%. While this is higher than reported in the reviewed literature from Section 2.3, we note that our study cohort is composed of all preterm infants rather than only extremely preterm or extremely low birth weight infants on supplemental oxygen which would yield different results.

We calculated that, at the NICU level, our study cohort spent a total of 4.9% of their collective stays below their target SpO_2 range; at the patient level, the median and interquartile range for an individual patient's % time below their target SpO_2 range was 1.5 (0.7 - 3.7)%.

We calculated that, at the NICU level, our study cohort spent a total of 5.1% of their collective stays above their target SpO_2 range; at the patient level, the median and interquartile range for an individual patient's total % time above their target SpO_2 range was 0 (0 - 0.21)%. We note that this considers patients that are not on supplemental oxygen which leads to a skew toward 0%.

Distributions of the % time an individual patient spent within, below, and above

their target SpO_2 range can be seen in Figure 5-9.

When we only consider the 406 patients on supplemental oxygen, we calculated that this subset spent a total of 7.5% of their collective stays above their target SpO_2 range; the median and interquartile range for an individual patient's total % time above their target SpO_2 range was 0.3 (0.03 - 3.8)%.

We also noted that when patients are not on supplemental oxygen, their SpO_2 upper limit remains at 100%; therefore, patients can only be above their target SpO_2 range when they are on supplemental oxygen. Therefore, to determine a more accurate representation of the % time spent above the target SpO_2 range, we took a subset of patients that were on supplemental oxygen and considered an alternative measurement of true % time spent above the target SpO_2 range which we defined as as

$$\frac{\text{number of seconds with valid } SpO_2 \text{ recorded in DWC above the target range}}{\text{number of seconds with valid } SpO_2 \text{ recorded in DWC while the patient is recorded on } FiO_2 \text{ in the clinical logs}}$$

Therefore, we recalculated the true % time spent above the target SpO_2 range for our subset of 406 patients on supplemental oxygen; 402 of these patients had SpO_2 recorded. We calculated that these patients spent 24.0% of their collective stays while on supplemental oxygen above their target SpO_2 range; the median and interquartile range for an individual patient's total % time above their target SpO_2 range was 16.8 (3.4 - 33.8)%. Distributions of the % time an individual spent above their target SpO_2 range for this subset of 406 patients can be seen in Figure 5-10.

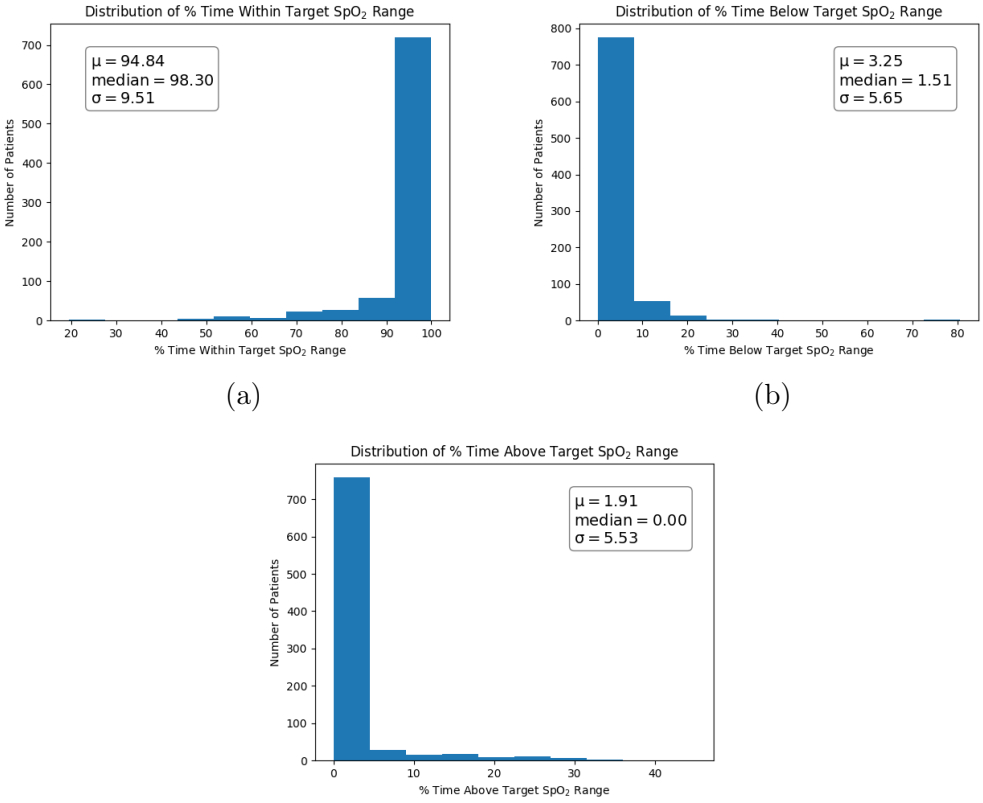

Figure 5-9: Distribution of % time an individual patient spent (a) within, (b) below, and (c) total above their target SpO_2 range.

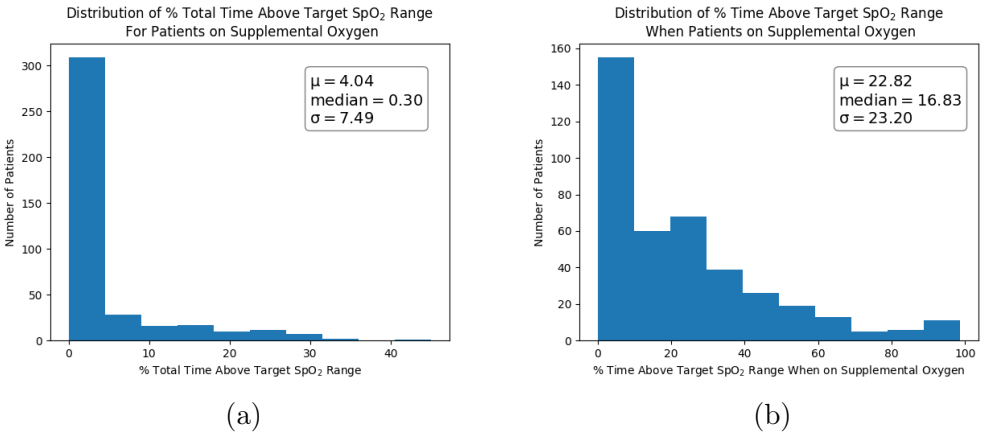

Figure 5-10: Distribution of % time an individual patient spent above their target SpO_2 out of (a) their total stay and (b) their total time on supplemental oxygen. This distribution represents a subset of patients that are recorded on supplemental oxygen at some point during their NICU stay.

5.4 Potential Factors Affecting Target SpO_2 Range Compliance

After determining the % time spent within, below and above the target SpO_2 range, we stratified the patients by various factors to gain insight into what may be affecting the target SpO_2 range compliance. These factors included gestational age, birth weight, gender, supplemental oxygen, postmenstrual age, time of day and bed space assignment.

5.4.1 Gestational Age

Regarding gestational age (GA), we categorized the 851 patients in our study cohort with SpO_2 data recorded into GA $<$ 28 weeks, 28 $<$ GA $<$ 30 weeks, 30 $<$ GA $<$ 32 weeks, 32 $<$ GA $<$ 35 weeks, and 35 $<$ GA $<$ 37 weeks; the number of patients in each of these categories can be seen in Figure 5-11. Patients less than 28 weeks are considered extremely preterm; patients between 28 and 32 weeks are very preterm; and patients between 32 and 37 weeks are moderate to late preterm [13]. We analyzed the % time spent within, below, and above the target SpO_2 range when patients are categorized by GA and these can be viewed in Figure 5-12; additionally, a table breaking down this analysis can be seen in Table 5.1.

For our % time above the target SpO_2 range calculation, we also considered the subset of 402 patients on supplemental oxygen; the number of these patients can also be seen in Figure 5-11. For this subset, we calculated the % total time above their target SpO_2 range as well as the % time above their target SpO_2 range when they are recorded on supplemental oxygen. This analysis can be viewed in Figure 5-13; additionally, a table breaking down this analysis can be seen in Table 5.2.

Through this, we observed that comparatively, extremely preterm ($<$28 weeks) infants spent less time within their target SpO_2 range. The % time within the target SpO_2 range for these patients seems to be higher than studies explored in Section 2.3 that focus on extremely preterm patients; however, these studies were focused

Table 5.1: Breakdown of % time individual patients spent within, below, and total above their target SpO_2 range binned by GA. Data are medians and interquartile ranges.

GA Category	Number of Patients	% Within	% Below	% Total Above
GA < 28 weeks	63	72.0 (59.2 - 88.5)	11.2 (5.7 - 16.6)	15.5 (0.4 - 24.3)
28 < GA < 30 weeks	62	98.8 (96.5 - 99.5)	1.1 (0.4 - 2.4)	0.1 (0.0 - 1.1)
30 < GA < 32 weeks	77	99.2 (97.8 - 99.6)	0.7 (0.4 - 1.6)	0.0 (0.0 - 0.2)
32 < GA < 35 weeks	390	99.1 (97.9 - 99.5)	0.9 (0.5 - 1.9)	0.0 (0.0 - 0.1)
35 < GA < 37 weeks	259	96.6 (94.1 - 98.2)	3.2 (1.7 - 5.2)	0.0 (0.0 - 0.2)

on patients on supplemental oxygen for limited periods of time. Our sub-analysis of extremely preterm patients on supplemental oxygen indicates these patients spend around 30% of their time above the target SpO_2 range when on supplemental oxygen; this aligns with the literature.

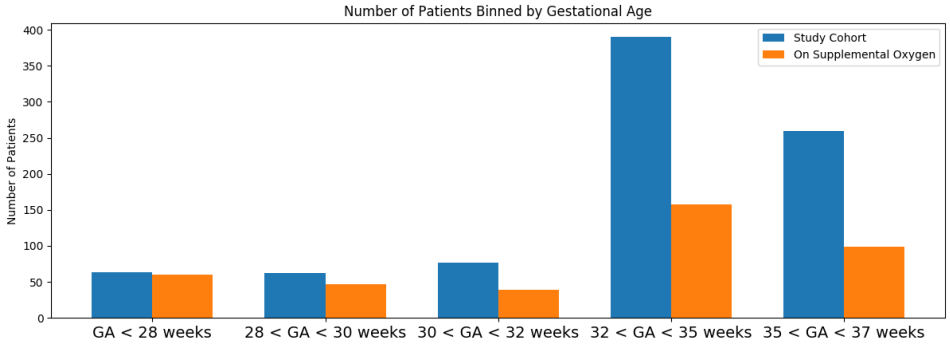

Figure 5-11: Number of patients in each GA category.

