

Vogel, Justus; Cordier, Johannes; Filipovic, Miodrag

Working Paper

Causal Effects and Optimal Policy Learning for Intensive Care Unit Discharge Decisions to Solve Hospital Process Bottlenecks: Approach, Methods, and First Results

Working Paper Series in Health Economics, Management and Policy, No. 2025-01

Provided in Cooperation with:

University of St.Gallen, School of Medicine, Chair of Health Economics, Policy and Management

Suggested Citation: Vogel, Justus; Cordier, Johannes; Filipovic, Miodrag (2025) : Causal Effects and Optimal Policy Learning for Intensive Care Unit Discharge Decisions to Solve Hospital Process Bottlenecks: Approach, Methods, and First Results, Working Paper Series in Health Economics, Management and Policy, No. 2025-01, University of St.Gallen, School of Medicine, Chair of Health Economics, Policy and Management, St.Gallen

This Version is available at:

<https://hdl.handle.net/10419/308439.2>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc-nd/4.0/>

Causal Effects and Optimal Policy Learning for Intensive Care Unit
Discharge Decisions to Solve Hospital Process Bottlenecks: Approach,
Methods, and First Results

Justus Vogel, Johannes Cordier, Miodrag Filipovic



Working Paper Series in Health Economics, Policy, and Management

Editor

Prof. Dr. Alexander Geissler
Professor and Chairholder
Chair of Health Economics, Policy, and Management
School of Medicine
University of St.Gallen

Editorial office

Jonas Subelack
Research Assistant
Chair of Health Economics, Policy, and Management
School of Medicine
University of St.Gallen

The entire series of publications is available on our website at: <https://med.unisg.ch/en/research/health-care-management/publications/>

Causal Effects and Optimal Policy Learning for Intensive Care Unit Discharge Decisions to Solve Hospital Process Bottlenecks: Approach, Methods, and First Results

Keywords: Causal Machine Learning, Intensive Care Unit Management, Hospital Operations, Policy Learning

JEL Classification: I10, C44

Authors:

Justus Vogel

Scientific Project Leader

Chair of Health Economics, Policy, and Management, School of Medicine, University of St.Gallen

justus.vogel@unisg.ch

Johannes Cordier

Research Assistant

Chair of Health Economics, Policy, and Management, School of Medicine, University of St.Gallen

Miodrag Filipovic

Department Head

Clinic for Operative Intensive Care Medicine, Cantonal Hospital St. Gallen

Recommended citation:

Vogel, Justus; Cordier, Johannes; Filipovic, Miodrag (2025): Causal Effects and Optimal Policy Learning for Intensive Care Unit Discharge Decisions to Solve Hospital Process Bottlenecks: Approach, Methods, and First Results. Working Paper Series in Health Economics, Management and Policy, No. 2025-01, University of St.Gallen, School of Medicine, Chair of Health Economics, Policy and Management, St.Gallen.

Causal Effects and Optimal Policy Learning for Intensive Care Unit Discharge Decisions to Solve Hospital Process Bottlenecks: Approach, Methods, and First Results

Dr. Justus Vogel*,^c Johannes Cordier*, and Prof. Dr. med. Miodrag Filipovic⁺

* Chair of Health Economics, Policy and Management, School of Medicine, University of St. Gallen, St.-Jakob-Strasse 21, CH-9000 St. Gallen, Switzerland

⁺ Cantonal Hospital of St. Gallen, Rorschacher Strasse 95, CH-9000 St. Gallen

^c Corresponding Author

Declarations

Acknowledgements

We thank Daria Bukanova-Berend for supporting us with the literature search which was necessary for this study. We thank Patrick Mürger for extracting all data necessary for this study. We thank our student assistant David Klug for his support in data cleaning and descriptive analyses.

Funding

This research is supported by a Stepping Stone Grant of the University of St. Gallen (project number 2300181).

Conflict of interest

None.

Author contributions

JV: Conceptualization, Methodology, Formal analysis, Interpretation, Writing – Original Draft, Visualization, Supervision, Project administration, Final approval, Funding acquisition

JC: Conceptualization, Methodology, Data curation, Formal analysis, Interpretation, Writing – Review & Editing, Final approval

MF: Conceptualization, Data acquisition, Interpretation, Writing – Review & Editing, Final approval

Ethics approval and consent to participate

This study was approved by the Ethical Commission of Eastern Switzerland (project number 2024-00829).

Consent for publication

Not applicable.

Data availability

Anonymized personal health data with restricted access.

Code availability

Available from the corresponding author upon request.

Abstract

Intensive care units (ICUs) operate with fixed capacities and face uncertainty such as demand variability, leading to demand-driven, early discharges to free up beds. These discharges can increase ICU readmission rates, negatively impacting patient outcomes and aggravating ICU bottleneck congestion. This study investigates how ICU discharge timing affects readmission risk, with the goal of developing policies that minimize ICU readmissions, managing demand variability and bed capacity.

To define a binary treatment, we randomly assign hypothetical discharge days to patients, comparing these with actual discharge days to form intervention and control groups. We apply two causal machine learning techniques (generalized random forest, modified causal forest). Assuming unconfoundedness, we leverage observed patient data as sufficient covariates. For scenarios where unconfoundedness might fail, we discuss an IV approach with different instruments.

We further develop decision policies based on individualized average treatment effects (IATEs) to minimize individual patients' readmission risk. Our sample comprises 12,950 ICU stays (11,873 unique cases) from the Department of Surgical Intensive Medicine of the Cantonal Hospital of St. Gallen admitted between January 01, 2016, and December 31, 2023. We find that for 72% of our sample discharge at point in time t as compared to $t + 1$ increases patients' readmission risk. Vice versa, 28% of cases profit from an earlier discharge in terms of readmission risk. The range of IATEs is quite large: For 91.4% of ICU stays, an earlier ICU discharge changes a patient's readmission risk between -0.05 and 0.05 percentage points (-55% and 55% relative change as compared to the average readmission rate of 9.04%).

To develop decision policies, we will exploit this treatment heterogeneity and rank patients according to their IATEs and compare IATEs of optimal and actual discharges across all decision points in our observation period. Finally, we outline how we will assess the potential reduction in readmissions and saved bed capacities under optimal policies in a simulation, offering actionable insights for ICU management.

We aim to provide a novel approach and blueprint for similar operations research and management science applications in data-rich environments.

1 Introduction

The intensive care unit (ICU) of a hospital treats critically ill patients often suffering from life-threatening diseases (Gopalan and Pershad, 2019; Milbrandt et al., 2008; Nates et al., 2016). An ICU is characterized by scarce, costly capacities such as high-end hospital beds and medical equipment, and specialized physicians and nurses. These capacities are fixed in the short- to mid-term. In addition, patient demand and the patients' complexity mix are uncertain (Dobson et al., 2010; Thirumalai et al., 2024). Patients arrive due to scheduled surgeries, e.g., for post-surgery observation, or as external or internal emergencies. Prediction of the volume of incoming patients is difficult, and rarely done in practice. Additionally, patients' length of stay in the ICU is also uncertain as it depends, among other factors, on patients' main diagnosis, co-morbidities, performed procedures, and the progression of patients' health status. Lastly, process times of prior process steps, i.e., emergency care or surgery, are much smaller (hours) than patients' length of stay in the ICU (days). Fixed capacities, uncertain and unscheduled incoming demand, uncertain process times, and faster throughput time of upstream processes make the ICU into a classic example of a process bottleneck (Bai et al., 2021; Chan et al., 2012). Unlike classic production settings, inventory cannot be built between process steps and areas, however: Arriving patients are in great need for care and thus need intensive care immediately. If the ICU's capacity is fully utilized, one option is to discharge a patient with a relatively low need for intensive care (Berk and Moinszadeh, 1998; Dobson et al., 2010). Discharging a patient in such a setting is referred to as demand-driven or early discharge, that is, if there were no newly arriving patient, the patient to be discharged would stay longer in the ICU, profiting from additional intensive care (Bai et al., 2021, 2018; Chan et al., 2012; Ouyang et al., 2020).

Demand-driven discharges, in turn, are linked to higher ICU readmission rates (Kramer et al., 2013, 2012; Niven et al., 2014). Indeed, the health status of patients discharged in demand-driven settings are more likely to worsen downstream, requiring a readmission to and additional stay in the ICU. Such readmissions not only gravely negatively impact patients' health and outcomes (Mcneill and Khairat, 2020; Rosa et al., 2020) and hospitals' bottom line, but they also additionally clog the ICU process

bottleneck – and might trigger additional demand-driven discharges, potentially setting off a vicious cycle (KC and Terwiesch, 2011).

We define an ICU discharge on a given day as the variable of interest (“treatment” in the econometric sense), causally linked to an outcome (i.e., ICU readmission). We formulate how a patient’s discharge on a given day compared to a discharge the following day causally effects a patient’s risk of readmission. Our goal is to discover a decision policy that minimizes ICU readmissions causally linked to (early) ICU discharge. To this end, we plan to show that in situations where there is more than one candidate that could be discharged to accommodate a newly arriving patient, readmission risk is minimized by discharging the patient for whom the effect of discharge on readmission risk is smallest.

The gold standard for causal inference is a randomized experiment or a randomized controlled trial. Such a study design is very difficult or rather impossible to implement – both from an ethical as well as operational point of view – for our research endeavor. Thus, we present an approach that uses observational data and causal inference under the selection on observables assumption, also referred to as unconfoundedness or conditional independence assumption. Our research questions for developing an empirical approach and methodology are:

- I How can the causal effect of an ICU discharge on a given day on the ICU readmission risk of a patient be estimated with observational data?
- II How can a decision policy be learned that minimizes ICU readmissions?

Research questions I and II structure our research objective in two concrete steps: First, estimating the causal effect and second, learning a decision policy based on the estimated causal relationship, exploiting effect heterogeneity. Note that both steps need to consider individual patient characteristics and individualized average treatment effects (IATEs) to enable decision makers to make a discharge decision in a setting of uncertainty.

There is a rich operations management (OM) and medical literature on ICU management (Bai et al., 2018; Gopalan and Pershad, 2019; Niven et al., 2014), some of which we will review in the next chapter. As for many operations research settings (Ho et al., 2017), OM models address the ICU process

bottleneck problem with normative mathematical models, aiming to develop decision policies from a theoretical framework. Commonly, studies then test and/ or calibrate their models with (small) samples from one or several hospitals (e.g., Bai et al., 2021; Chan et al., 2012), or perform a simulation study (Ouyang et al., 2020).

In this paper, we outline that due to the causal relationship of ICU discharge and ICU readmission and individual patient characteristics creating variation and uncertainty, a different approach might be more effective, generalizable, and better scalable than traditional OM methods. While we also develop our model from a theoretical start point, its main strength comes from its empirical application and practicality. We aim to show that machine learning methods from the family of doubly robust learners can answer our research questions yielding data-driven decisions that optimize medical quality and free up scarce capacity. We believe that our approach is not only valid for ICU management and other hospital operations settings but that our study can serve as a blueprint for a more general application of causal machine learning methods in operations research.

This section continues with a review of related work. Section 2 presents a formal problem statement. Section 3 discusses our approach for estimating causal effects and developing optimal decision policies. We also present first descriptive results, and first estimations of causal effects. Lastly, we outline a simulation study to assess the practical utility of our approach. Section 4 summarizes and gives an outlook.

Related work

We apply machine learning methods for causal inference and policy learning to a hospital operations problem, more specifically for resolving one reason for bottleneck congestion of a key hospital process area, the ICU. Thus, related work of our study are studies in the OM literature focusing on (1) machine learning applications, (2) causal inference and causal machine learning, (3) policy learning, and (4) hospital operations.

