

Optimal Dynamic Treatment Regime Estimation Using Information Extraction From Unstructured Clinical Text

Nina Zhou^{1,2}, Robert D. Brook³, Ivo D. Dinov², and Lu Wang^{1*}

¹ Department of Biostatistics, University of Michigan, Ann Arbor, USA

² Statistics Online Computational Resource, University of Michigan, Ann Arbor, MI, USA

³ Division of Cardiovascular Diseases, School of Medicine, Wayne State University, Detroit, MI, USA

* Corresponding author's email: luwang@umich.edu

The wide-scale adoption of Electronic Health Records (EHRs) provides extensive information to support precision medicine and personalized health care. In addition to structured EHRs, we leverage free-text clinical information extraction (IE) techniques to estimate optimal dynamic treatment regimes (DTR), a sequence of decision rules that dictate how to individualize treatments to patients based on treatment and covariate history. The proposed IE of patient characteristics closely resembles “The clinical Text Analysis and Knowledge Extraction System” and employs named entity recognition, boundary detection, and negation annotation. It also utilizes regular expressions to extract numerical information. Combining the proposed IE with optimal DTR estimation, we extract derived patient characteristics and use Tree-based Reinforcement Learning (T-RL) to estimate multi-stage optimal DTRs. Information extraction significantly improved the estimation in counterfactual outcome models compared to using structured EHR data alone, which often includes incomplete data, data entry errors, and other potentially unobserved risk factors. Moreover, including IE in optimal DTR estimation provides larger study cohorts and a broader pool of candidate tailoring variables. We demonstrate the performance of our proposed method via simulations and an application using clinical records to guide blood pressure control treatments among critically ill patients with severe acute hypertension. This joint estimation approach improves the accuracy of identifying the optimal treatment sequence by 14-24% compared to traditional inference without using information extraction, based on our simulations over various scenarios. In the blood pressure control application, we successfully extracted significant blood pressure predictors that are unobserved or partially missing from structured EHR.

Key words: Causal Inference; Clinical Decision-Making; Electronic Health Record; Precision Medicine; Text Mining

Supporting Information for this article is available from the author or on the WWW under <http://dx.doi.org/10.1022/bimj.XXXXXXX> (please delete if not applicable)

1 Introduction

Personalized medicine provides individualized patient treatment recommendations by tailoring specific treatments to heterogeneous patients with various characteristics, unlike the one-size-fits-all model of care. Dynamic treatment regime represents an effective vehicle under the umbrella of personalized medicine that offers adaptive treatment strategies (Chakraborty and Murphy, 2014). Especially for chronic conditions, a course of medical intervention containing multiple treatment stages is often needed for patients. A driving motivational challenge in this study is the time-varying treatment decision to control patients blood pressure during the first two days in intensive care units (ICU). We consider four potential classes of antihypertensive agents: angiotensin-converting enzyme inhibitors (ACEI), beta-blockers, calcium channel blockers (CCB), and diuretics. The first treatment will be assigned to patients based on medical history and clinical evidence immediately after the ICU admission. If the treatment gets the elevated blood pressure controlled, no further action is needed on day two. Otherwise, a subsequent treatment is recommended to non-responders in order to better manage the patients blood pressure. In this example, we have two

sequential treatment decision stages, and one possible DTR example is to treat patients younger than 55 years of age with ACEI and treat the rest of patients with CCB on day one; then provide CCB to non-responders on day two.

To guide evidence-based effective treatment decisions, researchers have developed various statistical methods to evaluate and identify the optimal DTR, tailoring the optimal treatment choice to each individual to maximize their expected clinical outcome given current disease status and medical history. Abundant patient information is needed to obtain accurate estimations for desired clinical outcomes and provide various candidate tailoring variables while constructing treatment rules.