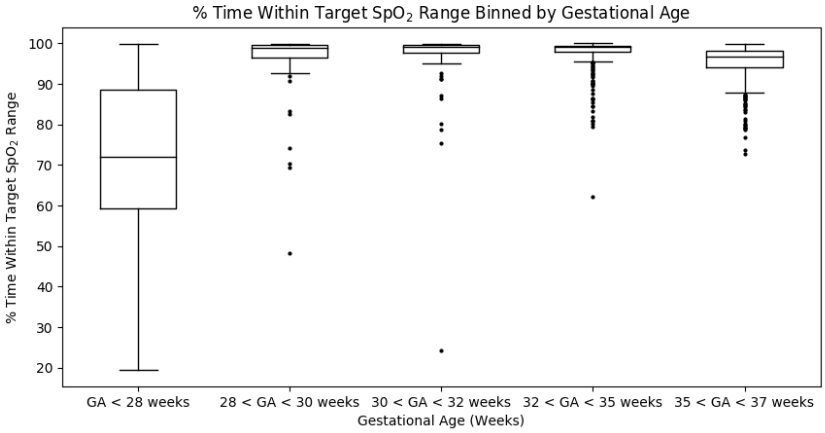

(a)

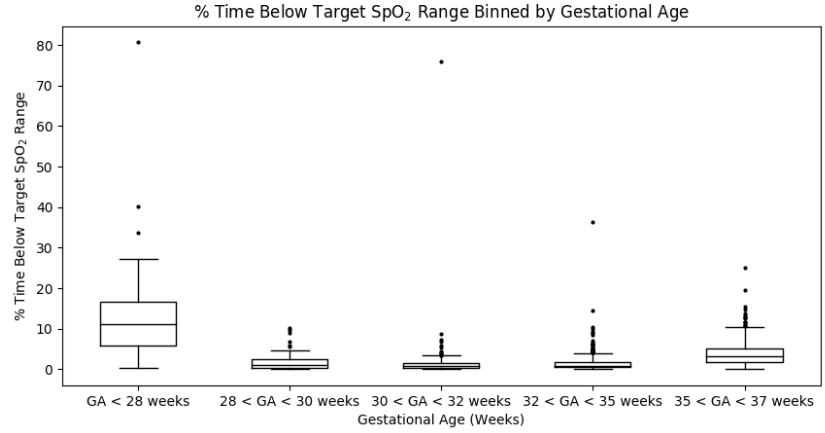

(b)

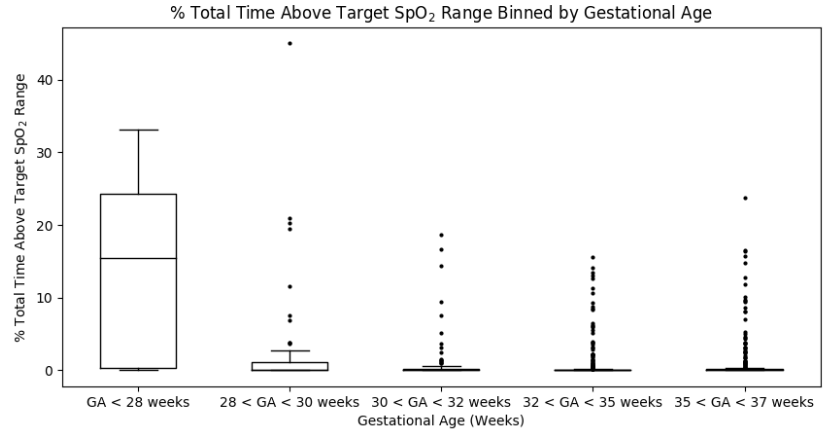

(c)

Figure 5-12: Distribution of % time an individual patient spent (a) within, (b) below, and (c) total above their target SpO_2 range binned by GA.

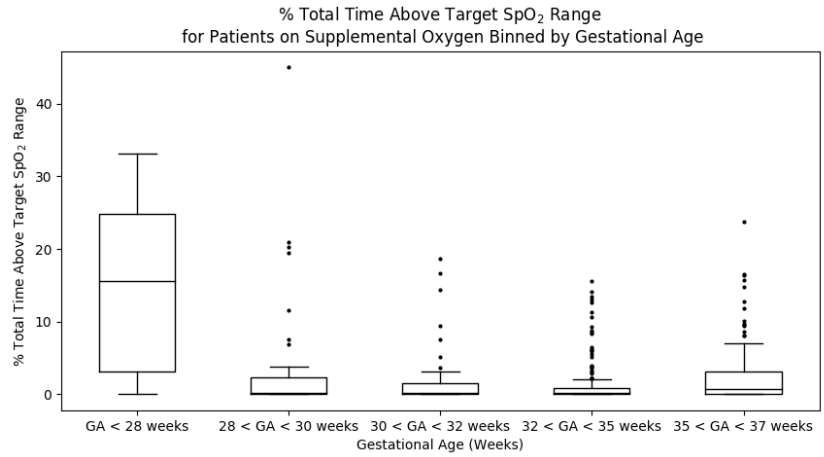

(a)

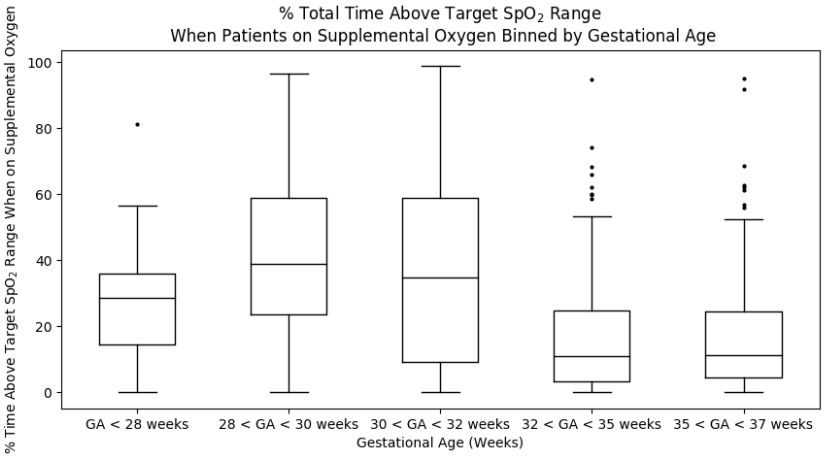

(b)

Figure 5-13: Distribution of % time an individual patient spent above their target SpO_2 range out of (a) their total stay and (b) their total time on supplemental oxygen binned by GA. This distribution represents a subset of patients that are recorded on supplemental oxygen at some point during their NICU stay.

Table 5.2: Breakdown of % time individual patients spent above their target SpO_2 range out of their total stay and total time on supplemental oxygen binned by GA. The patients included in this table were recorded on supplemental oxygen at some point during their NICU stay. Data are medians and interquartile ranges.

GA Category	**Number of Patients**	**% Above out of Total Stay**	**% Above out of Time on Supplemental Oxygen**
GA $<$ 28 weeks	60	15.6 (3.1 - 24.8)	28.5 (14.5 - 35.8)
28 $<$ GA $<$ 30 weeks	47	0.1 (0.0 - 2.3)	38.8 (23.6 - 58.8)
30 $<$ GA $<$ 32 weeks	39	0.2 (0.0 - 1.5)	34.7 (9.3 - 59.0)
32 $<$ GA $<$ 35 weeks	157	0.1 (0.0 - 0.8)	11.0 (3.3 - 24.9)
35 $<$ GA $<$ 37 weeks	99	0.8 (0.1 - 3.1)	11.4 (4.5 - 24.6)

5.4.2 Birth Weight

Regarding birth weight (BW), we categorized the 851 patients in our study cohort with SpO_2 recorded into BW $<$ 500 grams, 500 $<$ BW $<$ 1000 grams, 1000 $<$ BW $<$ 1500 grams, and BW $>$ 1500 grams; the number of patients in each of these categories can be seen in Figure 5-14. Patients less than 500 grams are considered micro-preemies; patients between 500 and 1000 grams are extremely low birth weight; and patients between 1000 and 1500 grams are very low birth weight [7]. We analyzed the % time spent within, below, and above the target SpO_2 range when patients are categorized by BW and these findings can be viewed in Figure 5-15; additionally, a table breaking down this analysis can be seen in Table 5.3.

For our % time above the target SpO_2 range calculation, we also considered the subset of patients on supplemental oxygen; the number of these patients can also be seen in Figure 5-14. For this subset, we calculated the % total time above their target SpO_2 range as well as the % time above their target SpO_2 range when they are recorded on supplemental oxygen. This analysis can be viewed in Figure 5-16; additionally, a table breaking down this analysis can be seen in Table 5.4.

Through this, we observed that comparatively, patients less than 1000 grams spent less time within their target SpO_2 range. The % time within the target SpO_2 range for these patients seems to be higher than studies explored in Section 2.3 focusing on low birth weight patients; however, as noted before, these studies focused on patients on supplemental oxygen for limited amounts of time. When patients are on supplemental oxygen, they spent around 30% of their stay above the target SpO_2 range; this seems to align with the reviewed literature.

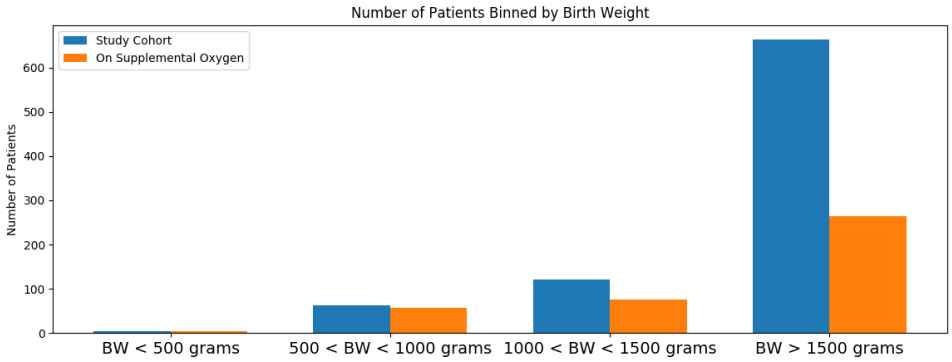

Figure 5-14: Number of patients in each BW category.

Table 5.3: Breakdown of % time individual patients spent within, below, and total above their target SpO_2 range binned by BW. Data are medians and interquartile ranges.

BW Category	**Number of Patients**	**% Within**	**% Below**	**% Total Above**
BW < 500 grams	5	58.8 (50.7 - 74.9)	15.8 (10.7 - 18.2)	25.4 (14.4 - 31.1)
500 < BW < 1000 grams	62	73.7 (63.3 - 96.1)	8.3 (3.3 - 15.6)	12.0 (0.0 - 22.9)
1000 < BW < 1500 grams	121	99.0 (96.1 - 99.6)	0.9 (0.4 - 2.5)	0.0 (0.0 - 0.2)
BW > 1500 grams	663	98.3 (96.4 - 99.3)	1.5 (0.7 - 3.3)	0.0 (0.0 - 0.1)

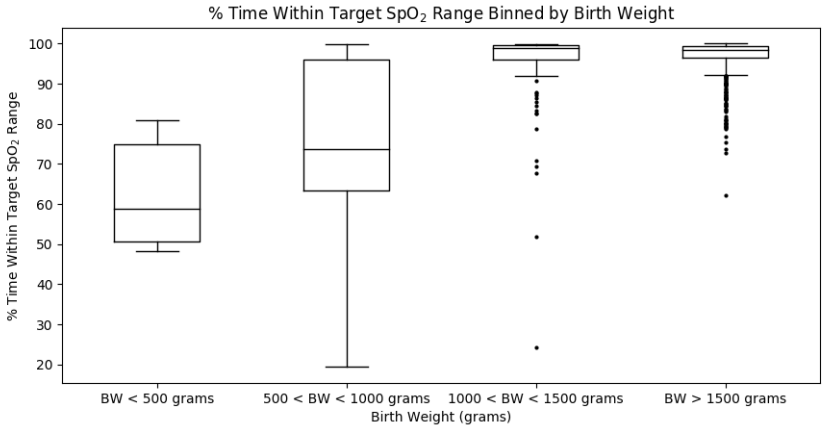

(a)

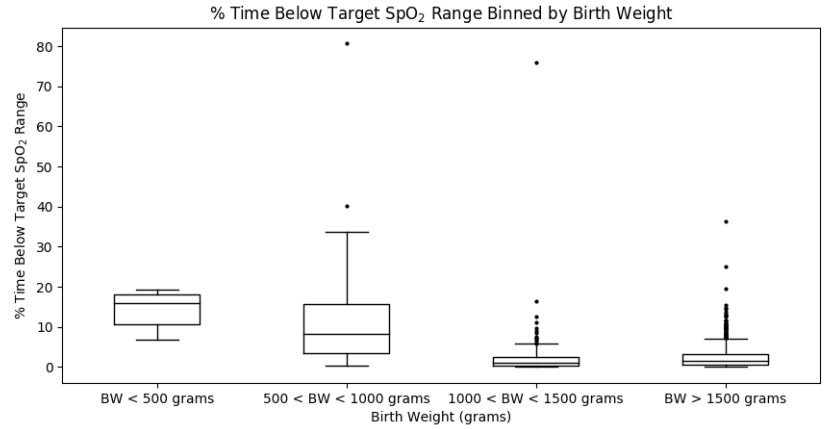

(b)

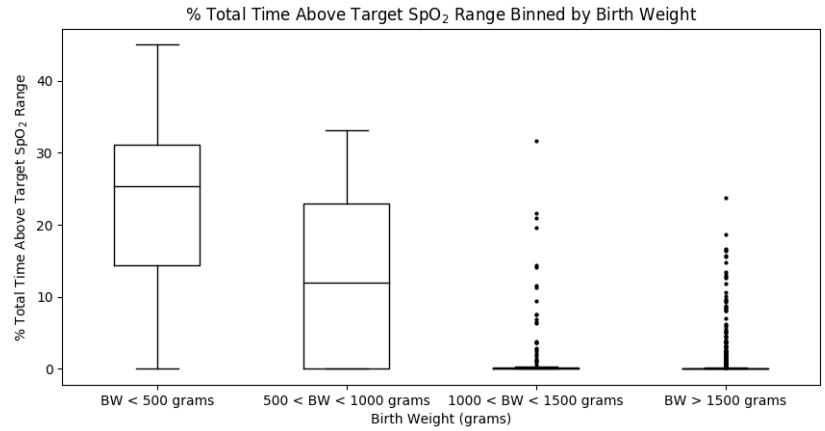

(c)

Figure 5-15: Distribution of % time an individual patient spent (a) within, (b) below, and (c) total above their target SpO_2 range binned by BW.