Machine Learning Applications

Regarding machine learning applications, many studies focus on prediction in supervised learning settings. Chou et al. (2023) provided a review of studies published in leading OM journals applying a variety of prediction algorithms such as random forest, gradient boosting, logistic or probit regression, or support vector machine. The authors note that the aim of the majority of studies is to show that machine learning prediction techniques work better than other (statistical) methods commonly applied or are more precise than status quo decision making (i.e., managerial judgement). In the healthcare context, Ang et al. (2015) employ a Q-Lasso technique to anticipate waiting times of low-acuity patients in emergency departments of four different hospitals. They show that their method is superior for the prediction and management of waiting times compared to rolling average approaches, “state-of-the-art” in practical emergency department management at the time of the study. Mukhopadhyay et al. (2021) compare the performance of six machine learning classifiers (e.g., neural networks, decision trees, and naïve-Bayes) with a linear programming model they built to detect diseases in four different datasets (two for breast cancer, one for diabetes, and one for diabetic retinopathy). Nenova and Shang (2022) develop an intelligent case-based reasoning approach to predict disease progression for patients suffering from chronic kidney disease and compare results to different machine learning and statistical models. The authors find that their approach outperforms other models. Lastly, Ko et al. (2019) employ machine learning methods to generate insights from textual patient reviews and use these insights to investigate the relationship between waiting time and patient satisfaction.

Causal Inference and Causal Machine Learning

In the last decade, an increasing number of studies employing methods for causal inference under the selection of observables assumption have been published in leading OM journals (Ho et al., 2017; Yilmaz et al., 2024). These studies evaluate effects of, for instance, introducing a new scheduling policy or production program, commonly using regression discontinuity designs or pre-post intervention study designs such as difference-in-differences or synthetic control methods. While we also aim to causally infer effects of a treatment (i.e., ICU discharge at a given point in time) on an outcome (i.e., patient’s readmission risk) under the selection on observables assumption, our setting differs to traditional policy

or program evaluation settings: We do not know the decision policy for selecting a patient for discharge yet but develop a decision policy minimizing overall readmission risk. To this end, we need to estimate conditional average treatment effects (CATEs), more specifically the effect of a discharge at point in time t for a patient group (group average treatment effect, GATE), and at the most detailed aggregation level of patient features (individualized average treatment effect (IATE)) (Lechner and Mareckova, 2024). In the last years, causal forests were refined to estimate heterogenous treatment effects (HTE), which can be used for GATE and IATE estimation (Athey et al., 2019; Athey and Imbens, 2016; Chernozhukov et al., 2018; Feuerriegel et al., 2024; Lechner and Mareckova, 2024; Wager and Athey, 2018; Wang et al., 2021). Causal forests are suitable for our research setting which we outline in more detail in section 3.

Causal machine learning methods for the estimation of heterogenous treatment effects have, to the best of our knowledge, only been applied to very few operation research issues. Cui and Davis (2022) employ causal forest (Athey et al., 2019; Wager and Athey, 2018) combined with a difference-in-differences framework to estimate HTEs of occupancy taxes for Airbnb bookings on Airbnb listings. Wang et al. (2021) expanded on advances in the machine learning literature, adding an instrumental variable forest (IVF) approach to address endogeneity issues in observational data. The authors apply this approach in another study (Wang et al., 2023) to investigate the heterogeneity of outcome differences between hospitals, to compare the quality ranking of hospitals according to patient-centric versus population-average information, and to simulate implications for pay-for-performance reimbursement schemes. In their first study, Wang et al. (2021) show with experiments, that in settings of endogeneity, their IVF approach is superior to a causal forest approach without instrumental variable to correct for biases. They also show that IVF outperforms causal forests with instrumental variable as part of the generalized random forests (GRF) (Athey et al., 2019), especially when the instrumental variable is continuous. The authors conclude that IVF should be used when unconfoundedness cannot be argued and only continuous instrumental variables are available. Endogeneity issues and potential strategies to offset endogeneity such as instrumental variables are well-known and discussed in the OM literature (Lu et

al., 2018). While we confidently argue for unconfoundedness, as sensitivity analysis, we will also employ an instrumental variable approach (see section 3).

Policy Learning

Recently, methodologies and applications for policy learning have received increasing attention in the OM literature. Policy learning generally means employing machine learning methods to optimally assign treatment decisions, such as whom to offer a discount or whom to target with a certain drug, based on individuals' characteristics (Zhan et al., 2024). Athey and Wager (2021) developed the policy tree approach, a double robust machine learning technique for policy learning, which has been applied, discussed, and refined by several authors. Zhan et al. (2024), for instance, study how to extend policy learning when employing adaptively collected data, focusing on the two key challenges of (1) dependent samples in ex post optimal policy learning with adaptive samples, and (2) low assignment rates of different types of treatment to different types of individuals when using adaptive assignment rules. Munro (2024) explores how, under personalization of treatment assignment, strategic behavior of individuals trying to increase their chances to obtain a (better) treatment, might prevent optimal assignment. The author conducts a sequential experiment with Bayesian optimization to showcase a solution for managing individuals' strategic behavior.

Regarding applications, in their working paper, Cordier et al. (2023) combine causal forest with policy tree to learn what patients to optimally select for rapid recovery, a specific post-surgery treatment path for joint replacement patients, under capacity constraints. The authors show that optimal assignment rules can increase overall patient welfare, measured by patients' expected improvement of joint functionality, by 17.56% without increasing hospital capacity. Based on their insights, the authors derive implications for surgery scheduling and clinical decision support. Yang et al. (2023) impute long-term outcomes, namely long-term revenue from online subscribers of *The Boston Globe*, as input to learn an optimal policy for whom to target with discounts. The authors show that their imputation and policy learning strategy is comparable to a policy learned from ground-truth long-term outcomes and superior to status quo policy learning using short-term proxies for long-term outcomes.

Hospital Operations

Healthcare has become one of the core fields in operations research (Browning, 2020; Green, 2012). One important part of this research field is hospital operations management (Keskinocak and Savva, 2019). Shi et al. (2021), for instance, design a Markov decision process to provide decision support for how many and what patients to discharge from a hospital on a given day, balancing readmission risk and inpatient congestion. They base these decisions on a personalized readmission prediction model, incorporating several statistical methods. Demir (2014) uses classification trees for predicting unplanned readmissions within 45 days after discharge for patients suffering from chronic obstructive pulmonary disease or asthma, and compares prediction results to other methods (e.g., logistic regression). Lastly, in their study, Berk and Moinzadeh (1998) find early discharges to be a viable option for increasing scarce capacities.

Within hospital operations management, many studies investigate ICU management (Bai et al., 2018). Chan et al. (2012) design an elaborate state and action space and dynamic optimization model to minimize the number of additional ICU inpatient days attributable to ICU readmissions linked to demand-driven discharges. Grand-Clément et al. (2023) design a Markov decision process model that uses predictions of patient deterioration supplied by machine learning models, to decide what patients should pro-actively be resubmitted to the ICU to prevent further patient deterioration, a longer length of stay, and deaths. They develop robust decision policies for transferring patients to the ICU, while stressing the effect of parameter uncertainty on average ICU occupancy, patients' length of stay, and mortality.

To reduce ICU process bottleneck congestion, apart from early discharges, hospitals could reject external or internal emergencies, and/ or cancel or reschedule elective surgeries in the case of surgical ICUs. Chan et al. (2012) consider these options in an extension to their model, and Bai et al. (2021) model a (stationary) discrete time Markov decision process considering these options and early discharge of a low- or high-severity patient as possible actions. While patient rejection can potentially alleviate bottleneck development (Scheulen et al., 2001), studies indicate several drawbacks (Chalfin et al., 2007;

McConnell et al., 2005; Renaud et al., 2009; Schull et al., 2004). For instance, rejecting patients means lost revenue for a hospital, potentially negative impact on patient satisfaction when cancelling or rescheduling elective surgeries, and worse patient outcomes for rejected or diverted patients.

Besides, Ouyang et al. (2020) develop a mathematical model to support bed allocation decisions in the ICU. They develop policies for minimizing the long-run average mortality rate, and propose heuristic decision methods with relaxed, more general assumptions. They show the effectiveness of their heuristic methods in a simulation study. Lastly, they extend their model to deliberately allow for queuing for ICU admission and for readmission from the general ward.

Other work in the field of hospital operations management include an investigation of bottleneck and congestion spillover of one hospital unit to other units (Kim et al., 2023), the effect of advanced electronic medical records on patient safety (Hydari et al., 2018), optimal and predictive bed assignment (Bertsimas and Pauphilet, 2023), development of ICU admission policies by estimating the effect of denying admission on different patient outcomes (Kim et al., 2014), prediction of congestion risk in the ICU (Bravo et al., 2024), reduction of alarm fatigue (Piri et al., 2022), and understanding and better serving the needs of patients not critical enough to be treated in the ICU but too critical to be treated on a general ward (Chan et al., 2018).

There are also several studies in the medical literature investigating ICU discharge and readmission decisions. There are, for instance, reviews regarding ICU decision making (Gopalan and Pershad, 2019), studies informing guideline development (Nates et al., 2016), qualitative studies exploring the experience of physicians and nurses involved in and/ or affected by these decisions (Ofoma et al., 2018), and machine learning models predicting ICU readmission (Hammer et al., 2020; Loreto et al., 2020), which has also been studied by data scientists (Barbieri et al., 2020).

With our approach and methods, we do not only add to the operations research literature, but also to the literature on precision medicine, clinical decision support and optimal treatment selection (Feuerriegel et al., 2024; Prosperi et al., 2020). Finally, we refer to Prosperi et al. (2020) who correctly point out that machine learning models used for prediction do not deliver causal treatment suggestions,

which is the main reason we are employing causal machine learning methods for causal inference rather than binary classifiers.

2 Formal Problem Statement

Conventional Models

Past OM studies have proposed Markov decision processes described by state and action space models to derive decision policies optimizing costs associated with demand-driven discharges and, by model extension (Chan et al., 2012) or as part of the main model (Bai et al., 2021), also rejection of arriving patients. They define the time period t with $t = \{1, 2, \dots, T\}$ to be small enough so only one patient can arrive in t . The arrival of a patient might require an action (e.g., demand-driven discharge), first and foremost in situations where the ICU bed capacity B is exceeded. Such an action, among other things such as changes of other patients' health status and "natural" discharges (service completion), changes the current state and defines the next state. The optimal policy, e.g., rejecting a patient, or early discharging one type rather than another type of patient, is then derived by minimizing the costs associated with the action associated with the policy (e.g., the additional ICU inpatient days "caused" by readmissions due to ICU discharge). While the authors show how such models can improve ICU management, we have three points of critique that impede transferability, scalability, and generalizability of existing models:

- 1) Causality: Costs or rewards of actions, i.e., the outcomes, are not causally linked to the policy, i.e., the treatment, in the econometric sense. Identifying assumptions, such as unconfoundedness or exogeneity, are not discussed and argued. Thus, it is unclear whether a change in costs is attributed to a policy (e.g., discharge patient x_1 rather than patient x_2 at point t) or to confounders.
- 2) Fixed and normative inputs: Costs or rewards are calculated with fixed inputs, e.g., the average length of stay of patients in the ICU, or normatively set money values. This reduces variation which empirically exists and could in fact be observed in the costs or rewards of actions. Such simplifications increase in sample performance but limit generalizability and scalability.