Many statistical methods have been developed for identifying optimal DTRs using observational data. Commonly used parametric and semi-parametric methods include marginal structural models with inverse probability weighting (Murphy *et al.*, 2001; Hernn, Brumback, and Robins, 2001; Wang *et al.*, 2012), G-estimation of structural nested mean models (Robins, 1994; Robins, 1997; Robins, 2004), and targeted maximum likelihood estimations (Laan and Rubin, 2006). These methods provide high interpretability but make specific modeling assumptions for a sequence of conditional models, which may be practically unattainable in certain cases. To alleviate strict modeling assumptions and maintain interpretability, Laber and Zhao proposed a tree-based method for estimating optimal treatment regimes (Laber and Zhao, 2015). Tao and Wang generalized the method using the doubly robust approach and developed Tree-based Reinforcement Learning (T-RL) that supports multi-stage treatment decision making (Tao and Wang, 2017; Tao, Wang, and Almirall, 2018).

Electronic Health Records (EHRs) are major data sources for observational studies in the biomedical field, including detailed information about patients medical histories (e.g., diagnoses, medications, and test results). As EHRs are adopted widely across different healthcare institutions, an increasing amount of health information is available to develop and fine-tune optimal DTRs. To the best of our knowledge, existing methods only consider using information from structured data. In EHR, structured data is designed for the management of care or billing purposes. Therefore, critical patient characteristics might not be available in the structured format when studying a specific disease. For example, tobacco use is a risk factor for many chronic diseases, including vascular diseases and lung cancer (Bartal, 2001; Doll, 1998), but some structured EHR archives may not include smoking status as a regular entry. Moreover, structured EHR data is often sparse and partial missingness can occur in many variables. In addition, manual transcription errors may occur in 1-10% of the structured EHR data (Mays and Mathias, 2019). For these reasons, it is challenging to collect various and accurate patient characteristics from structured EHR, which is essential for selecting the correct tailoring variables in optimal DTR. However, extracting additional information from unstructured EHR data may offer a solution.

Narrative content in EHR, like clinical notes, represents a supplementary data resource for patient characterization. Such text information is recorded by healthcare providers describing patient summary, reason of hospital visits, and activities arising from episodes of patient care. Information extraction techniques in clinical notes were well studied in many disease-specific investigations for identifying unique disease conditions (Wang *et al.*, 2018). Most of the information extraction techniques are rule-based, depending on regular expression matching. The clinical Text Analysis and Knowledge Extraction System (cTAKES) is the most popular Information extraction method facilitating biomedical studies based on clinical free-text data (Savova *et al.*, 2010). cTAKES includes many individual tools, including sentence boundary detection, tokenization, named entity recognition, and negation annotation. These individual components can be grouped to handle the extraction of patient-specific characteristics.

Nevertheless, no existing work has evaluated the benefit of performing information extraction using clinical notes to enhance the accuracy of the estimated optimal DTR systematically. In this paper, we developed a technique for extracting patient characteristics from the EHR narrative contents to improve the estimation of optimal DTR, which adopts T-RL as the treatment strategy estimating approach and employs cTAKES components as the IE tool. Our extended simulations demonstrate the benefits of the proposed method under different circumstances. Applying the proposed joint approach to the Medical Information Mart for Intensive Care III (MIMIC-III) database, we exhibit the effective use of information

extraction on estimating an optimal two-stage DTR that guides hypertensive drug use among critically ill patients with severe acute hypertension.

2 Methods

2.1 Notations and problem formalization

In this section, we formally define the problem of optimal DTR estimation using EHR data. Using the general DTR framework, let n denote the number of patients, T denote the number of treatment stages, and $\mathcal{A}_j = \{1, \dots, K_j\}$ represent the observed treatment options at the j^{th} treatment stage, where $j = 1, 2, \dots, T$ and $K_j \geq 2$. For the ICU blood pressure management example, we have $T = 2$, $K_1 = K_2 = 4$. In observational data, patients are observed to follow one of the treatments available at each stage. Let A_j denote the treatment at the j^{th} treatment stage that may take a value a_j , where $a_j \in \mathcal{A}_j = \{1, \dots, K_j\}$. Let $\bar{\mathbf{A}}_T = (A_1, \dots, A_T)$ denote the sequence of treatments until stage T . Similarly, we denote the observed treatment routes with $\bar{a}_T = (a_1, \dots, a_T)$. We use R_j to denote the clinical outcome observed following A_j , which varies under different patient characteristics \mathbf{X}_j and prior treatments received $\bar{\mathbf{A}}_{j-1}$. The overall clinical outcome at stage T is considered a functional of the reward history, $Y = f(R_1, \dots, R_T)$, where $f(\cdot)$ is a pre-specified function (e.g., sum). We assume Y is bounded and larger values of Y is associated with better health outcomes. Stage-wise individualized treatment recommendations are inferred from the observed final outcome Y , the current candidate treatments $a_j \in \mathcal{A}_j$, and the patient medical history $\mathbf{H}_j = (\bar{\mathbf{A}}_{j-1}, \mathbf{X}_j^T)^T \in \mathcal{H}_j$.