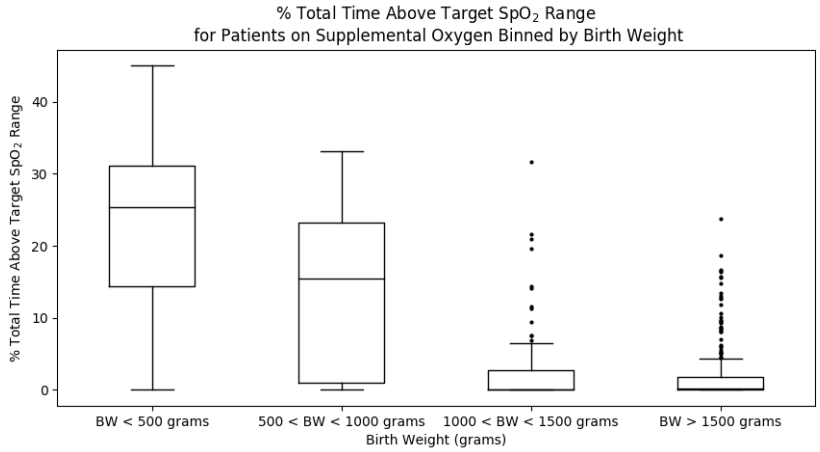

(a)

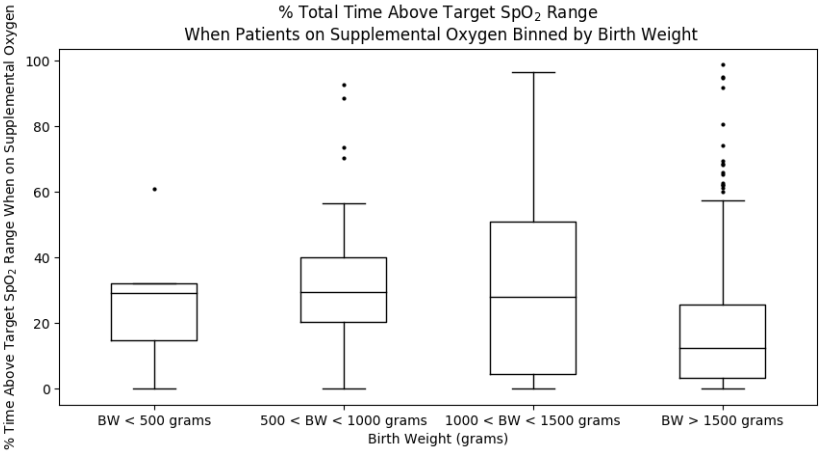

(b)

Figure 5-16: Distribution of % time an individual patient spent above their target SpO_2 range out of (a) their total stay and (b) their total time on supplemental oxygen binned by BW. This distribution represents a subset of patients that are recorded on supplemental oxygen at some point during their NICU stay.

Table 5.4: Breakdown of % time individual patients spent above their target SpO_2 range out of their total stay and total time on supplemental oxygen binned by BW. The patients included in this table were recorded on supplemental oxygen at some point during their NICU stay. Data are medians and interquartile ranges.

BW Category	**Number of Patients**	**% Above out of Total Stay**	**% Above out of Time on Supplemental Oxygen**
BW $<$ 500 grams	5	25.4 (14.4 - 31.1)	29.3 (14.7 - 32.3)
500 $<$ BW $<$ 1000 grams	56	15.5 (0.9 - 23.3)	29.6 (20.5 - 40.0)
1000 $<$ BW $<$ 1500 grams	76	0.1 (0.0 - 2.7)	27.9 (4.6 - 50.9)
BW $>$ 1500 grams	265	0.2 (0.0 - 1.8)	12.4 (3.5 - 25.8)

Table 5.5: Breakdown of % time individual patients spent within, below, and total above their target SpO_2 range binned by gender. Data are medians and interquartile ranges.

Gender	Number of Patients	% Within	% Below	% Total Above
Female	365	98.4 (95.8 - 99.3)	1.5 (0.7 - 3.6)	0.0 (0.0 - 0.1)
Male	457	98.2 (95.6 - 99.3)	1.6 (0.6 - 3.5)	0.0 (0.0 - 0.3)

5.4.3 Gender

Regarding gender, only 822 patients in our study cohort had gender and SpO_2 data recorded; there were 365 females and 457 males recorded. We analyzed the % time spent within, below, and above the target SpO_2 range when patients are categorized by gender and these findings can be viewed in Figure 5-17; additionally, a table breaking down this analysis can be seen in Table 5.5.

For our % time above the target SpO_2 range calculation, we also considered the subset of patients on supplemental oxygen of which there were 162 females and 220 males. For this subset, we calculated the % total time above their target SpO_2 range as well as the % time above their target SpO_2 range when they are recorded on supplemental oxygen. This analysis can be viewed in Figure 5-18; additionally, a table breaking down this analysis can be seen in Table 5.6.

Through this, we observed that females and males spend similar amounts of time within their target SpO_2 range. However, we noted that males were observed to be above their target SpO_2 range slightly more when on supplemental oxygen. One potential reason for this finding is that male infants have been found to face greater morbidity and mortality rates [5].

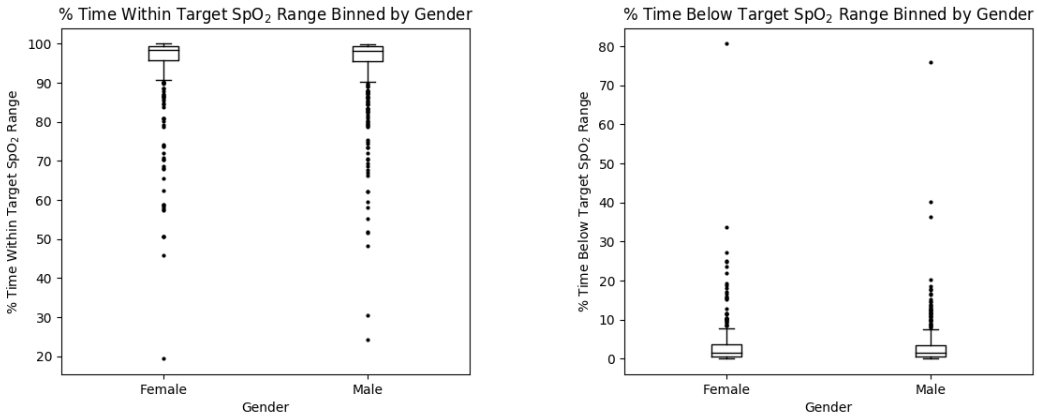

(a)

(b)

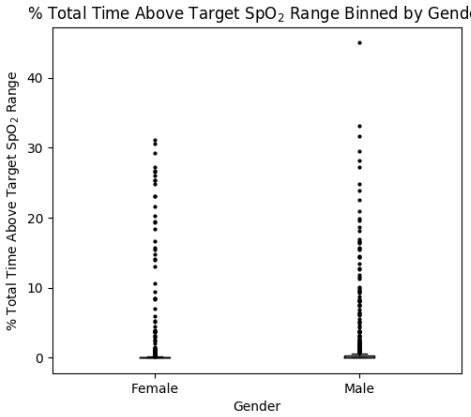

(c)

Figure 5-17: Distribution of % time an individual patient spent (a) within, (b) below, and (c) total above their target SpO_2 range binned by gender.

Table 5.6: Breakdown of % time individual patients spent above their target SpO_2 range out of their total stay and total time on supplemental oxygen binned by gender. The patients included in this table were recorded on supplemental oxygen at some point during their NICU stay. Data are medians and interquartile ranges.

Gender	**Number of Patients**	**% Above out of Total Stay**	**% Above out of Time on Supplemental Oxygen**
Female	162	0.2 (0.0 - 3.1)	14.3 (1.8 - 34.5)
Male	220	0.4 (0.0 - 4.4)	20.2 (6.0 - 35.4)

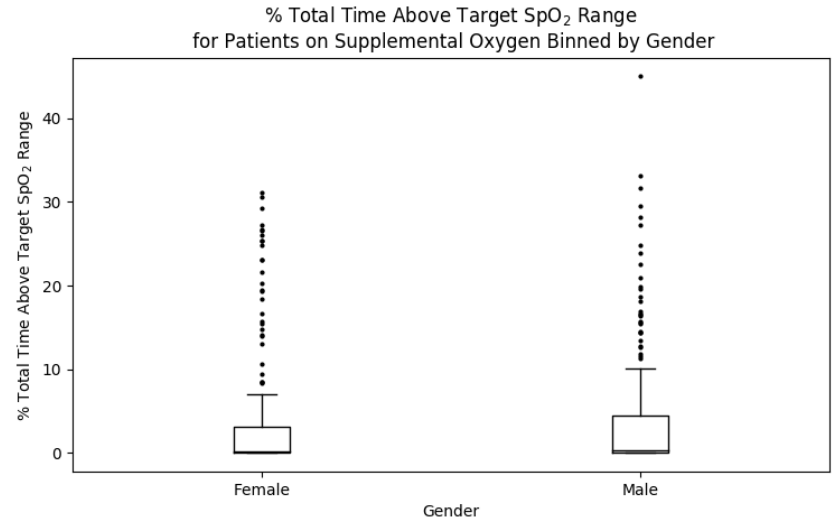

(a)

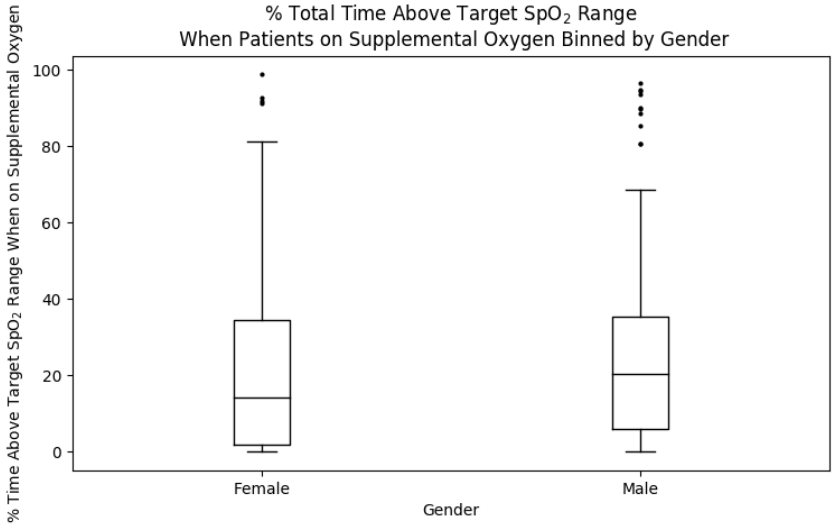

(b)

Figure 5-18: Distribution of % time an individual patient spent above their target SpO_2 range out of (a) their total stay and (b) their total time on supplemental oxygen binned by gender. This distribution represents a subset of patients that are recorded on supplemental oxygen at some point during their NICU stay.