- 3) Bias and uncertainty of transition probabilities: Markov decision processes require identifying probabilities for transitioning from one state to another. In dynamic and complex settings, such probabilities have to be identified for many parameters (Bertsekas, 2012). A single biased transition probability will bias the whole model. Additionally, the uncertainty of a point estimate of transition probabilities is usually not considered in Markov decision process models (Zhang et al., 2019). In other words, Markov decision process models implicitly assume point estimates from empiric data or expert estimates to be true probabilities. Indeed, uncertain transition probabilities have received attention in the literature for decades (e.g., Satia and Lave, 1973), and there are several approaches to address this issue (Delgado et al., 2011; Mastin and Jaillet, 2012; Zhang et al., 2019).


Practical ICU decision making: A novel problem statement

We propose a novel approach for addressing the limitations presented above. Firstly, we define a model based on practical ICU decision making which needs to resolve the bed capacity constraint $B_a \geq 0$ over the course of each day. Secondly, to help decision makers solve this constraint, we design policies that are generalizable, learnable with causal machine learning methods and re-learnable with the same algorithms and similar data from other hospitals, and thus scalable. Note also that state and action space models were shown to improve ICU decision making on a tactical level, e.g., when deciding whether and how much ICU capacity should be reserved to avoid more costly rejection and/ or early discharge (Bai et al., 2021). With our model, we will provide decision support for operational decision making in the ICU, i.e., for decisions that must be made routinely and daily.

We use the flow of a patient through a hospital as the starting point of our problem statement (cf. Figure 1, loosely based on Bai et al. (2021, 2018) and Litvak et al. (2008)). We are interested in reducing readmission flows (1), (2), and (3). Note, however, that at our partner hospital, we investigate a mixed ICU also accommodating intermediate care patients. Thus, we are specifically investigating readmission flows (1) and (3), i.e., the readmission of a patient to a higher level of care unit.

Kramer et al. (2013), for instance, report that the median readmission rate at the more than 100 ICUs they investigated was at 5.9% (interquartile range between 5.1% and 7.0%) and Hosein et al. (2014) found in a meta-analysis that readmission rates typically are between 4% and 6%. While a certain percentage of ICU readmissions appears to be nonpreventable (Al-Jaghbeer et al., 2016), preventable readmissions and especially those causally linked to discharges could be termed as rework and correction in Lean Management terms, adding to bottleneck congestion.

Figure 1: Hospital patient flow



Annotations: Patients can enter a hospital as scheduled, plannable cases (full arrows) through the outpatient clinic or as unscheduled/ emergency cases through the emergency department (dashed arrows). From the emergency department, patients are pushed onto the process area with free capacity and/ or where they need to receive care. Unscheduled transfers from other hospitals disrupt hospital system management further. Patients are pushed through the different process areas in a scheduled or oftentimes unscheduled manner. Demand planning and corresponding supply planning only regularly occur in the central operating room area, other process areas are commonly staffed corresponding to their full capacity (e.g., to serve all beds on a ward). We are particularly interested in how to minimize readmission flows (1) and (2), often causally linked to early, demand-driven discharges of patients in the ICU. In such situations, patients with the smallest risk of being readmitted due to the early discharge should be discharged, and not necessarily those with the lowest predicted readmission risk (cf. Athey (2017), Feuerriegel et al. (2024), and Prosperi et al. (2020) for a discussion of this topic).

The basis for our formal problem statement is the capacity constraint $B_a \geq 0$, denoting that the available ICU bed capacity B_a must always be equal to or greater than zero operatable ICU beds throughout a given day. In a naïve state, B_a is solely defined by the overall ICU bed capacity on a given day, and the sum of patients who will not be discharged and thus occupy the ICU $\sum_{i=1}^I x_i$ with $i = \{1, 2, \dots, I\}$ and $x = 1$ if x_i resides in the ICU at least for one more day, $x = 0$ otherwise:

$$B_a = B - \sum_{i=1}^I x_i \quad (1)$$

Consulting Figure 1, the first patient flow decision-makers in a surgical ICU, as at our partner hospital, must consider and incorporate into (1) is the sum of elective surgery patients with a planned post-

surgery ICU stay $\sum_{e=1}^E a_e$ on a given day, with $e = \{1, 2, \dots, E\}$, and $a = 1$ if a_e has a planned post-surgery ICU stay, $a = 0$ otherwise. In addition, planned discharges occur because a patient has reached ICU service completion (natural discharge), denoted by d_n with $n = \{1, 2, \dots, N\}$, and $d = 1$ if d_n is naturally discharged during the day, $d = 0$ otherwise. This expands the capacity constraint given in (1) into:

$$B_a = B - \left(\sum_{i=1}^I x_i + \sum_{e=1}^E a_e \right) + \sum_{n=1}^N d_n \quad (2)$$

If overall ICU capacity and natural discharges were sufficient to accommodate $\sum_{i=1}^I x_i$ and new arrivals $\sum_{e=1}^E a_e$, and if we stayed in this simplified decision framework, we could be satisfied with the definition of available ICU bed capacity as outlined in (2).

However, even if there were no unscheduled, external emergency patient arrivals, there are three additional considerations to make (cf. Figure 1). Given the capacity constraint $B_a \geq 0$, there will be days where $\sum_{e=1}^E a_e$ will require discharges in addition to $\sum_{n=1}^N d_n$, which we define as planned early, demand-driven discharges denoted by d_p^{early} with $p = \{1, 2, \dots, P\}$ and $d = 1$ if d_p^{early} is planned to be early discharged, $d = 0$ otherwise. Secondly, another lever to satisfy $B_a \geq 0$ is to reject patients, i.e. to cancel or postpone elective surgeries with a planned post-surgery ICU stay, denoted by $a_c^{elective}$ with $c = \{1, 2, \dots, C\}$, and $a = 1$ if $a_c^{elective}$ is cancelled, $a = 0$ otherwise. Thirdly, the capacity constraint must also hold when considering readmissions from downstream units, denoted by a_r with $r = \{1, 2, \dots, R\}$, and $a = 1$ if a_r is readmitted, $a = 0$ otherwise. These three considerations expand B_a to:

$$B_a = B - \left(\sum_{i=1}^I x_i + \sum_{e=1}^E a_e - \sum_{c=1}^C a_c^{elective} + \sum_{r=1}^R a_r \right) + \left(\sum_{n=1}^N d_n + \sum_{p=1}^P d_p^{early} \right) \quad (3)$$

In a setting without unscheduled arrivals, e.g., in an orthopedic hospital exclusively treating elective surgery patients, (3) would suffice to describe B_a and the discharge decision processes to satisfy $B_a \geq 0$. Note that there still is uncertainty in service time, i.e., the number of ICU inpatient days for each x_i , and thus prediction of $\sum_{i=1}^I x_i$ and $\sum_{n=1}^N d_n$ are prone to uncertainty. Further note that at the start of a day (e.g., between 07:00 a.m. to 08:00 a.m. at our partner hospital), ICU decision-makers might initiate their

decision process at (2), and if they foresee that the capacity constraint will not be satisfied at any point in time throughout the day, they will try to satisfy the capacity constraint by balancing $\sum_{c=1}^C a_c^{elective}$, and $\sum_{p=1}^P d_p^{early}$. While we appreciate that past studies must make this assumption for their models, cancellations and planned demand-driven discharges do not happen simultaneously to patient arrival. ICU decision-makers rather anticipate the number of planned ICU admissions according to the operating room schedule at the start of a day and compare these with the available ICU bed capacity after natural discharges (cf. (2)).

Considering unscheduled or emergency/ urgent patients coming from upstream units, i.e., directly from the emergency department or, by a roundabout route, from the central operating room area, or as transfer-ins from other hospitals, the decision process to satisfy $B_a \geq 0$ becomes more complex. Note that in practice, experienced ICU decision-makers might anticipate unscheduled ICU admissions and thus reserve some ICU capacity by discharging more patients than absolutely needed to satisfy $B_a \geq 0$ as defined in (3). Indeed, past studies such as Bai et al. (2021) quantitatively derive exactly how much capacity should be reserved balancing the costs of patient rejection, and demand-driven discharge. In fact, one could consider that the capacity constraint actually is $B_a \geq \sum_{u=1}^U a_u^{ant}$, that available ICU capacities must be at least the total number of unscheduled arrivals/ admissions the experienced ICU decision-maker anticipates on a given day ($u = \{1, 2, \dots, U\}$, and $a = 1$ if a_u^{ant} is anticipated as unscheduled arrival, $a = 0$ otherwise). As it is uncertain how many unscheduled patients will arrive exactly, we additionally define a_v^{add} as the unforeseen unscheduled arrivals in addition to $\sum_{u=1}^U a_u^{ant}$, with $v = \{1, 2, \dots, V\}$ and $a = 1$ if the arrival of a_v^{add} is unscheduled and is not anticipated, $a = 0$ otherwise. Note that decision makers might increase the sum of planned early discharges $\sum_{p=1}^P d_p^{early}$ and/ or the sum of cancellations of elective surgeries $\sum_{c=1}^C a_c^{elective}$ to offset $\sum_{u=1}^U a_u^{ant}$, yet to manage any

a_v^{add} , the only lever is to demand-driven discharge a patient in an unplanned manner, denoted by d_q^{early}

with $q = \{1, 2, \dots, Q\}$ and $d = 1$ if d_q^{early} is early discharged in an unplanned way, $d = 0$ otherwise.¹

To satisfy $B_a \geq 0$, decision makers then consider that

$$B_a = B - \left(\sum_{i=1}^I x_i + \sum_{e=1}^E a_e - \sum_{c=1}^C a_c^{elective} + \sum_{r=1}^R a_r + \sum_{u=1}^U a_u^{ant} \right) + \left(\sum_{n=1}^N d_n + \sum_{p=1}^P d_p^{early} \right) \quad (4)$$

$$+ \left(\sum_{q=1}^Q d_q^{early} - \sum_{v=1}^V a_v^{add} \right)$$


Note that planned discharges, both natural and demand-driven, happen at a later point in time than the discharge decision. At our partner hospital, $\sum_{n=1}^N d_n$ and $\sum_{p=1}^P d_p^{early}$ are decided in the early morning of a (week-) day, while discharges occur in the late morning to early afternoon, in close alignment with the downstream care units admitting the patient(s). Should a patient have worsened between initial discharge decision and actual discharge, decision-makers might re-think and change their initial decision.

Empirical evidence for formally defined decision constraint

In our data, there is empirical evidence for the dynamics of the decision constraint as formulated in equation (4). Figure 2 shows the typical arrival pattern at our partner hospital, depicting the share of patients arriving per hour of the day.

¹ Additionally, one could argue that if a_v^{add} arrived early in the day before surgery has started for the last a_e , decision-makers might also have the option to cancel one additional $a_c^{elective}$. For simplicity, we do not consider this option.

Figure 2: Arrival pattern





Annotations: The analysis includes all ICU stays of patients with a positive or waived general consent admitted at our partner hospital between January 01, 2016 and December 31, 2023, no exclusion criteria applied (14,121 stays and 12,932 unique cases; cf. Figure 4 below).

In sum, roughly 72% of all patients arrive between 10:00 a.m. and 09:00 p.m. Still, the volume of patient admissions (28% of all admissions) between 09:00 p.m. and 10:00 a.m. is not negligible. The share of discharges in the same time window is quite small (see Figure 3). This is evidence that ICU decision makers incorporate planned admissions after surgery $\sum_{e=1}^E a_e$ into their decision making process and that they do anticipate unscheduled admissions $\sum_{u=1}^U a_u^{ant}$ and possibly also readmissions $\sum_{r=1}^R a_r$.