When estimating optimal DTR using EHR data, patient characteristics can be observed from two data components: structured data elements and unstructured free-text. We denote the patient characteristics at stage j observed in structured EHR data by \mathbf{S}_j , which describes n_s patients with q_s covariates. \mathbf{S}_j might contain sporadic missing values and entry errors. Let \mathbf{T}_j denote the patient characteristics at stage j extracted from the free-text data, which contains n_t patients and q_t covariates. After combining the two components we obtain \mathbf{X}_j with n patients and q covariates having potentially more patients and covariates than the structured data \mathbf{S}_j . The addition of \mathbf{T}_j can help handle missingness and adding extra (derived) features that are not directly observed from \mathbf{S}_j and correct entry errors. With the observed data, we aim to find a sequence of personalized treatment rules $\mathbf{g} = (g_1, \dots, g_T)$ that maximizes the expectation of the counterfactual clinical outcome $Y^*(\mathbf{g})$ if the \mathbf{g} is followed to make treatment decisions, where g_j is the decision rule at stage j and it maps from patient history \mathbf{H}_j to potential treatments $a_j \in \mathcal{A}_j$. That is, we aim to find the optimal treatment regime \mathbf{g}^{opt} , such that

$$\mathbf{g}^{\text{opt}} = \underset{\mathbf{g}}{\operatorname{argmax}} E\{Y^*(\mathbf{g})\}.$$

2.2 Information extraction from EHRs

Below we describe the proposed technique to extract information and construct structured target variables from clinical free text for optimal DTR estimation. We consider using rule-based information extraction tools that rely on regular expressions (Karttunen et al., 1996), which provide a standard mechanism to select specific strings using a bit pattern. With regular expressions, we can successfully search for patient information consisting of specific keywords and punctuation. Specifically, the proposed protocol closely follows the cTAKES system (Savova et al., 2010) and employs the following individual components: (1) Named entity recognition, which is the core component to identify and locate the target information; (2) Boundary detection, which detects the start and end location of the desired informative substring; and (3) Negation annotation, which helps determine the binary status referring to an identified named entity. The extraction procedure adapts to different types of variables. Numeric variables can be extracted using (1) and (2), whereas (1) and (3) may be used to extract binary and categorical variables.

For each target variable, the method takes the following four steps to extract derived patient characteristics from EHR clinical notes using cTAKES components and regular expressions:

I. Identify the *type and section of the clinical note* that might contain the target information and limit the text materials for extraction accordingly. For example, we search among “discharge summary” or “physician notes” to extract information regarding hospital stay and among “pharmacy” sections for medication usage.

II. *Detect boundaries* of the target information using (2). Patient information is usually contained in one sentence or a number surrounded by words. Identifying the boundaries of the target information helps to further fine prune the target string. Punctuation, including periods and colons and units like “lb,” “cm,” and “kg,” can indicate the beginning of or the end of the target string.

III. Search for *named entities or target keywords*. After obtaining the candidate strings in specific sections of the clinical notes, we search for relevant keywords in those strings with (1). For example, if we are extracting patient height, we can search for “height,” “ht,” or “hgt.”

IV. *Convert target strings into structured data*. For numerical information, we use an ad hoc process by analyzing the possible structures of the target strings and use regular expressions to extract the numerical values and corresponding units. For binary or categorical information, we employ (3) to search for negative tags around (within five tokens of) the target keyword. The default status is truth. If a negative tag is identified, the status of the corresponding condition is mapped to false.

2.3 The estimation of DTR using Tree-based Reinforcement Learning (T-RL)

After combining structured and free-text EHR data, we utilize T-RL (Tao, Wang, and Almirall, 2018) to estimate the optimal DTR. T-RL is a non-parametric optimization method that outputs an unsupervised decision tree for treatment guidance at each stage, where each fork is a split in a tailoring variable, and each leaf node contains a recommended treatment for the corresponding patient subgroup.