5.4.4 Supplemental Oxygen

We also wanted to analyze the % time spent within, below, and above the target SpO_2 range when patients are on supplemental oxygen versus not on supplemental oxygen. In this case, a patient can be part of the supplemental oxygen category as well as the not on supplemental oxygen category, but portions of their stay will be split between these two categories. To account for this, we updated the formulas for % time spent within, below and above the target SpO_2 range. Here we redefine the equation for % time spent within the target SpO_2 range, and % time below and above have been updated similarly. The equation for % time spent within the target SpO_2 range for the not on supplemental oxygen category is

$$\frac{\text{number of seconds with valid } SpO_2 \text{ recorded in DWC within the target range while the patient is not on supplemental oxygen}}{\text{number of seconds with valid } SpO_2 \text{ recorded in DWC while the patient is not on supplemental oxygen}}$$

Similarly, the equation for % time spent within the target SpO_2 range for the on supplemental oxygen category is

$$\frac{\text{number of seconds with valid } SpO_2 \text{ recorded in DWC within the target range while the patient is on supplemental oxygen}}{\text{number of seconds with valid } SpO_2 \text{ recorded in DWC while the patient is on supplemental oxygen}}$$

These findings can be viewed in Figure 5-19; additionally, a table breaking down this analysis can be seen in Table 5.7. We note that patients not on supplemental oxygen will never be above their target SpO_2 range.

Through this, we observed that patients spend more time below their target SpO_2 range than when they are on supplemental oxygen; this finding along with inevitable increased time above their target SpO_2 lead to less time within the target SpO_2 range when on supplemental oxygen. Additionally, when patients are on supplemental oxygen, they spent more time above their target SpO_2 range rather than below; this aligns with the literature discussed in Section 2.3.

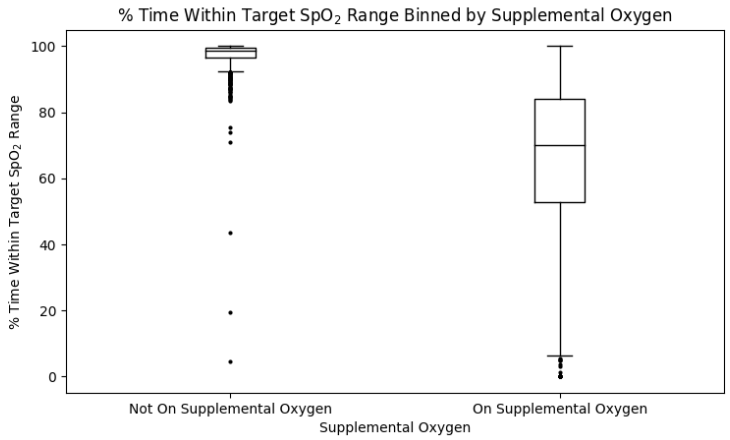

(a)

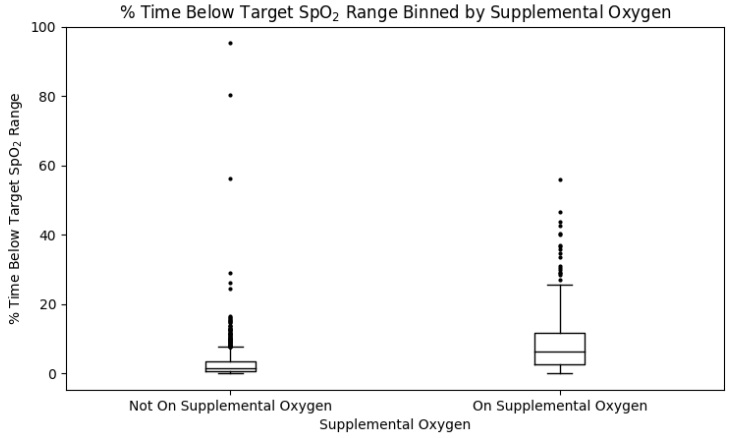

(b)

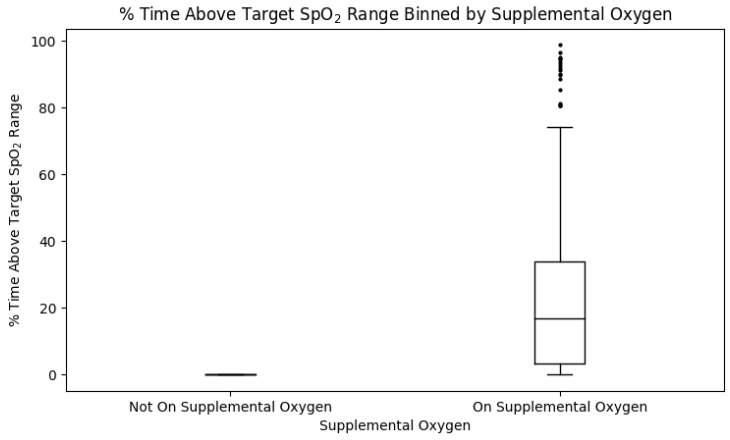

(c)

Figure 5-19: Distribution of % time an individual patient spent (a) within, (b) below, and (c) above their target SpO_2 range binned by supplemental oxygen usage.

Table 5.7: Breakdown of % time individual patients spent within, below, and total above their target SpO_2 range binned by supplemental oxygen usage. Data are medians and interquartile ranges.

Supplemental Oxygen Usage	**Number of Patients**	**% Within**	**% Below**	**% Total Above**
Not on Supplemental Oxygen	851	98.6 (96.5 - 99.4)	1.4 (0.6 - 3.5)	0.0 (0.0 - 0.0)
On Supplemental Oxygen	402	70.0 (52.8 - 84.0)	6.3 (2.5 - 11.8)	16.8 (3.4 - 33.8)

5.4.5 Postmenstrual Age

Regarding postmenstrual age (PMA), we considered a patient's SpO_2 over the course of their NICU stay by stratifying based on PMA (GA + chronological age) while the patient's PMA is $\leq$ 37 weeks. For example, if a patient had a GA of 30 weeks and stayed in the NICU for 3 weeks, then they would contribute to the postmenstrual bins of 30, 31, 32, and 33 weeks based on how much SpO_2 was recorded in DWC for each week. The number of patients contributing to each postmenstrual age binning can be see in Figure 5-20; the maximum number of patients occurs at 35 weeks. We analyzed the % time spent within, below, and above the target SpO_2 range when patients are stratified by PMA and these findings can be viewed in Figure 5-21.

For our % time above the target SpO_2 range calculation, we also considered the subset of patients on supplemental oxygen; the number of these patients can also be seen in Figure 5-20. For this subset, we calculated the % total time above their target SpO_2 range as well as the % time above their target SpO_2 range when they are recorded on supplemental oxygen. This analysis can be viewed in Figure 5-22.

Through this, we can see that it may become more difficult to maintain SpO_2 as infants mature as we observe an increase in % time spent above SpO_2 range as PMA increases. This aligns with a hypothesis in the literature which suggested that it may be difficult to maintain SpO_2 as infants mature due to patients requiring less invasive respiratory support and development of other complications [11].

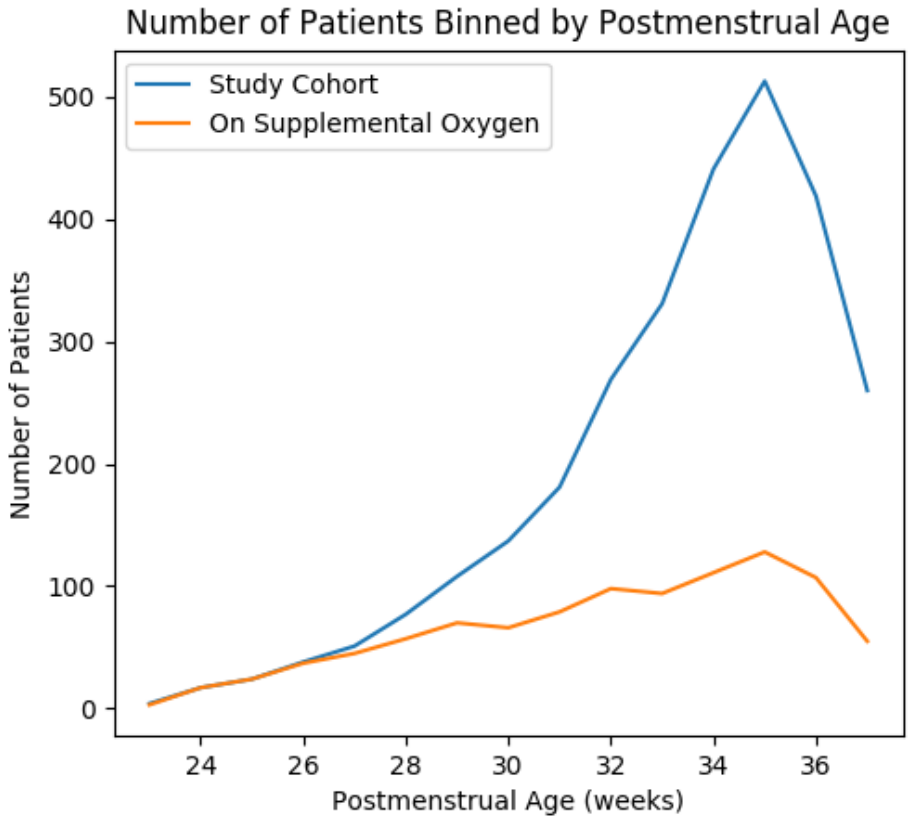

Figure 5-20: Illustrates the number of patients in each PMA category.

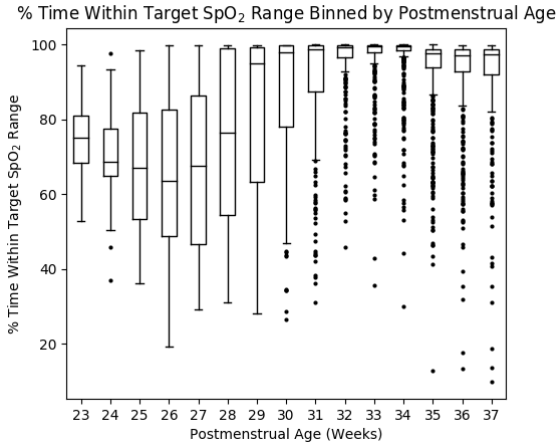

(a)

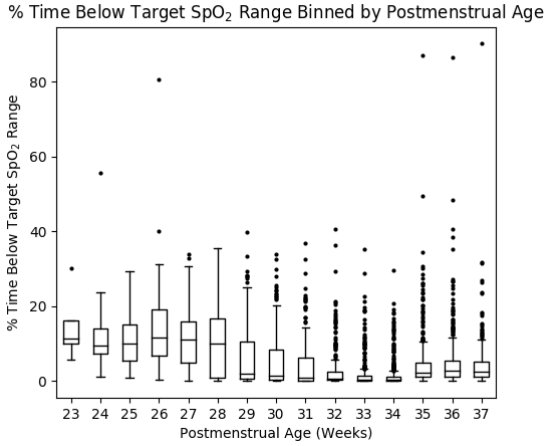

(b)

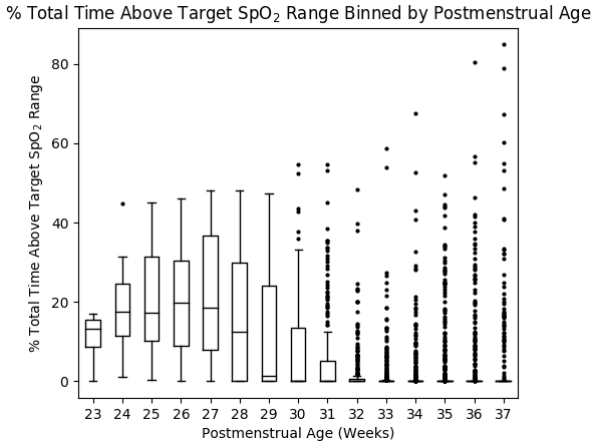

(c)

Figure 5-21: Distribution of % time an individual patient spent (a) within, (b) below, and (c) total above their target SpO_2 range binned by PMA.