Discharges typically occur between 09:00 a.m. and 04:00 p.m. (roughly 95% of all discharges for Panel A, 97% for Panel B), and most discharges happen in an even closer time window between 10:00 a.m. and 02:00 p.m. (roughly 86% for Panel A and 89% for Panel B). This is evidence for the planned discharge decisions $\sum_{n=1}^N d_n$ and $\sum_{p=1}^P d_p^{early}$ happening early in the morning with corresponding discharges a few hours later.

Figure 3: Discharge pattern



Annotations: Panel A includes all ICU stays of patients with a positive or waived general consent admitted at our partner hospital between January 01, 2016 and December 31, 2023, no exclusion criteria applied (14,121 stays and 12,932 unique cases; cf. Figure 4 below). In Panel B, we exclude patients who died during their ICU stay (589 stays, 518 unique cases), as the timing of these patients' "discharge" does not occur according to the decision process formalized in equation (4).

We also see that discharges in “off-hours” between 02:00 p.m. and 10:00 a.m. (14% (11%) of all discharges for Panel A (B)) and especially between 04:00 p.m. and 09:00 a.m. are rather rare (5.4% (2.9%) of all discharges for Panel A (B)). Especially for Panel B, where we exclude patients who died during their ICU stay, we may assume that a considerable share of these off-hour discharges are unplanned demand-driven discharges $\sum_{q=1}^Q d_q^{early}$ triggered by unforeseen unscheduled admissions $\sum_{v=1}^V a_v^{add}$. Figure 3 clearly shows that these situations occur comparatively rarely. At the same time, a considerable share of admissions does occur in these off-hours, namely roughly 57% between 04:00 p.m. and 09:00 a.m. This is evidence that – at least in our empirical setting – demand-driven discharges rarely occur as defined in conventional operations research studies.

Summary

In summary, we have ascertained three different types of points in time where ICU discharge decisions occur or are changed:

- 1) Planned discharges of any kind are made in the morning in the first hours of a physician’s shift, typically between 07:00 a.m. and 08:00 a.m. at our partner hospital
- 2) A discharge decision is changed should a patient’s health status change considerably between the time of the discharge decision and the planned discharge time
- 3) Unplanned discharge decisions are made at any point during the day (and night) if a_v^{add} occurs

In equation (4), we formulate the decision problem for satisfying the capacity constraint $B_a \geq 0$ that decision makers face at each of these three types of points in time. Our model supplies decision-makers at each of these types of points in time with the effect of a discharge on a patient’s readmission risk. Note that while our model can inform any kind of demand-driven discharge, it also provides decision support for natural discharges. Should the result of our model reveal that the change in readmission risk is too large due to the discharge, the “natural” discharge could be postponed.

3 Causal Effects and Optimal Policy Learning

We start with a short description of our dataset and descriptive analyses. Then, we outline a three-step approach to enable optimal ICU discharge decision making: (1) We present how to estimate the ATE,


CATEs, and IATEs of an ICU discharge at a point in time t on the ICU readmission risk in an observational study setting with double robust learners, (2) we outline how we plan to develop decision policies based on patients' IATEs, and (3) we present how we plan to apply these policies to our empiric data to gauge how many ICU bed capacities could have been saved with optimal decisions. In the first section, we also discuss identifying assumptions for causal inference in the context of our study.

Overview of Dataset and Descriptive Analyses

Dataset

Our dataset contains clinical and basic data of all ICU stays without a documented negative general consent admitted to the Department of Surgical Intensive Care Medicine between January 01, 2016, and December 31, 2023. Figure 4 shows the inclusion and exclusion criteria and corresponding samples for the different analyses we conducted.

Figure 4: Inclusion and exclusion criteria and corresponding samples



Annotations: Reasons for excluding cases that are not candidates for readmission are discharge to another ICU, to home, to a rehabilitation clinic or nursing home, or to a different hospital (cf. Kramer et al. (2013)).

Our data is recorded at ICU stay level. Each ICU stay has a unique identifier. One case, given by a unique case number, has at least one ICU stay. Theoretically, one patient could be admitted to a hospital

several times in a year. In such cases, a new unique case number is defined for each hospital admission. A case number for which two or more ICU stays are recorded thus indicates that a patient was readmitted to the ICU within the same hospital stay.

Sample A was used for analyses regarding arrival and discharge patterns (e.g., Figure 2 above). Sample B includes only those cases and stays that are useable for causal inference. To this end, we excluded patients who died in the ICU as these are recorded as discharges yet are not connected to a deliberate discharge decision, excluded cases that could not be readmitted to the ICU as they were discharged to a different ICU, to home, to a rehabilitation clinic or nursing home, or to a different hospital, and we excluded cases for whom the discharge reason was missing.

Overall, we can extract and engineer more than 4,600 features from our dataset (see Table 1). Most of these features are related to medication and drugs (more than 4,000 features). We include one feature per substance, dosage and volume unit and values represent the given volume (either 0 or a continuous number) within the last 24 hours before discharge (intervention group) or “simulated” discharge (control group, see Figure 6 below). Other major feature categories are laboratory diagnostic values (more than 200 features), vital signs (20 features), clinical (risk) scores (close to 40 features) and basic data (age, gender, weight, height, BMI).

Table 1: Overview of used features

Feature category	Number of features	Numeric type	Considered measurements	Description and examples
Medication	4,119	Continuous	Last 24 hours	Hundreds of different substances with at least one, oftentimes several dosages and volume units, e.g., Noradrenalin peripheral in microgram, Noradrenalin in microgram
Laboratory tests	208	Continuous	Last measurement	208 different laboratory values, e.g., Kalium, Creatinine, Cholesterol, HbA1c
Clinical scores	40	Continuous	Last 2 measurements	10 different scores, e.g., SAPS II, SAPS 3, NEMS, GCS, etc.
Diagnoses	114	Dummy	Time invariant	57 indication areas for first and follow-up diagnosis, e.g., "diseases of the liver and biliary tract"
Infusions	35	Continuous	Last 24 hours	22 different infusions, partially with more than one dosage, e.g., Glucose 5%, 10%, 20%, 40%, 50%; all in ml

Feature category	Number of features	Numeric type	Considered measurements	Description and examples
Ventilation mode	30	Dummy	Last 24 hours	30 different types of (non-) invasive ventilation
Therapies	23	Dummy	Time invariant	23 intervention/ surgery groups, e.g., orthopedic interventions
Infections	23	Dummy	Time invariant	For instance, Hepatitis A, B, and/ or C or HIV/ AIDS
Vital signs	20	Continuous	Last 4 hours	Mean, SD, and thresholds for blood pressure, body temperature, heart frequency, oxygen saturation
Blood gas analysis	20	Continuous	Last 24 hours	Mean and SD of ten different values
Basic data	5	Continuous (4), dummy (1)	Time invariant	Gender, age, BMI, height, weight
Blood products	3	Continuous	Last 24 hours	Fresh frozen plasma, platelet concentrate, erythrocyte concentrate; all in ml
Total	4,640			

Annotations: BMI = Body-Mass-Index; GCS = Glasgow Comma Scale; HbA1c = Hemoglobin A1c value; NEMS = Nine equivalents of nursing manpower use score; mg = milligram; ml = millimeters; SAPS = Simplified Acute Physiology Score; SD = Standard Deviation

Descriptive analyses

Table 2 shows descriptive results for patients' basic data.

Table 2: Descriptive results for patients' basic data


Feature	Number of ICU stays	Number of unique cases	Mean (SD)	Median (25 th and 75 th percentile)	Minimum, Maximum
Age	12,950	11,873	63.5 (15.8)	66 (55; 75)	18; 100
BMI	9,602	-	26.4 (5.6)	25.5 (22.7; 29.2)	11.7; 80.2
Height	9,602	8,976	169.9 (9.2)	170 (163; 176)	130; 207
Weight	9,602	-	76.3 (18.3)	75 (64; 86)	27; 275
Number of unique cases			Share		
Male patients	11,873		61.1 %		

Annotations: BMI = Body-Mass-Index; SD = Standard Deviation. Height is in centimeters, and weight in kilograms. For BMI, height, and weight, data was missing or 0 for 3,333 ICU stays and for 15 additional ICU stays, values were below or above meaningful thresholds (130 and 210 centimeters for height, 25 and 400 kilogram for weight) and thus omitted from analysis. In the causal forests, we applied special coding for these ICU stays (BMI, height, and weight all equal to 0). For BMI, and weight, all ICU stays were considered when calculating descriptive statistics as weight could change between the first and subsequent stays. All numbers with decimals are rounded to the shown number of decimals, shown integers are not rounded.

Patients in our sample are on average 64 years old, yet the age variation is considerable as shown by the standard deviation and other descriptive metrics. On average, patients in our sample are slightly overweight, roughly 170 centimeters high and weigh 76 kilograms. Most patients are male.

Figure 5 shows the distribution of number of arrivals and discharges per day.

Figure 5: Distribution of number of arrivals (Panel A) and discharges (Panel B) per day



Annotations:

No admissions occur only for roughly 2% of all days during the observation period. The same is true for discharges. On typical days, ICU staff evidently is considerably busy with admission and discharge processes as two to ten admissions (discharges) must be managed for 88.9% (90.8%) of all days. Eleven or more admissions (discharges) are completed for 1.8% (1.5%) of all days.

Average Treatment Effect Estimation

Potential outcomes framework

To estimate causal effects, we rely on the potential outcomes framework (Rubin, 1974). In an i.i.d. sample, there is a set of patients $i = 1, 2, \dots, n$, each described by a feature vector $X_i \in \mathbb{R}$, a treatment assignment W_i which, in a binary setting, equals 1 if i was treated and 0 otherwise, and an outcome $Y_i \in \mathbb{R}$. The causal effect of a treatment for patient i is given by $Y_i(1) - Y_i(0)$. Accordingly, the average treatment effect (ATE) can be expressed as:

$$\tau = E[Y_i(1) - Y_i(0)] \quad (5)$$

However, for each i only one state $Y_i = Y_i(W_i)$ can be observed. To make progress in observational studies, unconfoundedness (also referred to as selection on observables or conditional independence assumption (CIA) (Lechner, 2001)) is assumed and must be fulfilled (Rosenbaum and Rubin, 1983):

$$\{Y_i(0), Y_i(1)\} \perp W_i \mid X_i \quad (6)$$

With unconfoundedness, it is assumed that treatment W_i is quasi random, given that we control for X_i . In other words, the potential outcome Y_i is independent of the assignment of the treatment W_i , conditional on X_i .

With unconfoundedness, the ATE can be expressed as:

$$\tau = E[\mu_1(X_i) - \mu_0(X_i)] \quad (7)$$

where

$$\mu_w(x) = E[Y_i | X_i = x, W_i = w] \quad (8)$$

To estimate $\hat{\tau}$, we can learn $\hat{\mu}_1(x)$ and $\hat{\mu}_0(x)$ from our data and calculate:

$$\hat{\tau} = \frac{1}{n} \sum_{i=1}^n (\hat{\mu}_1(X_i) - \hat{\mu}_0(X_i)) \quad (9)$$

To estimate $\hat{\tau}$, one could employ a common parametric model such as Ordinary Least Squares (OLS) regression. OLS regression is accurate and consistent as well as unbiased if the underlying linear model is well specified and if its assumptions hold. If the conditional response functions $\mu_w(x)$ are of a non-linear form or if the model's functional form suffers from misspecification, however, parameter estimates will inhibit bias (Wooldridge, 2010).