When estimating the optimal DTR, we adopt the counterfactual framework for causal inference defined in Robins 1986. Under the three standard assumptions: consistency, no unmeasured confounding, and positivity, we can link the mean of counterfactual outcomes with the observed information. More specifically, the consistency assumption implies that the observed outcome agrees with the counterfactual outcome for any individual patient had the patient been given this specific observed treatment. The no unmeasured confounding assumption guarantees that the treatment assignment is independent of future outcomes given patient history. Due to the nature of observational studies, these assumptions are not testable when using EHR data to estimate optimal DTRs. However, they may be plausible if we have collected sufficient patient information and the data covers a wide variety of patients.

At stage T , we denote $Y^*(A_1, A_2, \dots, A_{T-1}, a_T)$ or $Y^*(a_T)$ as the counterfactual outcome for patients receiving treatment a_T given previous treatments. We aim to search for an optimal treatment regime g_T^{opt} that maximizes the expected counterfactual outcome. Using backward induction, at stage j ($j < T$), we maximize the expected counterfactual outcome when all future treatments are optimized, which is denoted by $Y^*(\bar{A}_{j-1}, a_j, g_{j+1}^{\text{opt}}, \dots, g_T^{\text{opt}})$. However, such counterfactual outcomes are not observable for all patients. We estimate stage-wise pseudo-outcome denoted as $PO_j = \hat{E}[Y(A_1, \dots, A_j, g_{j+1}^{\text{opt}}, \dots, g_T^{\text{opt}})]$ to approximate the target outcome and obtain g_j^{opt} . T-RL uses doubly robust augmented inverse probability weighted estimates to predict patients counterfactual outcomes under all possible treatments.

The T-RL algorithm seeks the optimal regime with a sequence of treatment decision rules at each stage by constructing a binary tree. At any stage, we use $\hat{E}[Y_i(a)]$ to denote the estimated pseudo outcome for patient $i = 1, 2, \dots, n$ given treatment $a \in \mathcal{A}$. For considering each split that separates patient group Ω into ω and ω^c , the T-RL compares the parent node expectation:

$$\mathcal{P}(\Omega, \phi) = \max_{a \in \mathcal{A}} \frac{1}{n} \sum_{i=1}^n \hat{E}[Y_i(a)]$$

with the aggregate expectation of the proposed split-pair of children nodes:

$$\mathcal{P}(\Omega, \omega) = \max_{a', a'' \in \mathcal{A}n} \frac{1}{n} \left(\sum_{i=1}^n \hat{E}[Y_i(a')] I(i \in \omega) + \sum_{i=1}^n \hat{E}[Y_i(a'')] I(i \in \omega^c) \right),$$

where ϕ denotes no partition of the patient group Ω . This comparison determines if the proposed (parent) node split is beneficial. When $P(\Omega, \omega) - P(\Omega, \phi) \gg 0$, the algorithm will make the partition with the corresponding tailoring variable and optimal treatments. To avoid overfitting by pruning the tree, we use stopping rules that consider the minimal node size, minimal improvement for $P(\Omega, \omega) - P(\Omega, \phi)$ and the maximum depth of the tree.

3 Simulation

We conducted two simulation studies in this section. First, we simulated a two-stage ($T = 2$) with three treatments per stage ($K_1 = K_2 = 3$) observational study and its corresponding EHR data to compare the performance of our method with the traditional tree-based reinforcement learning method without information extraction in estimating optimal DTRs. The full data contained stage-wise rewards (R_1, R_2), treatments received (A_1, A_2), 3 complete structured patient characteristics (X_1, X_2, X_3), and 2 patient characteristics that can be extracted from the clinical narratives: weight in pounds (X_4) and smoking status (X_5), where X_5 is binary. We considered cases where X_4 contained missing values or had entry errors and X_5 was not observed in the structured entries. Additionally, we conducted a three-stage ($T = 3$) study that has $K_1 = K_2 = 3$ and $K_3 = 2$ for stage-wise treatments. The patient characteristics, stage-wise rewards and treatments set up for its first two stages were identical to the first study. The third stage with potential treatments A_3 and reward R_3 added complexity to optimal DTR estimation. In both studies, we compared the performance of the estimated optimal DTRs with and without information extraction under the above cases for X_4 and X_5 .