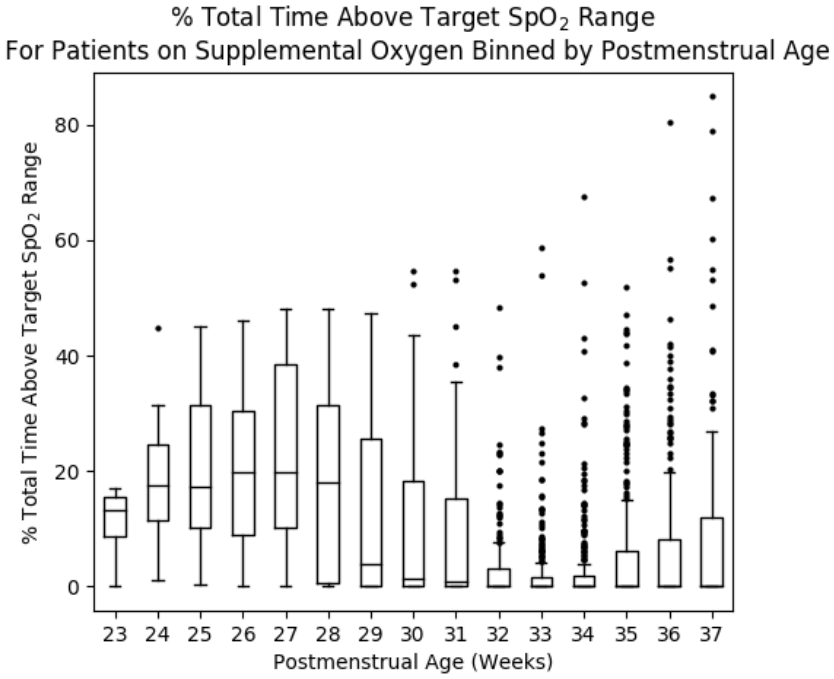

(a)

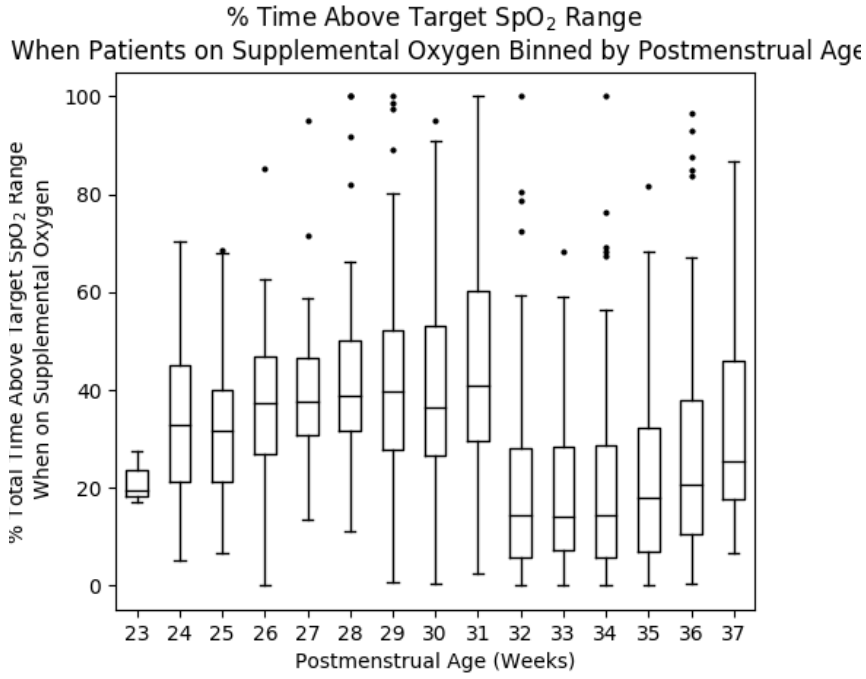

(b)

Figure 5-22: Distribution of % time an individual patient spent above their target SpO_2 range out of (a) their total stay and (b) their total time on supplemental oxygen binned by PMA. This distribution represents a subset of patients that are recorded on supplemental oxygen at some point during their NICU stay.

5.4.6 Time of Day

Regarding time of day, we considered a patient's SpO_2 over the course of the day by stratifying each day by hour and aggregating all days together. We analyzed the % time spent within, below, and above the target SpO_2 range when patients are stratified by hour and these findings can be viewed in Figure 5-23.

For our % time above the SpO_2 range calculation, we also considered the subset of patients on supplemental oxygen. For this subset, we calculated the % total time above their SpO_2 range as well as the % time above their target SpO_2 range when they are recorded on supplemental oxygen. This analysis can be viewed in Figure 5-24.

Through this analysis, we found that all hours had approximately the same % time within, below, and above the target SpO_2 range indicating consistent care over time. We noted that the nurses' shift change at 7AM and 7PM; these graphs indicate that this shift change may not be causing a disruption in the NICU when considering patient oxygenation.

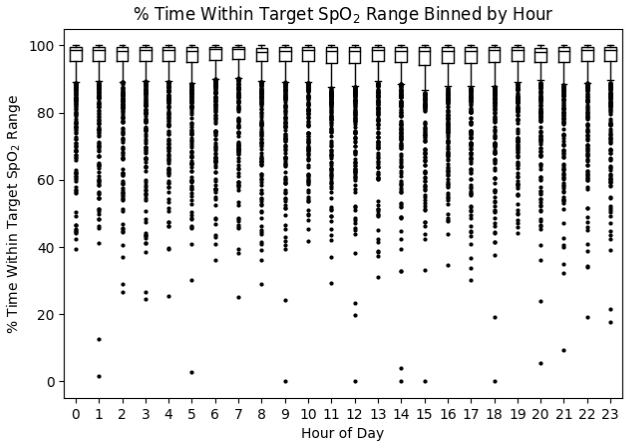

(a)

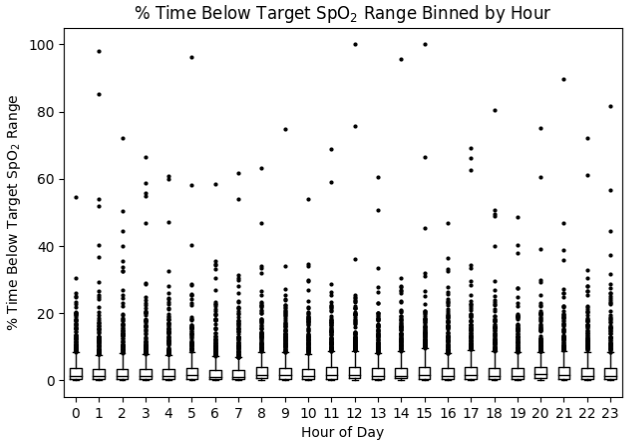

(b)

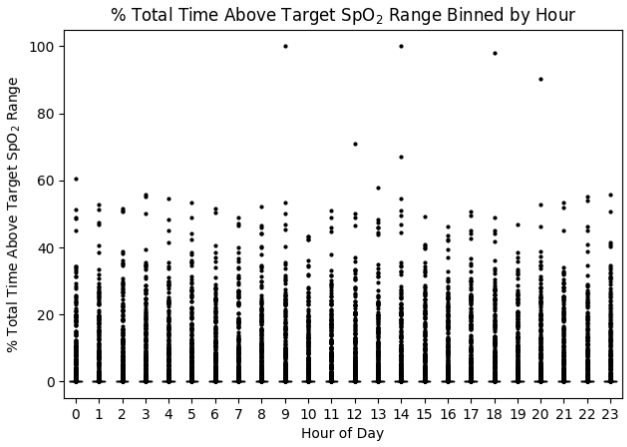

(c)

Figure 5-23: Distribution of % time an individual patient spent (a) within, (b) below, and (c) above their target SpO_2 range binned by hour.

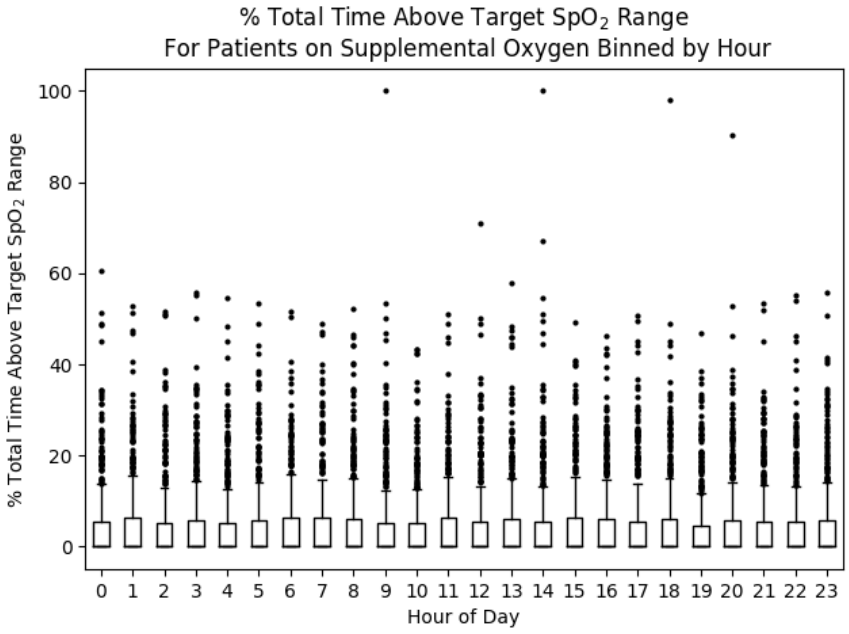

(a)

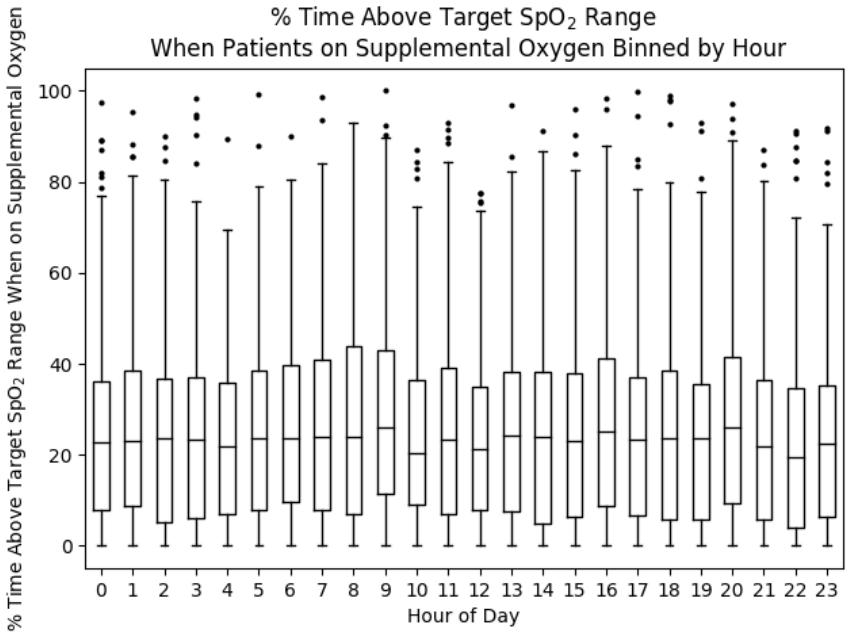

(b)

Figure 5-24: Distribution of % time an individual patient spent above their target SpO_2 range out of (a) their total stay and (b) their total time on supplemental oxygen binned by hour. This distribution represents a subset of patients that are recorded on supplemental oxygen at some point during their NICU stay.

5.4.7 Bed Space Assignment

Regarding bed space assignment, we considered a patient's SpO_2 stratified by specific bed spaces; the total number of patients recorded in each bed space can be seen in Figure 5-25. We analyzed the % time spent within, below, and above the target SpO_2 range when patients are stratified by bed space and these findings can be viewed in Figure 5-26. Since bed space 3 is an extension bed space and is often not occupied, we removed this bed space from our analysis.

For our % time above the SpO_2 range calculation, we also considered the subset of patients on supplemental oxygen. For this subset, we calculated the % total time above their SpO_2 range as well as the % time above their target SpO_2 range when they are recorded on supplemental oxygen. This analysis can be viewed in Figure 5-27.

It is hard to uncover any trends from this analysis due to the large number of bed spaces. Since we were aware that certain bed spaces are used for sicker patients due to their location being closer to the main central monitoring station, we worked with clinicians to group the bed spaces into high and low acuity. Bed spaces in rooms 952-957 and 977-981 were deemed low acuity; in Figure 5-26 and Figure 5-27, this represents the first 12 and last 10 bed spaces.