Non-parametric estimators and (augmented) inverse propensity weighting

Causal forest is a non-parametric machine learning method (Athey et al., 2019; Athey and Imbens, 2016, 2019, 2017; Chernozhukov et al., 2018; Wager and Athey, 2018). Specifically, a random forest (Breiman, 2001) is used to predict $\hat{\mu}_w(x)$, and thus ultimately $\hat{\tau}$. Random forests converge slowly, however, and thus performance with finite samples might still inhibit some bias. Thus, propensity scores, i.e., the probability of receiving treatment based on X_i , are used to weigh $\mu_w(x)$ (inverse-propensity weighting, IPW).

The propensity score $e(x)$ is defined as (Rosenbaum and Rubin, 1983):

$$e(x) = p(W_i = 1 | X_i = x) \quad (10)$$

Under unconfoundedness, the ATE from (5) can then be restated as:

$$\tau = E[Y_i(1) - Y_i(0)] = E \left[\frac{W_i Y_i}{e(X_i)} - \frac{(1 - W_i) Y_i}{1 - e(X_i)} \right] \quad (11)$$

While both outcomes $Y_i(1)$ and $Y_i(0)$ can never be observed for the same unit, the propensity score can and so can (11) be calculated for any i . The unbiased oracle IPW estimator for the ATE is then given by:

$$\hat{\tau}_{IPW}^* = \frac{1}{n} \sum_{i=1}^n \left(\frac{W_i Y_i}{e(X_i)} - \frac{(1 - W_i) Y_i}{1 - e(X_i)} \right) \quad (12)$$

The different author groups listed above have shown (and/ or provided solutions) that the rate of convergence when using machine learning methods, e.g., random forest, to predict $\hat{e}(X_i)$, yielding $\hat{\tau}_{IPW}$, will not be $1/\sqrt{n}$ and thus such estimates will still inhibit some bias. A solution for this is double-robustness. With double-robustness, the estimated ATE $\hat{\tau}$ as formulated in (9) is corrected by the inverse-propensity weighted difference of the observed outcome Y_i and the estimated outcome $\hat{\mu}_w(x)$, also referred to as augmented inverse probability weighting (AIPW) (Chernozhukov et al., 2018; Robins et al., 1994):

$$\hat{\tau}_{AIPW} = \frac{1}{n} \sum_{i=1}^n \left(\hat{\mu}_1(X_i) - \hat{\mu}_0(X_i) + \frac{W_i}{\hat{e}(X_i)} (Y_i - \hat{\mu}_1(X_i)) - \frac{(1 - W_i)}{1 - \hat{e}(X_i)} (Y_i - \hat{\mu}_0(X_i)) \right) \quad (13)$$

$$\hat{\tau}_{AIPW} = \frac{1}{n} \sum_{i=1}^n \hat{\Gamma}_i \quad (14)$$

with $\hat{\Gamma}_i$ denoting each patient's individual double-robust score.

In causal forest applications, such as GRF by Wager and Athey (2018) and Athey et al. (2019), $\hat{\mu}_w(x)$ and $\hat{e}(X_i)$ are estimated from data using random forest. For our research purpose, a key advantage of this method is that it has proven efficient in high-dimensional settings, rendering pre-selection of confounders unnecessary (Wager and Athey, 2018). This is especially important for our research endeavor as our dataset includes thousands of detailed features per ICU stay (cf. Table 1 above). Discarding parts of these data based on expert or domain knowledge would inhibit the risk that the “wrong” features are excluded from analysis, while a data-driven feature selection decreases computational efficiency and increases complexity.

Heterogeneous treatment effects

To analyze treatment effect heterogeneity, commonly, one could try to find strata that reduce variation in ATE estimation, i.e., that help to better predict a potential outcome for a sub-group (Imbens and Rubin, 2015). Stratification has limitations, however. In our context, the two most important limitations

are (1) that we could not possibly know which of the thousands of features in our dataset might be the best or even among the best strata, (2) that stratification's data use is rather inefficient, especially if one needs to consider interaction terms, and – most importantly for our study – (3) that two discharge candidates might belong to the same stratified group, rendering a decision whom to discharge of the two, based on the discharge's effect on readmission, impossible.

With their causal forests, Wager and Athey (2018) and Lechner and Mareckova (2024) have provided causal machine learning methods not only for ATE estimation with observational data under the selection on observables assumptions, but also a flexible and efficient approach to estimate IATEs. IATEs are formally defined by:

$$\tau(x) = E[Y_i(1) - Y_i(0) | X_i = x, Z_i = z] \quad (15)$$

where X_i describes the confounders, and Z_i comprises the features defining the patients or group of patients an average effect is estimated for. We will estimate the IATEs with the same strategy and approach as the ATE described above.


Treatment definition

Each patient admitted to the ICU must be discharged at some point. Thus, a patient's treatment status at t might be different than at $t + 1$, $t + 2$, etc. To stay in a binary treatment framework, we follow a four-step approach.

Firstly, we define exclusion criteria which indicate that a patient cannot be discharged for medical reasons. These criteria are invasive ventilation and medication with a catecholamine as substance (e.g., noradrenalin) (Heidegger et al., 2005). For each patient in our dataset, we identify the point in time when the last invasive ventilation has stopped and the last dosage of catecholamines has been given. We add four hours to these identified points in time to approximate the typical time these patients could (theoretically) be discharged. If ventilation has stopped after the last dosage of catecholamines, the ventilation stop plus four hours is marked as the soonest possible discharge and vice versa if the last dosage of catecholamines was later than the ventilation stop.

Secondly, we count how many times a patient passed the regular time when senior physicians decide for a discharge, approximated by 07:00 a.m., after the possible discharge point in time, denoted by n . In our dataset, $n = 1$ for 54% of discharges, $n = 2$ for 20%, and approximately 10% of discharges occur after passing the regular discharge decision time more than five times (cf. Figure 6).

Figure 6: Share of discharges per n



Annotations: n is defined as the number of times a patient passes the regular discharge decision time after becoming a potential candidate for discharge. The plot is based on Sample B (12,950 ICU stays and 11,873 unique cases), limited to all observations with $n \leq 20$ ($n > 20$ is true for 106 ICU stays in our dataset).

Thirdly, we define a random theoretic discharge point in time d of patient i with $d_i = \{\max(n_i - 1, 0), n_i\}$ with equal probability for i to be discharged at $d_i = \max(n_i - 1, 0)$ and $d_i = n_i$.

Lastly, we assign a positive treatment status $W_i = 1$ to all i for which $d_i = n_i$ and a negative treatment status $W_i = 0$ to all i for which $d_i = \max(n_i - 1, 0)$.


For our current results, we only run the last two steps once per ICU stay. For our final results, we plan to run these two steps multiple times for separate (I)ATE estimations. Results from each run will be averaged in the end to receive the final results.

Outcome definition

There is an ongoing discussion in the medical literature regarding a meaningful measurement and use of ICU readmission rates (Hosein et al., 2014; Kramer et al., 2013; Woldhek et al., 2017). In our main model, we use readmission regardless of the time between discharge and readmission as outcome. We plan to perform one to two sensitivity analyses, using more narrowly defined readmission rates such as readmission within 48 hours and 96 hours (two to four days).

In our dataset, we observe a raw readmission rate of roughly 9.0% while readmission rates with defined times between discharge and readmission are between 2.8% and 6.1% (see Figure 7).

Figure 7: Readmission rates according to time between discharge and readmission



Annotations: The plot is based on Sample B (12,950 ICU stays and 11,873 unique cases).

Identifying assumptions

Identifying assumptions for causal inference are (Angrist et al., 1996; Imbens, 2000; Lechner, 2001): (1) Unconfoundedness, (2) Common Support (CS) or overlap, (3) Stable-Unit-Treatment-Value Assumption (SUTVA), and (4) exogeneity.

With unconfoundedness, we assume that we observe all variables that might influence both the treatment selection (i.e., to be or not to be discharged at point t) and the potential outcome of a discharge. Discharge decisions are made by senior physicians (Nates et al., 2016). We observe all data and decision variables that are available to these physicians when they make discharge decisions. This includes most variables that were judged by most Swiss ICUs as relevant for ICU discharge decisions (Heidegger et al., 2005). Thus, we argue for unconfoundedness in our study setting. There are three lines of counter-argumentation, however.

The first argument is that situations may occur in which a discharge decision (i.e., our treatment assignment) and thus discharge are determinate. Unconfoundedness would be violated if such a treatment determination was dependent on patient characteristics (X_i). (i) The decision must be positive every time when there is only one patient candidate for discharge at t . We may still assume unconfoundedness, as the number of potential discharge candidates is exogenous, i.e., we can view treatment assignment in these situations still as random. (ii) A discharge decision must be negative if there is a contraindication, e.g., a patient is intubated or a patient receives a certain drug or substance (e.g., catecholamines) (Heidegger et al., 2005). This does not violate unconfoundedness in our context, however, as we always observe the discharge at a point in time when it is not contraindicated (cf. treatment definition).

The second argument is that in practice, physicians actually consider more variables than “only” the thousands of clinical parameters available to them (and to us): Physicians might additionally collect “soft” data, e.g., during daily patient visiting rounds by visual control of the patient and discussions with nurses (Nates et al., 2016; Ofoma et al., 2018). Data obtained this way might include degree of paleness, sweating, communicated pain, mental confusion and disorientation, or therapy compliance.

Soft data is not available to us (or readable by any machine). Unconfoundedness would be violated if these soft data were not correlated with the thousands of clinical parameters available to us. This is rather unlikely.

The third argument is that a discharge, and possibly also a later readmission, are influenced by the available downstream skill mix and/ or capacity. Physicians might decide to discharge a certain patient on a given day if they know that there are experienced, well-qualified physicians and nurses available in downstream units who can manage the patient, even if this patient is sicker given clinical parameters than another patient who was not discharged another day when downstream skill mix and capacity were (allegedly) inadequate. In such situations, physicians might decide against a discharge as they expect a worsening of the patient's health status downstream, rendering an ICU readmission more likely.

To address such doubts that might still remain after controlling for several thousand patient characteristics, we could implement an instrumental variable (IV) approach. An IV approach is available for binary instruments in GRF (Athey et al., 2019) and was developed by Wang et al. (2021) for continuous IVs in causal forests. As IV, we have several options: (1) the daily number of admitted patients as a degree of ICU busyness both as continuous variable (only useable if we can incorporate a continuous IV in GRF) and as dummy (1=busy, 0=not busy; threshold to be determined, e.g., one third of overall ICU capacity), as this has shown to increase the number of ICU discharges (Nates et al., 2016), (2) the number of potential discharge candidates throughout ICU day shifts (e.g., 06:00 a.m. to 08:00 p.m.), or (3) proxies for downstream unit capacity (and skill-mix), e.g., weekday vs. weekend, or days until weekend. Regarding (3), actual utilization data would be preferable, yet such data is not collected on a daily level in a digital manner or highly unreliable if collected manually. Note that all three IV options are exogenous of treatment assignment.

For future studies, collecting data on daily downstream skill mix and capacity might be perceivable. This comes with considerable effort, however, as such data is not available in a structured manner for all IMCUs and NCUs of a hospital. While staff schedules could be available (often only available in

analog form or in e-mails, manually administered Microsoft Excel tables, or similar), these alone will not explain the skill mix perceived by the ICU decision maker. To accurately account for this perceived skill mix, we would have to label all physician and nurse staff members of all downstream units according to the experience and qualification of all ICU decision makers. This might result in different labels for individual downstream physicians and/ or nurses depending on what ICU decision maker is asked to label. Indeed, if experience and skill level are perceived differently by physicians, this might in turn supply an argument that unconfoundedness does hold, as downstream skill mix would then not systematically influence treatment assignment and outcome in the same (perceived) way. Lastly, another argument that unconfoundedness still holds is that it is unlikely that a physician will be able to judge the experience and skill level of all treating nurses and physicians active in downstream units, or even that an ICU physician always has full transparency of downstream unit staff schedules. Granted, if a discharge decision for a particularly complex patient were made, a physician could take the time to get transparency over staff schedules and also experience and skill level. A systematic influence still is unlikely, however.