To create a scenario with simulated clinical text, we first sampled 27,707 discharge summaries from the MIMIC-III data (Johnson et al., 2016) that contained information about patients weight and had no information about smoking. According to original MIMIC-III clinical notes, patient smoking status were often described in short phrases throughout the discharge summaries. Thus, we randomly inserted smoking status information in a subset of text documents to simulate smokers and non-smokers. The inserted information included smoker tags like “10 pack-year smoking” and “heavy smoking,” and non-smoker tags like “tobacco: denies.” As a result, we simulated 45% of smokers in our study population.

For all simulation cases, X_1, X_2, X_3 were sampled independently from $N(0, 1)$. When we observed weight information in text, the true X_4 values agreed with the text information; otherwise, it was sampled from $N(195, 51)$, which approximates the weight distribution observed in the original text. The true X_5 is 1 when a smoker tag was assigned and 0 otherwise. We set $\mathcal{A}_1 = \mathcal{A}_2 = \{0, 1, 2\}$. In the first stage, the treatments A_1 followed a Multinomial($\pi_{01}, \pi_{11}, \pi_{21}$) distribution where

$$\begin{aligned} \pi_{01} &= \frac{1}{\exp(0.005X_4 + 0.5X_5) + \exp(0.5X_3 - 0.5X_5) + 1}, \\ \pi_{11} &= \frac{\exp(0.005X_4 + 0.5X_5)}{\exp(0.005X_4 + 0.5X_5) + \exp(0.5X_3 - 0.5X_5) + 1}, \\ \pi_{21} &= \frac{\exp(0.5X_3 - 0.5X_5)}{\exp(0.005X_4 + 0.5X_5) + \exp(0.5X_3 - 0.5X_5) + 1}. \end{aligned}$$

We considered a tree-structured true optimal regime such that

$$g_1^{opt}(\mathbf{H}_1) = I(X_5 > 0) \times I(X_1 > -0.5) + I(X_1 > 0.5).$$

Correspondingly, the stage-wise reward was generated by

$$R_1 = \exp \left\{ 1.5 + 0.003X_4 - |1.5X_5 - 2| \times [A_1 - g_1^{opt}(\mathbf{H}_1)]^2 \right\} + \epsilon_1,$$

where $\epsilon_1 \sim N(0, 1)$. For stage 2, the treatments were distributed as Multinomial($\pi_{02}, \pi_{12}, \pi_{22}$), where

$$\begin{aligned} \pi_{02} &= \frac{1}{\exp(0.2R_1 - 0.5) + \exp(0.5X_1) + 1}, \\ \pi_{12} &= \frac{\exp(0.2R_1 - 0.5)}{\exp(0.2R_1 - 0.5) + \exp(0.5X_1) + 1}, \\ \pi_{22} &= \frac{\exp(0.5X_1)}{\exp(0.2R_1 - 0.5) + \exp(0.5X_1) + 1}. \end{aligned}$$

The optimal decision rule for the second stage depends on the first treatment response

$$g_2^{opt}(\mathbf{H}_2) = I(X_2 > -1) \times [I(R_1 > 0) + I(R_1 > 2)].$$

The second stage reward was generated by

$$R_2 = \exp \{ 1.18 + 0.2X_1 - |1.5X_2 + 2| \times [A_2 - g_2^{opt}(\mathbf{H}_2)]^2 \} + \epsilon_2,$$

where $\epsilon_2 \sim N(0, 1)$. For the two-stage simulation study, the target clinical outcome is the sum of two stage rewards $Y = R_1 + R_2$.

In the 3-stage scenario, data for the first two stages was simulated using the above setting, and treatments $A_3 \in \{0, 1\}$ in the third stage followed *Bernoulli*(π_{13}), where

$$\pi_{13} = \frac{\exp(0.1R_2 + 0.3X_3)}{\exp(0.1R_2 + 0.3X_3) + 1}$$

was the probability of getting $A_3 = 1$. We define the third stage true optimal as

$$g_3^{opt}(\mathbf{H}_3) = I(R_2 > 0).$$

The corresponding reward function was generated by

$$R_3 = \exp \{ 0.1X_2 - [A_3 - g_3^{opt}(\mathbf{H}_3)]^2 \} + \epsilon_3,$$

where $\epsilon_3 \sim N(0, 1)$. The target clinical outcome for the three-stage simulation study is the sum of stage-wise rewards $Y = R_1 + R_2 + R_3$.