We also note that bed space 1 is closer to the door while bed space 2 is closer to the window in each NICU room; we considered this bed position when we grouped these bed spaces. The number of patients found in these groups can also be seen in Figure 5-25.

With this new grouping, we reanalyzed the % time spent within, below, and above the target SpO_2 range; this can be viewed in Figure 5-28 and Table 5.8. Additional analysis regarding the % time spent above the target SpO_2 range can be found in Figure 5-29 and Table 5.9.

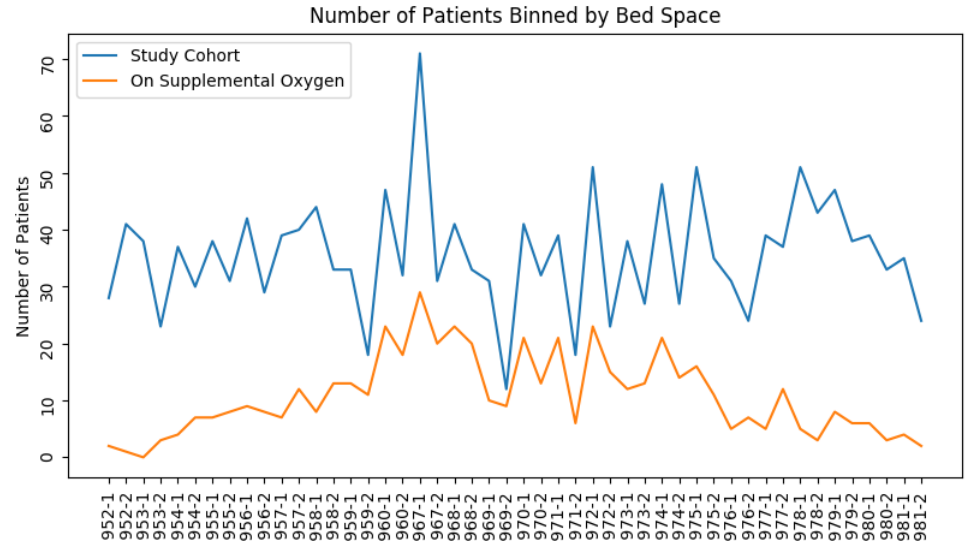

(a)

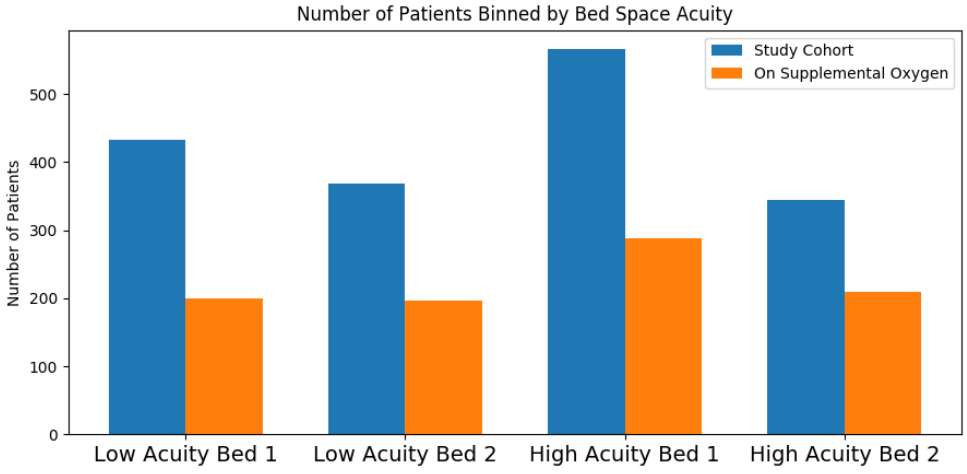

(b)

Figure 5-25: Illustrates the number of patients in each (a) bed space and (b) bed space acuity grouping.

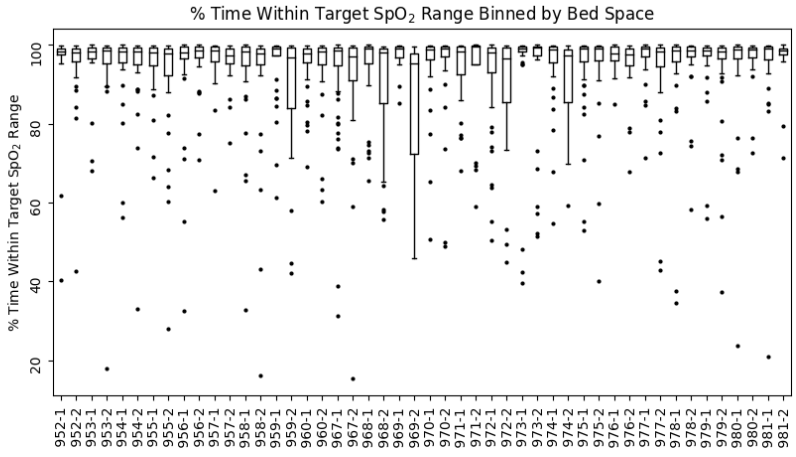

(a)

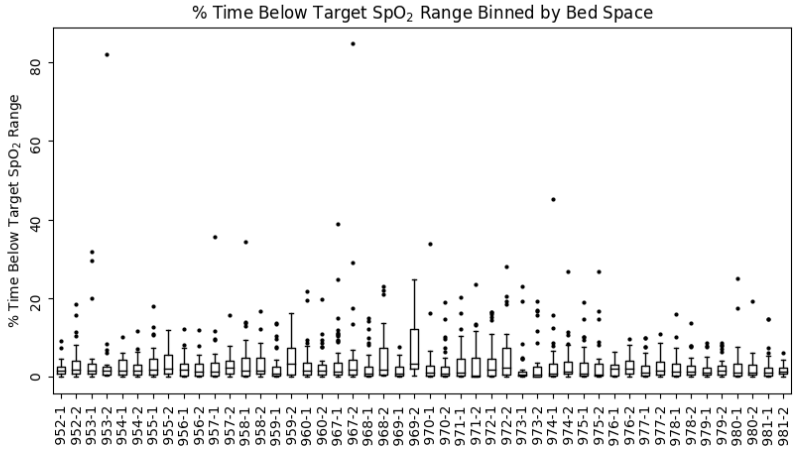

(b)

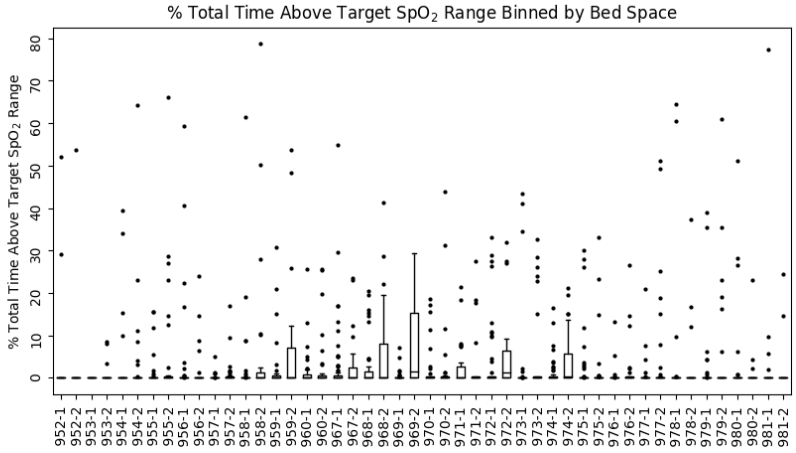

(c)

Figure 5-26: Distribution of % time an individual patient spent (a) within, (b) below, and (c) above their target SpO_2 range binned by bed space.

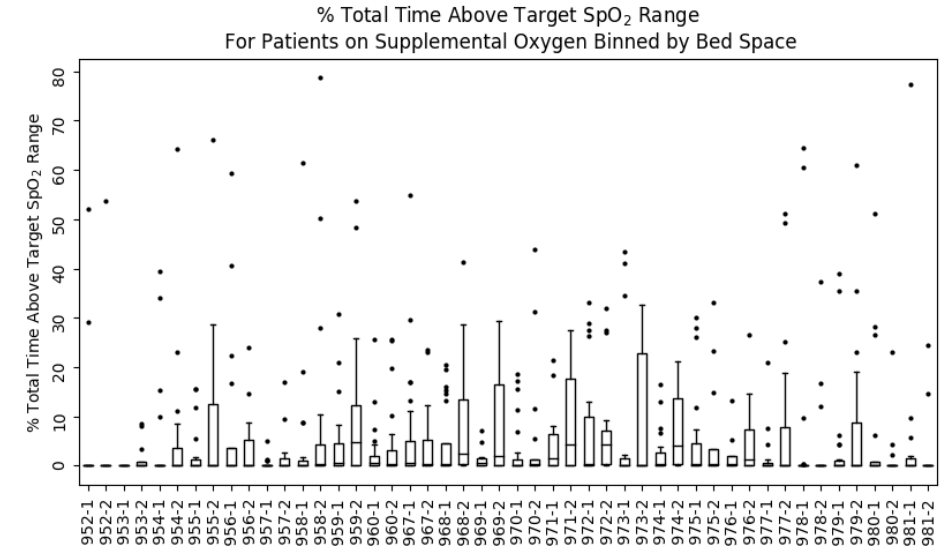

(a)

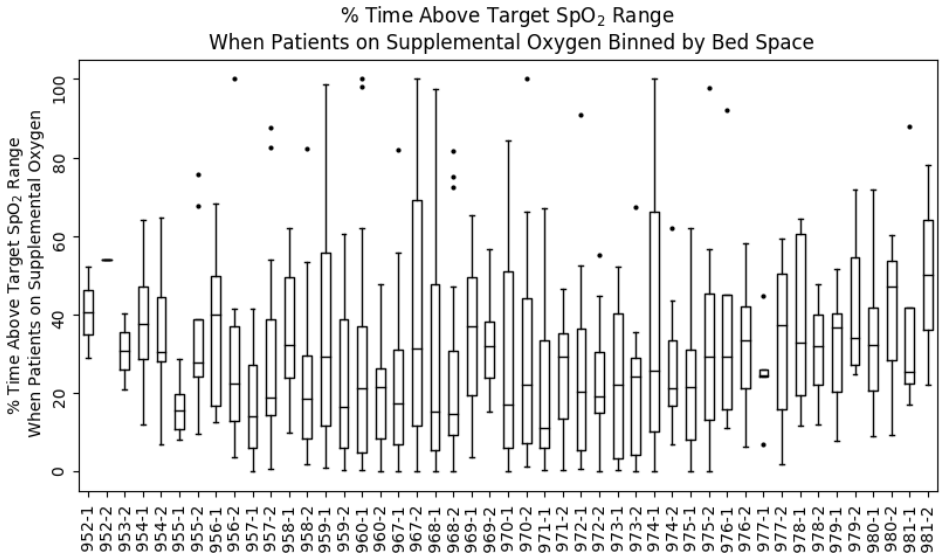

(b)

Figure 5-27: Distribution of % time an individual patient spent above their target SpO_2 range out of (a) their total stay and (b) their total time on supplemental oxygen binned by bed space. This distribution represents a subset of patients that are recorded on supplemental oxygen at some point during their NICU stay.