Another possibility could be to measure experience by the number of years a medical professional has been actively working, and qualification by academic degrees, further education certificates, and scientific publications. This approach would also pose a major challenge in terms of required effort, and at least some of these data will not be available in a structured form. Lastly, measuring experience and skill level in this way also has limitations. For instance, one might argue that experience is in fact built by being exposed to adverse events and difficult situations. This should correlate with the number of years a professional has been working but this must not necessarily be true.

In summary, we believe we can viably argue that unconfoundedness holds (1) as we control for all confounders also available to senior physicians making discharge decisions, (2) patient characteristics do not systematically influence treatment assignment and outcome and if they do, serve as exclusion criteria (e.g., mechanical ventilation), and (3) while downstream capacity and skill-mix might influence ICU decision makers, it is not perceivable that they have transparency over both for all downstream


units for every discharge decision, or even the majority of decisions, thus impeding systematic influence. Still, we to acknowledge any concerns possibly left and the observational setting of our study, we will implement an IV approach as sensitivity analysis in our final manuscript.

We may assume CS, if we can show that propensity scores “overlap”, i.e., that each patient could be observed with or without a discharge at point t :

$$0 < p(W_i = 1|X_i = x) < 1 \quad \forall x \in X \quad (16)$$

In Figure 8, we plot the distribution of propensity scores for our sample (cf. Wager and Athey, 2018).

Figure 8: Common support and overlap analysis of propensity scores



Annotations: Propensities were estimated with GRF, and for Sample B (12,950 ICU stays and 11,873 unique cases).

The plot shows that for our sample and treatment definition, there is a medium to strong selectivity into and out of treatment, respectively. Lechner and Mareckova (2024) show that their Modified Causal Forest (MCF), delivers more robust estimates than GRF in cases of (medium and) strong selectivity, especially when estimating IATEs. Thus, we estimate ATEs with both algorithms, one example of CATEs for exemplification with GRF, and IATEs with MCF. For our final results, we will estimate all results with both MCF and GRF.

The SUTVA requires that spillover effects between discharged patients are absent. More concretely, the discharge of one patient can only affect the outcome of the same patient and not the outcome of another patient. This is given as discharging one patient does not directly influence the readmission risk of another patient. One might argue that once the capacity constraint $B_a \geq 0$ holds, no additional patient(s) is discharged any more, even if this were medically possible, and thus the treatment state of the discharged patient(s) to satisfy $B_a \geq 0$ might influence the treatment state of some of the patients remaining in the ICU. Still, the treatment state of the discharged patient(s) will not directly influence the outcome of the patient(s) remaining in the ICU.

Exogeneity stipulates that patient characteristics used as confounders (X_i) are not influenced by the discharge at point in time t (the treatment). This assumption holds as we observe X_i before the discharge occurs.

In summary, we should be able to fulfill all four identifying assumptions or at least be able to implement an empirical strategy and methods to handle potential violations, ensuring robust ATE, CATE, and IATE estimations.

First results

In Table 3, we present the ATE estimation, employing both MCF and GRF.

Table 3: Estimated Average Treatment Effects

Algorithm	Estimate	Standard Error	Absolute change of readmission risk (in %-pts.)	Relative change of readmission risk (in %)
MCF	0.00767	0.0083903	0.77	8%
GRF	-0.00179	0.0047603	-0.18	-2%

Annotations: Estimates were made with Sample B (12,950 ICU stays and 11,873 unique cases), >4,600 included features (cf. Table 1). The relative change of readmission risk is calculated in comparison to the raw readmission rate of 9.04%.

The ATE of discharging an ICU patient to a downstream unit at t as compared to one decision cycle later increases the readmission risk by 8% according to MCF but decreases the readmission risk by -2% according to GRF. Evidently, both point estimates are insignificant, while the standard error for GRF is much larger as compared to the estimate than for MCF. Besides, the magnitude of the effect, regardless of significance and sign of the coefficient, is more than four times higher for MCF than for GRF. Still, simply put, according to both MCF and GRF, on average, there is no effect. It is important to note,

however, that we are not interested in the ATE as it does not enable discriminative discharge decisions between two or more individual ICU discharge candidates (cf. below). In fact, it is sensible that on average, discharging a dischargeable, comparatively stable patient one day (or rather decision cycle) later will not affect the readmission risk significantly.

Table 4 shows the estimated CATEs for four groups, stratified by the number of times a patient passes the regular discharge decision time after becoming a potential candidate for discharge, denoted by n . The CATEs show that there is ample treatment heterogeneity. While there is no effect for the group $n = 2$, the magnitude of the estimate is quite large for all other n and standard errors are smaller for $n = 0$ and $n = 1$, but effects are still insignificant.

Table 4: Estimated Conditional Average Treatment Effects according to n

n	Estimate	Standard Error	Absolute change of readmission risk (in %-pts.)	Relative change of readmission risk (in %)
0	-0.01504	0.006106	-1.50	-17%
1	0.01168	0.007846	1.17	13%
2	-0.00097	0.018317	-0.10	-1%
3	0.03302	0.036234	3.30	37%

Annotations: Estimates were made with GRF and with Sample B (12,950 ICU stays and 11,873 unique cases), >4,600 included features (cf. Table 1). The relative change of readmission risk is calculated in comparison to the raw readmission rate of 9.04%. n denotes the number of times a patient passes the regular discharge decision time after becoming a potential candidate for discharge.


Note that neither the ATE, nor any kind of CATEs will be of use for decision support. For the ATE, this is obvious as all cases have the same ATE and it thus cannot suggest what patient should rather be discharged. CATEs encompass still too many patients, so that on a given day, two or more discharge candidates might fall into the same group with the same CATE. Then, the same issue as with using an ATE applies.

Therefore, IATEs need to be estimated, forming the basis for developing decision policies that can truly support daily, operational ICU decision making. Given the results in Table 3 and our common support analysis in Figure 8, we chose MCF as causal machine learning method to estimate IATEs. Figure 9 shows the density of estimated IATEs.

Almost the entire sample (97.8%, 12,669 ICU stays) shows IATE estimates between -0.05 and 0.10. Note that this variation is actually quite large: In percentage points, readmission risk is influenced between

negative 5 and 10 points, amounting to a relative change of readmission risk between -55% and 111% when compared to the average readmission rate of 9.04%. Considering a narrower IATE range between -0.05 and 0.05 (91.4% of sample, 11,833 ICU stays), the relative change of readmission risk still amounts to -55% to 55%. This range shows the existing treatment heterogeneity and showcases that these results could prove effective when developing discharge policies based on IATEs.

Figure 9: Density of Individualized Average Treatment Effects



Annotations: IATEs were estimated with MCF, and for Sample B (12,950 ICU stays and 11,873 unique cases).

Still, the figure also shows that the IATE estimate for large share of ICU stays is between 0.00 and 0.03 (58.5%, 11,833 ICU stays). This means that for the majority of patients, a discharge one decision cycle later increases their readmission risk relatively mildly, between 0% and 33% compared to the readmission rate of 9.04%. Lastly, CATEs in Table 4 already indicate that there are certain patients that might not benefit from additional time on the ICU. Our results for the IATE estimation show that IATEs are indeed negative for roughly 28% of all patients.

Still, when interpreting IATEs, we must be cautious insofar that we do not show estimated standard errors of patients' IATEs which would be needed to generalize statements about the effectiveness of additional intensive care for individual ICU patients. Accordingly, the above statements have a descriptive and not necessarily causal intention.

For our research objective, i.e., supporting clinical decision making, standard errors are of a lesser importance, which we discuss in the next section.

Development of Decision Policies and Simple Simulation Study

We aim to design policies offering decision support for all ICU discharge types satisfying the capacity constraint $B_a \geq 0$, with B_a as defined in equation (4):

$$B_a = B - \left(\sum_{i=1}^I x_i + \sum_{e=1}^E a_e - \sum_{c=1}^C a_c^{elective} + \sum_{r=1}^R a_r + \sum_{u=1}^U a_u^{ant} \right) + \left(\sum_{n=1}^N d_n + \sum_{p=1}^P d_p^{early} \right) + \left(\sum_{q=1}^Q d_q^{early} - \sum_{v=1}^V a_v^{add} \right)$$

Specifically, policies should offer support when planning discharges $\sum_{n=1}^N d_n + \sum_{p=1}^P d_p^{early}$, and when having to decide in unplanned demand-driven situations described by $\sum_{q=1}^Q d_q^{early} - \sum_{v=1}^V a_v^{add}$. Note that in practice, both d_n and d_p^{early} are planned discharges and distinguishing clearly between what discharge is natural and what discharge is rather demand-driven due to unplanned but anticipated admissions a_u^{ant} , anticipated readmissions a_r , or planned admissions due to elective surgeries a_e , might be difficult.

To support decision making, in a first step, we will rank all patients according to their IATE estimate. In a second step, we determine that patients should be selected for discharge based on their spot in this ranking as compared to other patients who are discharge candidates at the same regular decision cycle. More concretely, if the ICU decision maker were to decide between two patients, the patient with the lower change of readmission risk due to the discharge at point in time t , i.e., the smaller IATE, would be selected. The same applies if multiple patients were selected for discharge. For instance, if there are

seven discharge candidates and to comply with the capacity constraint $B_a \geq 0$, five patients would need to be discharged, the five patients with the lowest IATEs would be selected for discharge.

Statistical significance, e.g., at the 5%-level of estimated IATEs would indicate the robustness of the magnitude of individual effects, i.e., how certain we can be about individual patients' point estimate. We argue that statistical significance is of lesser importance for selecting discharge candidates as even insignificant IATEs can optimize decisions, at least as long as confidence intervals do not overlap. For the simulation and final version of our manuscript, we will address this potential challenge in more detail.

To show the practical utility of our decision policy approach in terms of avoided readmissions and saved ICU capacity, we will start a simulation at the first day of our observation period. We then apply our decision policy and change the discharge decision for all patients where the policy would discharge a different patient than the patient that was actually discharged. We continue to apply this decision policy across all days of our observation period. To receive the total number of avoidable readmissions, we sum the difference between the minimal (i.e., optimal) IATE, $IATE_{d,p}$, and empirical (i.e., actual) IATE, $IATE_{d,a}$, across all discharge decisions D . Note that this difference is zero where the optimal and actual discharge decisions are the same.

$$R = \sum_{d=1}^D IATE_{d,p} - IATE_{d,a} \quad (17)$$

Additionally, we can compute the number of saved ICU bed capacities B according to our decision policy with

$$B = \frac{R \times \overline{LOS}}{365 \times \text{bed occupancy rate}} \quad (18)$$

where $\overline{LOS}$ represents the average length of stay of readmitted cases. We will assume different bed occupancy rates typical for ICUs, e.g., 90% and 95%, to receive a range for our avoidable ICU bed capacity estimate.

Note that this is a simple simulation. A more sophisticated approach would be to account for changes in future ICU demand due to changes in readmissions. More concretely, optimizing ICU discharges in t would influence the number of ICU arrivals in $t + 1$, as some patients are less likely to be readmitted, i.e., $\sum_{r=1}^R a_r$ will be lower. A change in the number of ICU arrivals then will lead to different (expected) fulfilment of the capacity constraint $B_a \geq 0$, potentially requiring more or less changes discharges as empirically observed. We refrain from conducting such a simulation, however, since it would be prone to many of the limitations of Markov Decision Process Models, as described above.