We considered 2 possible cases of the observed data. For *Case 1*, some of the X_4 entries contained errors that the observed values were 100 times larger than the true values. The errors were introduced with the following probability

$$P(X_4 \text{ contains entry error}) = 0.1 \times I(X_5 = 0) \times I(\text{weight observed in text}).$$

In *Case 1*, X_4 contained entry errors and X_5 was not observed in structured EHR. When using information extraction, we obtained the full complete data via combining information. For *Case 2*, the structured data contained missing values such that X_4 was missing at random (MAR). The probability of missing was assigned based on X_5 and if weight information was included in the clinical notes:

$$P(X_4 = NA) = 0.1 + 0.5 \times I(X_5 = 0) \times I(\text{weight observed in text}).$$

In Case 2, X_4 was MAR in the structured data but missing completely at random with 10% missing when combining information extracted from clinical text. Also, X_5 was not observed in structured data. Since imputation might alleviate the missing data problem, we also considered imputing X_4 using missForest (Stekhoven and Bühlmann, 2012) under the default setting with all other variables as potential predictors, which created a single complete dataset under each setting. Thus, under Case 2, we also compare the performances between the imputed dataset and the subset of complete observations.

We replicated the simulation 1,000 times with training samples of size $n = 500$ or $n = 1,000$ and test samples of size $n = 1,000$ and applied both methods (with or without information extraction). The estimated DTRs (\hat{g}^{opt}) were evaluated by the percentage of optimal two-stage treatment recommendations (opt%) given to the patients in the test set and the estimated expected counterfactual outcome ($\hat{E}\{Y^*(\hat{g}^{opt})\}$) in the test set using the true rewards model.

Table 1 summarizes the simulation results. Our IE method successfully extracted the true smoking status from the text messages for all patients. Thus, all datasets with IE observed the true X_5 whereas the datasets without information extraction did not contain X_5 . Case 1 mimicked the scenario when entry errors are present for some structured variables. As shown in Table 1, when $n = 500$, for both simulation studies, our proposed method with information extraction under Case 1 greatly improved the accuracy to identify the optimal treatment for study subjects (91.5% vs. 58.8% in the two-stage and 86.8% vs 54.2% in the three-stage study). The estimated mean counterfactual outcome had increased over 20% compared to the traditional method without information extraction. Increasing sample size to 1,000 had improved the performance of both the proposed method and the traditional method, while the advantage of using our proposed method remained similar. For Case 2, where missingness was involved in X_4 , applying our proposed method on the imputed data provided similar opt% and $\hat{E}\{Y^*(\hat{g}^{opt})\}$ compared to the full data without missingness (Case 1 with information extraction). We also noticed that the improvements compared to the naive method were similar between the two-stage and three-stage simulation scenarios. Not surprisingly, the three-stage example required slightly larger sample sizes to achieve the same accuracy level using our method. In general, simulation cases using the imputed datasets outperformed the cases using a subset with only complete observations. However, our proposed method with information extraction always significantly outperformed the naive method across different sample sizes and scenarios.

The empirical distribution of $\hat{E}\{Y^*(\hat{g}^{opt})\}$ is shown in Figure 1. We observed that using the additional information extracted from clinical text, most of $\hat{E}\{Y^*(\hat{g}^{opt})\}$ were closer to their optimal values. With a larger sample size ($n = 1,000$), the values were more centralized towards the optimal counterfactual outcome mean. When using the method without information extraction, the values of $\hat{E}\{Y^*(\hat{g}^{opt})\}$ were more disperse and failed to approach the optimal counterfactual outcome even with a reasonably large sample size $n = 1,000$.

4 Clinical case-study

Now we consider a specific clinical case-study involving personalized antihypertensive agents for critically ill patients with severe acute arterial hypertension.