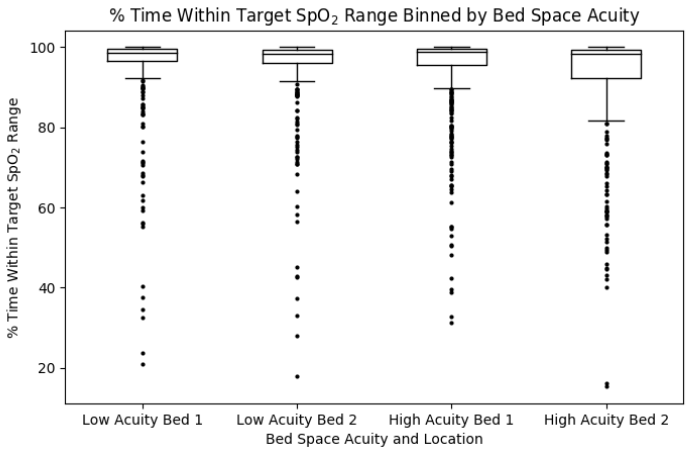

(a)

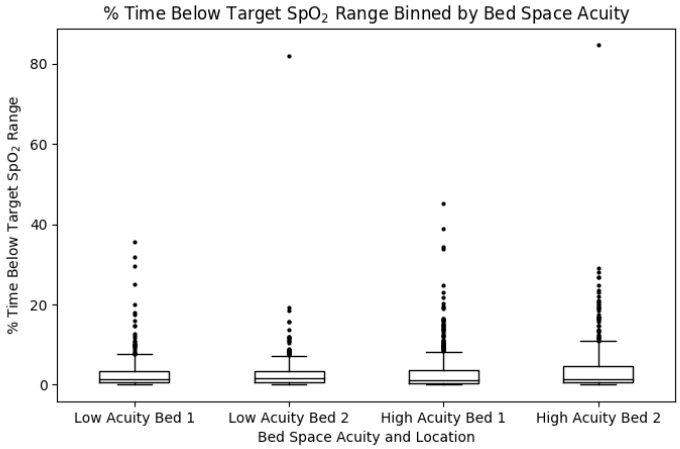

(b)

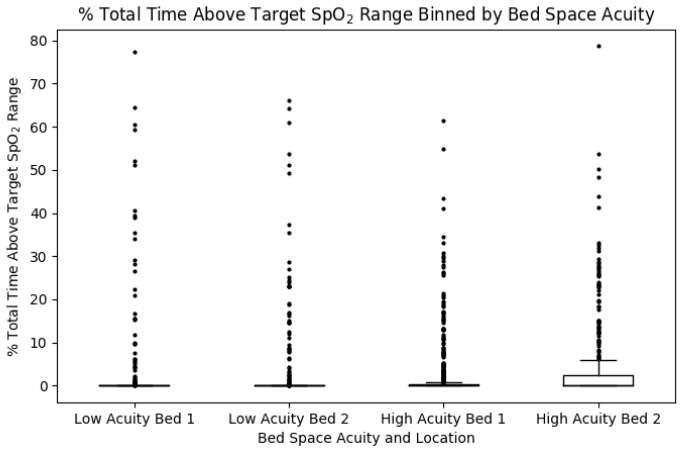

(c)

Figure 5-28: Distribution of % time an individual patient spent (a) within, (b) below, and (c) above their target SpO_2 range binned by bed space acuity.

Table 5.8: Breakdown of % time individual patients spent within, below, and total above their target SpO_2 range binned by bed space acuity. Data are medians and interquartile ranges.

Bed Space Acuity	**Number of Patients**	**% Within**	**% Below**	**% Total Above**
Low Acuity: Bed 1	433	98.6 (96.5 - 99.5)	1.4 (0.5 - 3.4)	0.0 (0.0 - 0.0)
Low Acuity: Bed 2	369	98.3 (96.0 - 99.3)	1.6 (0.7 - 3.3)	0.0 (0.0 - 0.0)
High Acuity: Bed 1	566	98.7 (95.6 - 99.5)	1.1 (0.4 - 3.6)	0.0 (0.0 - 0.3)
High Acuity: Bed 2	345	98.4 (92.3 - 99.5)	1.4 (0.5 - 4.7)	0.0 (0.0 - 2.4)

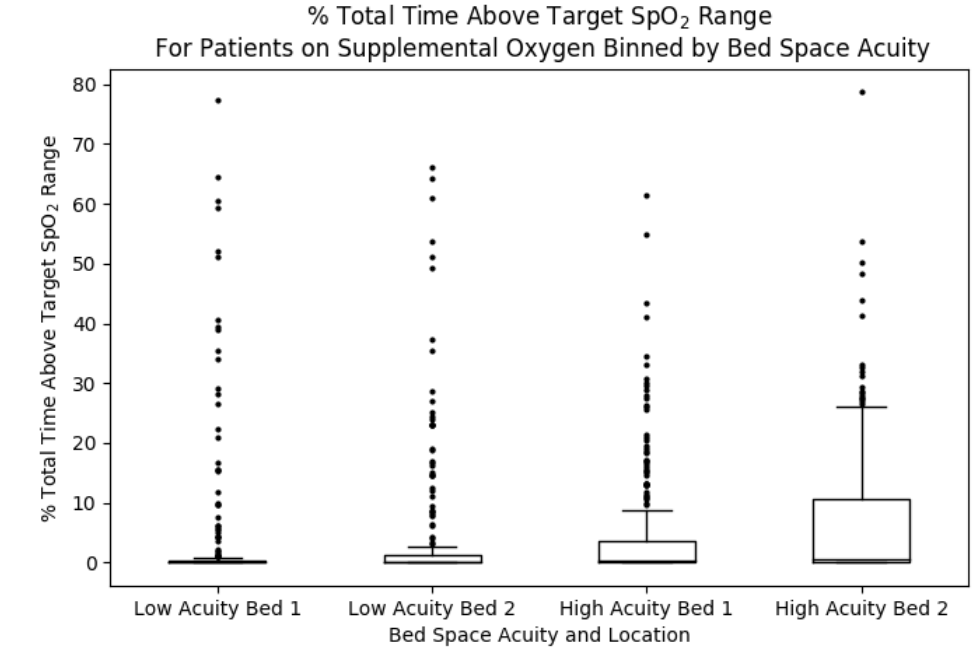

(a)

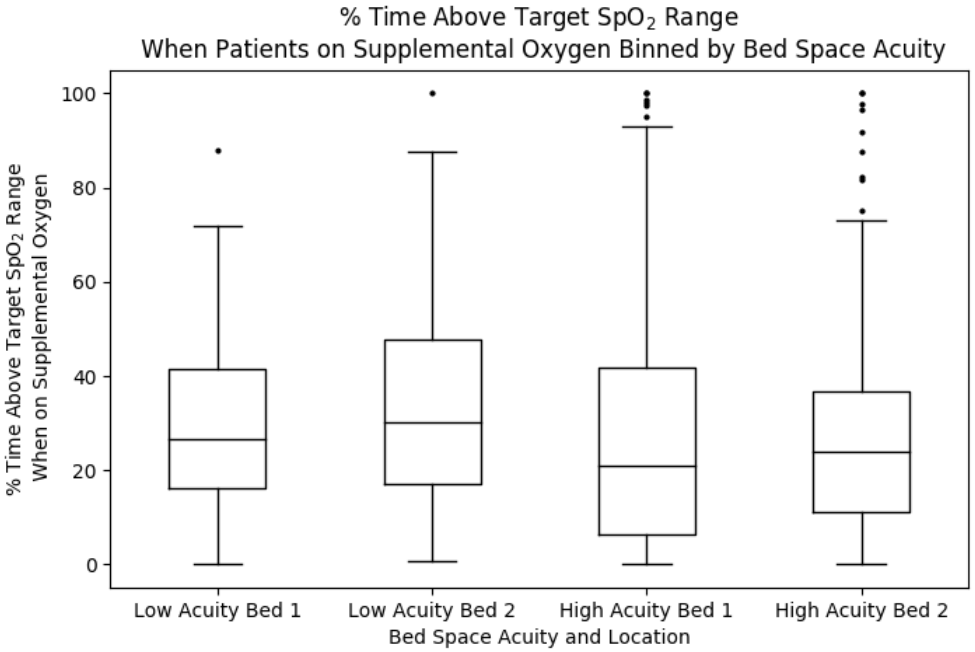

(b)

Figure 5-29: Distribution of % time an individual patient spent above their target SpO_2 range out of (a) their total stay and (b) their total time on supplemental oxygen binned by bed space acuity. This distribution represents a subset of patients that are recorded on supplemental oxygen at some point during their NICU stay.

Table 5.9: Breakdown of % time individual patients spent above their target SpO_2 range out of their total stay and total time on supplemental oxygen binned by bed space acuity. The patients included in this table were recorded on supplemental oxygen at some point during their NICU stay. Data are medians and interquartile ranges.

Bed Space Acuity	**Number of Patients**	**% Above out of Total Stay**	**% Above out of Time on Supplemental Oxygen**
Low Acuity: Bed 1	199	0.0 (0.0 - 0.0)	26.5 (16.3 - 41.6)
Low Acuity: Bed 2	196	0.0 (0.0 - 1.2)	30.3 (17.0 - 47.8)
High Acuity: Bed 1	288	0.3 (0.0 - 3.7)	20.9 (6.4 - 41.8)
High Acuity: Bed 2	210	0.6 (0.0 - 10.5)	23.9 (11.1 - 36.6)

5.5 Summary

We were able to take advantage of the massive amounts of SpO_2 data provided by DWC to determine the % time patients are within, below, and above their target SpO_2 range; we also began exploring potential factors which may be affecting this compliance.

To ensure the usability of the SpO_2 data, we performed multiple data integrity checks which led us to the conclusion that we had 90.3% data coverage; this accounted for throwing out SpO_2 data from hypoxemic episodes that had missing data from the start of the episode.

We determined that our study cohort spent a total of 90.0% of their collective stay within their target SpO_2 range; 4.9% of their stay below their target SpO_2 range; and 5.1% of their stay above their target SpO_2 range. We took a subset of 406 patients on supplemental oxygen and found that 7.5% of their total stay was spent above their target SpO_2 range while 24.0% of their time on supplemental oxygen was spent above their target SpO_2 range.

Finally, we investigated various factors and their relationship with SpO_2; factors included gestational age, birth weight, gender, supplemental oxygen, postmenstrual age, time of day, and bed space assignment. Overall, we determined that younger patients with lower birth weights spend the least time within their target SpO_2 range. Additionally, male patients spent a greater % time above their target SpO_2 range while on supplemental oxygen than female patients.

Chapter 6

Conclusion

6.1 Summary of Findings

Studies have made it clear that keeping preterm infants' oxygen saturations (SpO_2) in a pre-defined target range for a high fraction of time is important to prevent life threatening complications for these tiniest of patients; however, it can be difficult to maintain infants' SpO_2 within the set target range due to continuously evolving pathophysiology, the intricacies of various respiratory support modalities and possibly complex workflow factors. We worked to leverage clinical, demographic, and workflow information as well as physiological monitoring data streams to identify factors that may place preterm infants at risk for hypoxemia (low SpO_2) and hyperoxemia (high SpO_2).

We collected data from a study cohort of 865 preterm (gestational age $<$ 37 weeks) neonates that were in the NICU for more than 24 hours from January 1, 2018 to September 12, 2019 which granted us access to a larger study cohort with more recorded data than previous studies. Through Data Warehouse Connect (DWC), we had access to patient information including bed space assignment, alarms, and SpO_2 data and we analyzed the integrity of this data to ensure that patients could be correctly be mapped to patients recorded in the NICU's electronic medical record.

After confirming the integrity of the data, we determined we had 90.3% SpO_2 data coverage and began analyzing the time spent within, above, and below a patient's

target SpO_2 range and factors that may be affecting this. We determined that overall patients spent 90.0% of their collective stay within their target SpO_2 range; 4.9% of their stay below their target SpO_2 range; and 5.1% of their stay above their target SpO_2 range. Lastly, we determined that younger patients with lower birth weights spent the least time within their target SpO_2 range; additionally, male patients spent a greater % time above their target SpO_2 range while on supplemental oxygen.

Overall, we were able to explore a massive data source that provided continuous SpO_2 measurements for a large study cohort and explore factors affecting patients' in-target oxygenation.

6.2 Future Work

Through this work, we were able to familiarize ourselves with the DWC infrastructure and validate the massive amounts of data it provides. We leveraged this data to begin preliminary exploration of factors that may be affecting the % time spent within, below, and above the target SpO_2 range. From here, there are many directions that future work could explore. There are several additional factors that could be affecting % time spent within the target SpO_2 range that would be interesting to explore including a patient's acuity score (i.e. SNAP score), the NICU census, respiratory support (i.e. mechanical ventilation, high flow nasal cannula) and nurse to patient ratio. We could also analyze effects of out-of-target oxygenation by collecting records of patient outcomes and comparing the frequency of these outcomes with out-of-target oxygenation. Being able to explore these factors and outcomes with a large study cohort and continuous SpO_2 monitoring could provide great insight.