4 Summary and Outlook

We outline how optimal decisions can be learned from observational data under the unconfoundedness assumption, leveraging more than 4,600 features to estimate IATEs of ICU discharge at point in time t on ICU readmission. Our model addresses a classic hospital operations problem: Decision-making under uncertainty, coupled with fixed capacities, leads to sub-par medical outcomes and bottleneck congestion.

We reviewed the OM literature with regards to machine learning, causal inference, hospital operations, and policy learning, yielding that causal machine learning has rarely been employed in OM, especially in the context of hospital operations. In this working paper, we formulated the decision problem, given by the bed capacity constraint $B_a > 0$, of the ICU discharge process. Utilizing two causal machine learning algorithms, we estimated ATEs (GRF and MCF), CATEs (only GRF) and IATEs (only MCF). We described how ranked IATEs can be used in a simple simulation to show the practical utility of our methodological approach. For the final manuscript of this study, we will implement this simulation and discuss our results in the context of the operations research and medical literature. Furthermore, we will explore different IV approaches as sensitivity analyses acknowledging the (unlikely) situations where unconfoundedness might not hold. In addition, we will refine our feature engineering, e.g., by exploring sub-dimensions of clinical scores. Lastly, as additional scenario, we will define the outcome as readmission within 48 or 96 hours as compared to “raw” readmissions.

Bibliography

- Al-Jaghbeer, M.J., Tekwani, S.S., Gunn, S.R., Kahn, J.M., 2016. Incidence and Etiology of Potentially Preventable ICU Readmissions. *Crit Care Med* 44, 1704–1709. <https://doi.org/10.1097/CCM.0000000000001746>
- Ang, E., Kwasnick, S., Bayati, M., Plambeck, E.L., Aratow, M., 2015. Accurate Emergency Department Wait Time Prediction. *Manufacturing & Service Operations Management* 18, 141–156. <https://doi.org/10.1287/MSOM.2015.0560>
- Angrist, J.D., Imbens, G.W., Rubin, D.B., 1996. Identification of Causal Effects Using Instrumental Variables. *J Am Stat Assoc* 91, 444–455. <https://doi.org/10.1080/01621459.1996.10476902>
- Athey, S., 2017. Beyond prediction: Using big data for policy problems. *Science* (1979) 355, 483–485. <https://doi.org/10.1126/SCIENCE.AAL4321>
- Athey, S., Imbens, G., 2016. Recursive partitioning for heterogeneous causal effects. *Proc Natl Acad Sci U S A* 113, 7353–7360. <https://doi.org/10.1073/pnas.1510489113>
- Athey, S., Imbens, G.W., 2019. Machine Learning Methods That Economists Should Know about. *Annu Rev Econom* 11, 685–725. <https://doi.org/10.1146/ANNUREV-ECONOMICS-080217-053433/CITE/REFWORKS>
- Athey, S., Imbens, G.W., 2017. The State of Applied Econometrics: Causality and Policy Evaluation. *The Journal of Economic Perspectives* 32, 3–32.
- Athey, S., Tibshirani, J., Wager, S., 2019. Generalized random forests. *Ann Stat* 47, 1148–1178. <https://doi.org/10.1214/18-AOS1709>
- Athey, S., Wager, S., 2021. Policy Learning With Observational Data. *Econometrica* 89, 133–161. <https://doi.org/10.3982/ECTA15732>
- Bai, J., Fügener, A., Gönsch, J., Brunner, J.O., Blobner, M., 2021. Managing admission and discharge processes in intensive care units. *Health Care Manag Sci* 24, 666–685. <https://doi.org/10.1007/S10729-021-09560-6>
- Bai, J., Fügener, A., Schoenfelder, J., Brunner, J.O., 2018. Operations research in intensive care unit management: a literature review. *Health Care Manag Sci* 21, 1–24. <https://doi.org/10.1007/s10729-016-9375-1>
- Barbieri, S., Kemp, J., Perez-Concha, O., Kotwal, S., Gallagher, M., Ritchie, A., Jorm, L., 2020. Benchmarking Deep Learning Architectures for Predicting Readmission to the ICU and Describing Patients-at-Risk. *Scientific Reports* 2020 10:1 10, 1–10. <https://doi.org/10.1038/s41598-020-58053-z>
- Berk, E., Moinzadeh, K., 1998. The Impact of Discharge Decisions on Health Care Quality. *Manage Sci* 44, 400–415. <https://doi.org/10.1287/MNSC.44.3.400>

- Bertsekas, D.P., 2012. *Dynamic Programming and Optimal Control*, 4th ed. Athena Scientific, Belmont, Massachusetts.
- Bertsimas, D., Pauphilet, J., 2023. Hospital-Wide Inpatient Flow Optimization. *Manage Sci* 70, 4893–4911. <https://doi.org/10.1287/MNSC.2023.4933>
- Bravo, F., Rudin, C., Shaposhnik, Y., Yuan, Y., 2024. Interpretable Prediction Rules for Congestion Risk in Intensive Care Units. *Stochastic Systems* 14, 111–130. <https://doi.org/https://doi.org/10.1287/stsy.2022.0018>
- Breiman, L., 2001. Random forests. *Mach Learn* 45, 5–32. <https://doi.org/10.1023/A:1010933404324>
- Browning, T.R., 2020. Operations management writ large. *Journal of Operations Management* 66, 494–500. <https://doi.org/10.1002/JOOM.1094>
- Chalfin, D.B., Trzeciak, S., Likourezos, A., Baumann, B.M., Dellinger, R.P., 2007. Impact of delayed transfer of critically ill patients from the emergency department to the intensive care unit. *Crit Care Med* 35, 1477–1483. <https://doi.org/10.1097/01.CCM.0000266585.74905.5A>
- Chan, C.W., Farias, V.F., Bambos, N., Escobar, G.J., 2012. Optimizing Intensive Care Unit Discharge Decisions with Patient Readmissions. *Oper Res* 60, 1323–1341. <https://doi.org/10.1287/OPRE.1120.1105>
- Chan, C.W., Green, L. V., Lekwijit, S., Lu, L., Escobar, G., 2018. Assessing the Impact of Service Level When Customer Needs Are Uncertain: An Empirical Investigation of Hospital Step-Down Units. *Manage Sci* 65, 751–775. <https://doi.org/10.1287/MNSC.2017.2974>
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., Robins, J., 2018. Double/debiased machine learning for treatment and structural parameters. *Econom J* 21, C1–C68. <https://doi.org/10.1111/ectj.12097>
- Chou, Y.C., Chuang, H.H.C., Chou, P., Oliva, R., 2023. Supervised machine learning for theory building and testing: Opportunities in operations management. *Journal of Operations Management* 69, 643–675. <https://doi.org/10.1002/JOOM.1228>
- Cordier, J., Salvi, I., Steinbeck, V., Geissler, A., Vogel, J., 2023. Is rapid recovery always the best recovery? - Developing a machine learning approach for optimal assignment rules under capacity constraints for knee replacement patients (No. WP 23/08), Health, Econometrics and Data Group (HEDG) Working Papers. York.
- Cui, Y., Davis, A.M., 2022. Tax-Induced Inequalities in the Sharing Economy. *Manage Sci* 68, 7202–7220. <https://doi.org/10.1287/MNSC.2021.4277>
- Delgado, K.V., Sanner, S., De Barros, L.N., 2011. Efficient solutions to factored MDPs with imprecise transition probabilities. *Artif Intell* 175, 1498–1527. <https://doi.org/10.1016/J.ARTINT.2011.01.001>

- Demir, E., 2014. A Decision Support Tool for Predicting Patients at Risk of Readmission: A Comparison of Classification Trees, Logistic Regression, Generalized Additive Models, and Multivariate Adaptive Regression Splines. *Decision Sciences* 45, 849–880. <https://doi.org/10.1111/DECI.12094>
- Dobson, G., Lee, H.H., Pinker, E., 2010. A Model of ICU Bumping. *Oper Res* 58, 1564–1576. <https://doi.org/10.1287/OPRE.1100.0861>
- Feuerriegel, S., Frauen, D., Melnychuk, V., Schweisthal, J., Hess, K., Curth, A., Bauer, S., Kilbertus, N., Kohane, I.S., van der Schaar, M., 2024. Causal machine learning for predicting treatment outcomes. *Nat Med* 30, 958–968. <https://doi.org/10.1038/s41591-024-02902-1>
- Gopalan, P.D., Pershad, S., 2019. Decision-making in ICU – A systematic review of factors considered important by ICU clinician decision makers with regard to ICU triage decisions. *J Crit Care* 50, 99–110. <https://doi.org/10.1016/J.JCRC.2018.11.027>
- Grand-Clément, J., Chan, C.W., Goyal, V., Escobar, G., 2023. Robustness of Proactive Intensive Care Unit Transfer Policies. *Oper Res* 71, 1653–1688. <https://doi.org/10.1287/opre.2022.2403>
- Green, L. V., 2012. OM Forum—The Vital Role of Operations Analysis in Improving Healthcare Delivery. *Manufacturing & Service Operations Management* 14, 488–494. <https://doi.org/10.1287/MSOM.1120.0397>
- Hammer, M., Grabitz, S.D., Teja, B., Wongtangman, K., Serrano, M., Neves, S., Siddiqui, S., Xu, X., Eikermann, M., 2020. A Tool to Predict Readmission to the Intensive Care Unit in Surgical Critical Care Patients—The RISC Score. *J Intensive Care Med* 36, 1296–1304. <https://doi.org/10.1177/0885066620949164>
- Heidegger, C.P., Treggiari, M.M., Romand, J.A., 2005. A nationwide survey of intensive care unit discharge practices. *Intensive Care Med* 31, 1676–1682. <https://doi.org/10.1007/S00134-005-2831-X/TABLES/5>
- Ho, T.-H., Lim, N., Reza, S., Xia, X., 2017. OM Forum-Causal Inference Models in Operations Management. *Manufacturing & Service Operations Management* 19, 509–525. <https://doi.org/10.1287/msom.2017.0659>
- Hosein, F.S., Roberts, D.J., Turin, T.C., Zygun, D., Ghali, W.A., Stelfox, H.T., 2014. A meta-analysis to derive literature-based benchmarks for readmission and hospital mortality after patient discharge from intensive care. *Crit Care* 18, 1–12. <https://doi.org/10.1186/S13054-014-0715-6/TABLES/2>
- Hydari, M.Z., Telang, R., Marella, W.M., 2018. Saving Patient Ryan—Can Advanced Electronic Medical Records Make Patient Care Safer? *Manage Sci* 65, 2041–2059. <https://doi.org/10.1287/MNSC.2018.3042>

- Imbens, G.W., 2000. The role of the propensity score in estimating dose-response functions. *Biometrika* 87, 706–710. <https://doi.org/10.1093/BIOMET/87.3.706>
- Imbens, G.W., Rubin, D.B., 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge University Press. <https://doi.org/10.1017/CBO9781139025751>
- KC, D.S., Terwiesch, C., 2011. An Econometric Analysis of Patient Flows in the Cardiac Intensive Care Unit. *Manufacturing & Service Operations Management* 14, 50–65. <https://doi.org/10.1287/MSOM.1110.0341>
- Keskinocak, P., Savva, N., 2019. A Review of the Healthcare-Management (Modeling) Literature Published in *Manufacturing & Service Operations Management*. *Manufacturing & Service Operations Management* 22, 59–72. <https://doi.org/10.1287/MSOM.2019.0817>
- Kim, S.H., Chan, C.W., Olivares, M., Escobar, G., 2014. ICU Admission Control: An Empirical Study of Capacity Allocation and Its Implication for Patient Outcomes. *Manage Sci* 61, 19–38. <https://doi.org/10.1287/MNSC.2014.2057>
- Kim, S.H., Zheng, F., Brown, J., 2023. Identifying the Bottleneck Unit: Impact of Congestion Spillover in Hospital Inpatient Unit Network. *Manage Sci* 70, 4200–4218. <https://doi.org/10.1287/MNSC.2023.4887>
- Ko, D.G., Mai, F., Shan, Z., Zhang, D., 2019. Operational efficiency and patient-centered health care: A view from online physician reviews. *Journal of Operations Management* 65, 353–379. <https://doi.org/10.1002/JOOM.1028>
- Kramer, A.A., Higgins, T.L., Zimmerman, J.E., 2013. The association between ICU readmission rate and patient outcomes. *Crit Care Med* 41, 24–33. <https://doi.org/10.1097/CCM.0B013E3182657B8A>
- Kramer, A.A., Higgins, T.L., Zimmerman, J.E., 2012. Intensive care unit readmissions in U.S. hospitals: Patient characteristics, risk factors, and outcomes. *Crit Care Med* 40, 3–10. <https://doi.org/10.1097/CCM.0B013E31822D751E>
- Lechner, M., 2001. Identification and estimation of causal effects of multiple treatments under the conditional independence assumption, in: Lechner, M., Pfeiffer, F. (Eds.), *Econometric Evaluation of Labour Market Policies*. Physica, Heidelberg, pp. 43–58. https://doi.org/10.1007/978-3-642-57615-7_3
- Lechner, M., Mareckova, J., 2024. *Comprehensive Causal Machine Learning*. <https://doi.org/10.48550/arXiv.2405.10198>
- Litvak, N., van Rijsbergen, M., Boucherie, R.J., van Houdenhoven, M., 2008. Managing the overflow of intensive care patients. *Eur J Oper Res* 185, 998–1010. <https://doi.org/10.1016/J.EJOR.2006.08.021>

- Loreto, M., Lisboa, T., Moreira, V.P., 2020. Early prediction of ICU readmissions using classification algorithms. *Comput Biol Med* 118, 103636. <https://doi.org/10.1016/J.COMPBIOMED.2020.103636>
- Lu, G., Ding, X. (David), Peng, D.X., Hao-Chun Chuang, H., 2018. Addressing endogeneity in operations management research: Recent developments, common problems, and directions for future research. *Journal of Operations Management* 64, 53–64. <https://doi.org/10.1016/J.JOM.2018.10.001>
- Mastin, A., Jaillet, P., 2012. Loss bounds for uncertain transition probabilities in Markov decision processes. *Proceedings of the IEEE Conference on Decision and Control* 6708–6715. <https://doi.org/10.1109/CDC.2012.6426504>
- McConnell, K.J., Richards, C.F., Daya, M., Bernell, S.L., Weathers, C.C., Lowe, R.A., 2005. Effect of increased ICU capacity on emergency department length of stay and ambulance diversion. *Ann Emerg Med* 45, 471–478. <https://doi.org/10.1016/j.annemergmed.2004.10.032>
- Mcneill, H., Khairat, S., 2020. Impact of Intensive Care Unit Readmissions on Patient Outcomes and the Evaluation of the National Early Warning Score to Prevent Readmissions: Literature Review. *JMIR Perioper Med* 3, e13782. <https://doi.org/10.2196/13782>
- Milbrandt, E.B., Kersten, A., Rahim, M.T., Dremsizov, T.T., Clermont, G., Cooper, L.M., Angus, D.C., Linde-Zwirble, W.T., 2008. Growth of intensive care unit resource use and its estimated cost in Medicare. *Crit Care Med* 36, 2504–2510. <https://doi.org/10.1097/CCM.0B013E318183EF84>
- Mukhopadhyay, S., Samaddar, S., Solis, A.O., Roy, A., 2021. Disease Detection Analytics: A Simple Linear Convex Programming Algorithm for Breast Cancer and Diabetes Incidence Decisions. *Decision Sciences* 52, 661–698. <https://doi.org/10.1111/DECI.12348>
- Munro, E., 2024. Treatment Allocation with Strategic Agents. *Manage Sci.* <https://doi.org/10.1287/MNSC.2022.01629>
- Nates, J.L., Nunnally, M., Kleinpell, R., Blosser, S., Goldner, J., Birriel, B., Fowler, C.S., Byrum, D., Miles, W.S., Bailey, H., Sprung, C.L., 2016. ICU Admission, Discharge, and Triage Guidelines: A Framework to Enhance Clinical Operations, Development of Institutional Policies, and Further Research. *Crit Care Med* 44, 1553–1602. <https://doi.org/10.1097/CCM.0000000000001856>
- Nenova, Z., Shang, J., 2022. Chronic Disease Progression Prediction: Leveraging Case-Based Reasoning and Big Data Analytics. *Prod Oper Manag* 31, 259–280. <https://doi.org/https://doi.org/10.1111/poms.13532>
- Niven, D.J., Bastos, J.F., Stelfox, H.T., 2014. Critical care transition programs and the risk of readmission or death after discharge from an ICU: A systematic review and meta-analysis. *Crit Care Med* 42, 179–187. <https://doi.org/10.1097/CCM.0B013E3182A272C0>

- Ofoma, U.R., Dong, Y., Gajic, O., Pickering, B.W., 2018. A qualitative exploration of the discharge process and factors predisposing to readmissions to the intensive care unit. *BMC Health Serv Res* 18, 6. <https://doi.org/10.1186/s12913-017-2821-z>
- Ouyang, H., Argon, N.T., Ziya, S., 2020. Allocation of Intensive Care Unit Beds in Periods of High Demand. *Oper Res* 68, 591–608. <https://doi.org/10.1287/OPRE.2019.1876>
- Piri, H., Huh, W.T., Shechter, S.M., Hudson, D., 2022. Individualized Dynamic Patient Monitoring Under Alarm Fatigue. *Oper Res* 70, 2749–2766. <https://doi.org/10.1287/OPRE.2022.2300>
- Prosperi, M., Guo, Y., Sperrin, M., Koopman, J.S., Min, J.S., He, X., Rich, S., Wang, M., Buchan, I.E., Bian, J., 2020. Causal inference and counterfactual prediction in machine learning for actionable healthcare. *Nature Machine Intelligence* 2:7 2, 369–375. <https://doi.org/10.1038/s42256-020-0197-y>
- Renaud, B., Santin, A., Coma, E., Camus, N., Van Pelt, D., Hayon, J., Gurgui, M., Roupie, E., Hervé, J., Fine, M.J., Brun-Buisson, C., Labarère, J., 2009. Association between timing of intensive care unit admission and outcomes for emergency department patients with community-acquired pneumonia. *Crit Care Med* 37, 2867–2874. <https://doi.org/10.1097/CCM.0B013E3181B02DBB>
- Robins, J.M., Rotnitzky, A., Zhao, L.P., 1994. Estimation of Regression Coefficients When Some Regressors are not Always Observed. *J Am Stat Assoc* 89, 846–866. <https://doi.org/10.1080/01621459.1994.10476818>
- Rosa, R.G., Falavigna, M., Robinson, C.C., Sanchez, E.C., Kochhann, R., Schneider, D., Sganzerla, D., Dietrich, C., Barbosa, M.G., De Souza, D., Rech, G.S., Dos Santos, R. da R., Da Silva, A.P., Santos, M.M., Dal Lago, P., Sharshar, T., Bozza, F.A., Teixeira, C., 2020. Early and Late Mortality Following Discharge From the ICU: A Multicenter Prospective Cohort Study*. *Crit Care Med* 48, 64–72. <https://doi.org/10.1097/CCM.0000000000004024>
- Rosenbaum, P.R., Rubin, D.B., 1983. The Central Role of the Propensity Score in Observational Studies for Causal Effects. *Biometrika* 70, 41–55. <https://doi.org/10.1093/biomet/70.1.41>
- Rubin, D.B., 1974. Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies. *J Educ Psychol* 66, 688–701.
- Satia, J.K., Lave, R.E., 1973. Markovian Decision Processes with Uncertain Transition Probabilities. *Oper Res* 21, 728–747. <https://doi.org/10.1287/OPRE.21.3.728>
- Scheulen, J.J., Li, G., Kelen, G.D., 2001. Impact of ambulance diversion policies in urban, suburban, and rural areas of Central Maryland. *Academic Emergency Medicine* 8, 36–40. <https://doi.org/10.1111/J.1553-2712.2001.TB00546.X>

- Schull, M.J., Vermeulen, M., Slaughter, G., Morrison, L., Daly, P., 2004. Emergency department crowding and thrombolysis delays in acute myocardial infarction. *Ann Emerg Med* 44, 577–585. <https://doi.org/10.1016/j.annemergmed.2004.05.004>
- Shi, P., Helm, J.E., Deglise-Hawkinson, J., Pan, J., 2021. Timing It Right: Balancing Inpatient Congestion vs. Readmission Risk at Discharge. *Oper Res* 69, 1842–1865. <https://doi.org/10.1287/OPRE.2020.2044>
- Thirumalai, S., Devaraj, S., Browning, T.R., 2024. Uncertainty in healthcare operations: How hospitals weather the perfect storm. *Journal of Operations Management* 1–19. <https://doi.org/10.1002/JOOM.1327>
- Wager, S., Athey, S., 2018. Estimation and Inference of Heterogeneous Treatment Effects using Random Forests. *J Am Stat Assoc* 113, 1228–1242. <https://doi.org/10.1080/01621459.2017.1319839>
- Wang, G., Li, J., Hopp, W.J., 2023. Personalized Healthcare Outcome Analysis of Cardiovascular Surgical Procedures. *Manufacturing & Service Operations Management* 25, 1567–1584. <https://doi.org/10.1287/MSOM.2023.1227>
- Wang, G., Li, J., Hopp, W.J., 2021. An Instrumental Variable Forest Approach for Detecting Heterogeneous Treatment Effects in Observational Studies. *Manage Sci* 68, 3399–3418. <https://doi.org/10.1287/MNSC.2021.4084>
- Woldhek, A.L., Rijkenberg, S., Bosman, R.J., van der Voort, P.H.J., 2017. Readmission of ICU patients: A quality indicator? *J Crit Care* 38, 328–334. <https://doi.org/10.1016/J.JCRC.2016.12.001>
- Wooldridge, J.M., 2010. *Econometric Analysis of Cross Section and Panel Data*, 2nd ed. The MIT Press.
- Yang, J., Eckles, D., Dhillon, P., Aral, S., 2023. Targeting for Long-Term Outcomes. *Manage Sci* 70, 3841–3855. <https://doi.org/10.1287/MNSC.2023.4881>
- Yilmaz, Ö., Son, Y., Shang, G., Arslan, H.A., 2024. Causal inference under selection on observables in operations management research: Matching methods and synthetic controls. *Journal of Operations Management* 70, 831–859. <https://doi.org/10.1002/JOOM.1318>
- Zhan, R., Ren, Z., Athey, S., Zhou, Z., 2024. Policy Learning with Adaptively Collected Data. *Manage Sci* 70, 5270–5297. <https://doi.org/10.1287/MNSC.2023.4921>
- Zhang, Y., Wu, H., Denton, B.T., Wilson, J.R., Lobo, J.M., 2019. Probabilistic sensitivity analysis on Markov models with uncertain transition probabilities: an application in evaluating treatment decisions for type 2 diabetes. *Health Care Manag Sci* 22, 34–52. <https://doi.org/10.1007/S10729-017-9420-8>