Additionally, as the developed code is generalizable and scalable, researchers could create a dashboard that is updated weekly with various statistics including analysis of the % time spent within, below, and above the target SpO_2 range in addition to the relationship with all the factors we explored in Chapter 5. This would help close the loop between researchers analyzing the data streams and the BIDMC clinical staff monitoring the NICU and allow clinicans to have an interactive tool to explore the

SpO_2 data and become more aware of the oxygenation of patients and what factors may be driving it.

While our current study focused on patients that spend time below their target SpO_2 range, it would be interesting to explore intermittent hypoxemia specifically. This could involve analysis of the amount the patient is deviating from their lower SpO_2 limit and determine how many of these deviations cause hypoxemic episodes ($Spo_2 < 80\%$). Through this, we may be able to determine an ideal lower SpO_2 limit that minimizes both hypoxemic episodes and alarms (and associated alarm fatigue).

Related to alarms, DWC also provides a rich dataset regarding NICU alarms which relate to SpO_2 as well as many other measurements such as HR and respiratory rate. These alarms can be analyzed in parallel with the SpO_2 data to also attempt to reduce alarm fatigue in the NICU.

DWC provides a great opportunity to explore many various aspects of a patient's stay in the NICU and the validation of the data accomplished through this study confirmed the untapped potential of this rich data source.

Appendix A

Tables

Table A.1: Columns in DWC tables used in our study.

Table Name	Column	Description
Patient	id	Unique ID that represents a patient.
	time_stamp	The time the row is initially stored in the database.
	bed_label	Bed space label.
	alias	Alias.
	category	Indicates the category of the patient (adult, pediatric or neonatal).
	height	Patient height.
	height_unit	Unit of measure for height.
	weight	Patient weight.
	weight_unit	Unit of measure for weight.
	pressure_unit	Unit of measure for pressure.
	paced_mode	Indicates whether the patient has a pacemaker.
	resuscitation_state	Patient's resuscitation state.
	admit_state	Patient's admission state (0 = Not admitted, 1 = Admitted, 2 = Discharged)
	clinical_unit	Patient's clinical unit.
	gender	Patient's gender

Table Name	Column	Description
Patient String Attribute	patient_id	Unique ID that represents a patient.
	time_stamp	The time the row is initially stored in the database.
	name	Name of the string attribute (i.e. FirstName, LastName, LifetimeId).
	value	Value of the string attribute.

Table Name	Column	Description
Alert	time_stamp	The time the row is initially stored in the database.
	sequence_number	A number that correlates to the time the alarm is announced at the monitoring central station (substitute for time_stamp).
	alert_id	Unique ID that represents a specific alarm.
	source	Source of the alarm.
	code	Generic type of the alarm.
	label	Alarm text.
	severity	Alarm severity (Red, Yellow, Cyan etc.).
	kind	Alarm category (Patient, Technical, Notification or Status).
	is_silenced	Indicates if the alarm is silenced.
	subtype_id	ID for the alarm to uniquely identify its type.
	onset_time	Time the alarm condition begins.
	announce_time	Time the alarm is announced.
	end_time	Time the alarm ends.
	mapping_id	ID of the row that corresponds to the patient ID in the patient_mapping table.

Table Name	Column	Description
Numeric Value	numeric_id	ID of the row that corresponds to this value in the numeric table.
	time_stamp	The time the row is initially stored in the database.
	sequence_number	A number that correlates to the time the numeric occurs at the monitoring central station.
	is_trend_uploaded	Indicates if the numeric came from a trend upload.
	compound_value_id	Unique ID to group compound values.
	value	Raw value of the numeric element.
	mapping_id	ID of the row that corresponds to the patient ID in the patient_mapping table.
Patient Mapping	id	Alternative unique mapping ID that represents a patient.
	patient_id	Unique ID that represents a patient.
	time_stamp	The time the row is initially stored in the database.
	is_mapped	Indicates if there is a patient ID mapped to this ID.

Table Name	Column	Description
Numeric	id	ID that represents the set of numeric attributes.
	base_physio_id	Generic physiological numeric type.
	physio_id	PhysioID of the corresponding numeric value.
	label	Numeric label.
	is_periodic	Indicates whether the numeric is aperiodic or periodic.
	unit_label	Unit of measure label.
	validity	Indicates the validity of the numeric value.
	lower_limit	Lower limit of the alarm set for the numeric value.
	upper_limit	Upper limit of the alarm set for the numeric value.
	is_alarming_off	Indicates whether alarms for this numeric are turned off.
	sub_physio_id	Generic physiological type.
	sub_label	Numeric element label.
	color	Hex value of the color.
	is_manual	Indicates whether the numeric is manual or automatic.
	max_values	The compound of the numeric.
	scale	The number of decimal places to be displayed for values and limits.

Appendix B

Figures

Post Menstrual Age (PMA)	Oxygen	Low Saturation Limit	High Saturation Limit	Target Saturation
$\leq 31^{6/7}$	Yes	87%	94%	90%
	No	87%	100%	90%
$32^{0/7} - 34^{6/7}$	Yes	87%	97%	92%
	No	87%	100%	92%
$\geq 35^{0/7}$	Yes	92%	98%	96%
	No	92%	100%	96%

Figure B-1: BIDMC guidelines illustrating target oxygen saturation ranges for patients of different PMA and supplemental oxygen.

Bibliography

[1] Hallenberger A, Poets CF, et al. Closed-loop automatic oxygen control (clac) in preterm infants: a randomized controlled trial. *Pediatrics*, 133(2):e379–385, February 2014.

[2] Sola A, Golombek SG, Montes Bueno MT, et al. Safe oxygen saturation targeting and monitoring in preterm infants: can we avoid hypoxia and hyperoxia? *Acta Paediatrica*, 103(10):1009–1018, October 2014.

[3] Laptook AR, Salhab W, et al. Pulse oximetry in very low birth weight infants: can oxygen saturation be maintained in the desired range? *Journal of Perinatology*, 26(6):337–341, June 2006.

[4] Mills BA, Davis PG, et al. Improving compliance with pulse oximetry alarm limits for very preterm infants. *Journal of Paediatrics and Child Health*, 46(5):255–258, May 2010.

[5] Townsel CD, Emmer SF, et al. Gender differences in respiratory morbidity and mortality of preterm neonates. *Frontiers in Pediatrics*, 5:6, January 2017.

[6] Poets CF, Roberts RS, et al. Association between intermittent hypoxemia or bradycardia and late death or disability in extremely preterm infants. *JAMA*, 314(6):595–603, August 2015.

[7] Cutland CL, Lackritz EM, et al. Low birth weight: Case definition guidelines for data collection, analysis, and presentation of maternal immunization safety data. *Vaccine*, 35(48):6492–6500, December 2017.

[8] Sink DW, Hope SAE, et al. Nurse:patient ratio and achievement of oxygen saturation goals in premature infants. *Archives of Disease in Childhood - Fetal and Neonatal Edition*, 96(2):93–98, October 2010.

[9] Wellington G, Campbell A, and Elder D. Intermittent hypoxia in preterm infants: Measurement using the desaturation index. *Pediatric Pulmonology*, 54(6):865–872, June 2019.

[10] Arawiran J, Curry J, et al. Sojourn in excessively high oxygen saturation ranges in individual, very low-birthweight neonates. *Acta Paediatrica*, 104(2):e51–56, February 2015.

[11] Armbruster J, Schmidt B, et al. Nurses' compliance with alarm limmits for pulse oximetry: qualitative study. *Journal of Perinatology*, 30(8):531–534, August 2010.

[12] Zapata J, Gomez JJ, et al. A randomised controlled trial of an automated oxygen delivery algorithm for preterm neonates receiving supplemental oxygen without mechanical ventilation. *Acta Paediatrica*, 103(19):928–933, September 2014.

[13] Quinn JA, Munoz FM, et al. Preterm birth: Case definition guidelines for data collection, analysis, and presentation of immunisation safety data. *Vaccine*, 34(49):6047–6056, December 2016.

[14] Hagadorn JI, Furey AM, et al. Achieved versus intended pulse oximeter saturation in infants born less than 28 weeks' gestation: The aviox study. *Pediatrics*, 118(4):1574–1582, October 2006.

[15] Cummings JJ, Polin RA, et al. Oxygen targeting in extremely low birth weight infants. *Pediatrics*, 138(2), August 2016.

[16] Di Fiore JM, Kaffashi F, et al. The relationship between patterns of intermittent hypoxia and retinopathy of prematurity in preterm infants. *Pediatric Research*, 72(6):606–612, December 2012.

[17] Di Fiore JM and Vento M. Intermittent hypoxemia and oxidative stress in preterm infants. *Respiratory Physiology and Neurobiology*, 266:121–129, August 2019.

[18] Lim K, Wheeler KI, et al. Oxygen saturation targeting in preterm infants receiving continuous positive airway pressure. *Journal of Pediatrics*, 164(4):730–736, April 2014.

[19] Askie L, Darlow B, et al. Association between oxygen saturation targeting and death or disability in extremely preterm infants in the neonatal oxygenation prospective metaanalysis collaboration. *JAMA*, 319(21):2190–2201, June 2018.

[20] Clucas L, Doyle LW, et al. Compliance with alarm limits for pulse oximetry in very preterm infants. *Pediatrics*, 119(6):1056–1060, June 2007.

[21] Williams LZJ, McNamara D, and Alsweiler JM. Intermittent hypoxemia in infants born late preterm: A prospective cohort observational study. *The Journal of Pediatrics*, 204:89–95, January 2019.

[22] Bizzarro MJ. Optimizing oxygen saturation targets in extremely preterm infants. *JAMA*, 319(21):2173–2174, June 2018.

[23] Urschitz MS, Horn W, et al. Automatic control of the inspired oxygen fraction in preterm infants: a randomized crossover trial. *American Journal of Respiratory and Critical Care Medicine*, 170(10):1095–1100, November 2004.

[24] Claure N, D'Ugard C, and Bancalari E. Automated adjustment of inspired oxygen in preterm infants with frequent fluctuationos in oxygenation: a pilot clinical trial. *Journal of Pediatrics*, 155(5):640–645, November 2009.

[25] Claure N, Bancalari E, et al. Multicenter crossover study of automated control of inspired oxygen in ventilated preterm infants. *Pediatrics*, 127(1):e76–83, January 2011.

[26] Claure N, Gerhardt T, et al. Closed-loop controlled inspired oxygen concentration for mechanically ventilated very low birth weight infants with frequent episodes of hypoxemia. *Pediatrics*, 107(5):1120–1124, May 2001.

[27] Committee on Injury, Poison Prevention, Committee on Fetus, and Newborn. Safe transportation of premature and low birth weight infants. *Pediatrics*, 97(5):758–760, May 1996.

[28] Martin RJ, Wang K, et al. Intermittent hypoxic episodes in preterm infants: Do they matter? *Neonatology*, 100(3):303–310, 2011.

[29] Li T, Matsushima M, et al. Epidemiology of patient monitoring alarms in the neonatal intensive care unit. *Journal of Perinatology*, 38(8):1030–1038, August 2018.

[30] van der Eijk AC, Dankelman J, et al. An observational study to quantify manual adjustments of the inspired oxygen fraction in extremely low birth weight infants. *Acta Paediatrica*, 101(3):e97–104, March 2012.

[31] van Zanten HA, Tan RN, et al. Compliance in oxygen saturation targeting in preterm infants: a systematic review. *European Journal of Pediatrics*, 174(12):1561–1572, December 2015.

[32] Tarnow-Mordi W and Kirby A. Current recommendations and practice of oxygen therapy in preterm infants. *Clinics in Perinatology*, 46(3):621–636, September 2019.

[33] Tin W, Carlo W, and Gupta S. Oxygenation targeting and outcomes in preterm infants: The new evidence. *The Newborn Lung: Neonatology Questions and Controversies*, pages 301–329, 2012.

[34] Lau YY, Tay YY, et al. Maintaining optimal oxygen saturation in premature infants. *The Permanente Journal*, 15(1):e108–113, 2011.