Severe acute arterial hypertension can cause significant consequences on various organs, including the heart, kidneys, brain, and lungs (Styron et al., 2009; Szczech et al., 2010), leading to life-threatening complications. Acute arterial hypertension is commonly encountered in ICU and acute care settings (Shafi, 2004). Patients with a marked increase in blood pressure and acute severe target-organ injuries (hypertensive emergencies) often require hospitalization in an ICU for immediate blood pressure control (Salgado, Silva, and Vincent, 2013). In “hypertensive emergencies”, the therapeutic strategy requires achieving careful and staged blood pressure lowering goals within 24 hours in order to avoid sudden, excessive reductions. “Hypertensive urgencies” describes the scenario when patients have severely elevated blood pressure but are not in danger of immediate acute end-organ injury. In this scenario, while blood pressure reduction is

Table 1 Simulation results from 1,000 replications. The weight variable (X_4) in structured EHR contained entry errors in Case 1 and had missing values in Case 2. Under both cases, the current smoking status (X_5) was not observed in structured EHR data. $\text{opt}\%$ is the percentage of correctly identifying the optimal treatment combinations among the test sample. $\hat{E}\{Y^*(\hat{\mathbf{g}}^{opt})\}$ shows the estimated expected counterfactual outcome. Standard deviations are recorded in parenthesis.

Two-stage Study				
Scenario	n = 500		n = 1,000	
	opt %	$\hat{E}\{Y^*(\hat{\mathbf{g}}^{opt})\}$	opt %	$\hat{E}\{Y^*(\hat{\mathbf{g}}^{opt})\}$
Case 1 with text	91.5 (13.4)	11.10 (0.76)	97.2 (7.4)	11.35 (0.44)
Case 1 without text	58.8 (8.3)	9.08 (0.39)	63.1 (7.1)	9.31 (0.30)
Case 2 with text, complete observations only	90.1(13.9)	10.93 (0.93)	96.3 (8.8)	11.23 (0.62)
Case 2 with text, missing imputed	91.7 (13.1)	11.11 (0.75)	97.2 (7.4)	11.35 (0.44)
Case 2 without text, complete observations only	57.4 (10.2)	8.75 (0.77)	64.5 (7.0)	9.25 (0.53)
Case 2 without text, missing imputed	61.0 (7.8)	9.22 (0.31)	65.4 (5.6)	9.43 (0.24)
Three-stage study				
Scenario	n = 500		n = 1,000	
	opt %	$\hat{E}\{Y^*(\hat{\mathbf{g}}^{opt})\}$	opt %	$\hat{E}\{Y^*(\hat{\mathbf{g}}^{opt})\}$
Case 1 with text	86.8 (16.1)	12.03 (0.80)	95.8 (8.1)	12.39 (0.47)
Case 1 without text	54.2 (11.5)	9.63 (0.80)	63.7 (7.1)	10.25 (0.49)
Case 2 with text, missing imputed	86.7 (15.9)	12.05 (0.91)	95.3 (8.3)	12.32 (0.56)
Case 2 without text, missing imputed	57.5 (9.8)	10.10 (0.34)	64.1 (6.0)	10.30 (0.24)

warranted, there are no specific evidence-based guidelines on treatment goals. As such, clinical recommendations typically suggest lowering blood pressure less aggressively and to aim for control over the ensuing few days. However, there are no absolute blood pressure thresholds that define hypertensive emergencies or urgencies, as the actual levels differ among individuals depending on a number of characteristics. In general, a clinically ill patient with systolic blood pressure (SBP) levels greater than 180 mmHg or diastolic blood pressures greater than 110 mmHg may require intervention (Marik and Varon, 2007).

Four classes of antihypertensive agents, including angiotensin-converting enzyme inhibitors (ACEI), beta-blockers, calcium channel blockers (CCB), and diuretics, are commonly used to treat hypertension. Studies suggest that antihypertensive drug responses are heterogeneous across patients (Mahmud and Feely, 2007). We estimated a two-stage dynamic treatment regime to guide antihypertensive treatment for critically ill patients with severe acute hypertension using IE and T-RL. The DTR was constructed using The Medical Information Mart for Intensive Care III data (Johnson et al., 2016), a de-identified EHR data with over 40,000 patients who stayed in critical care units at a large tertiary care hospital.

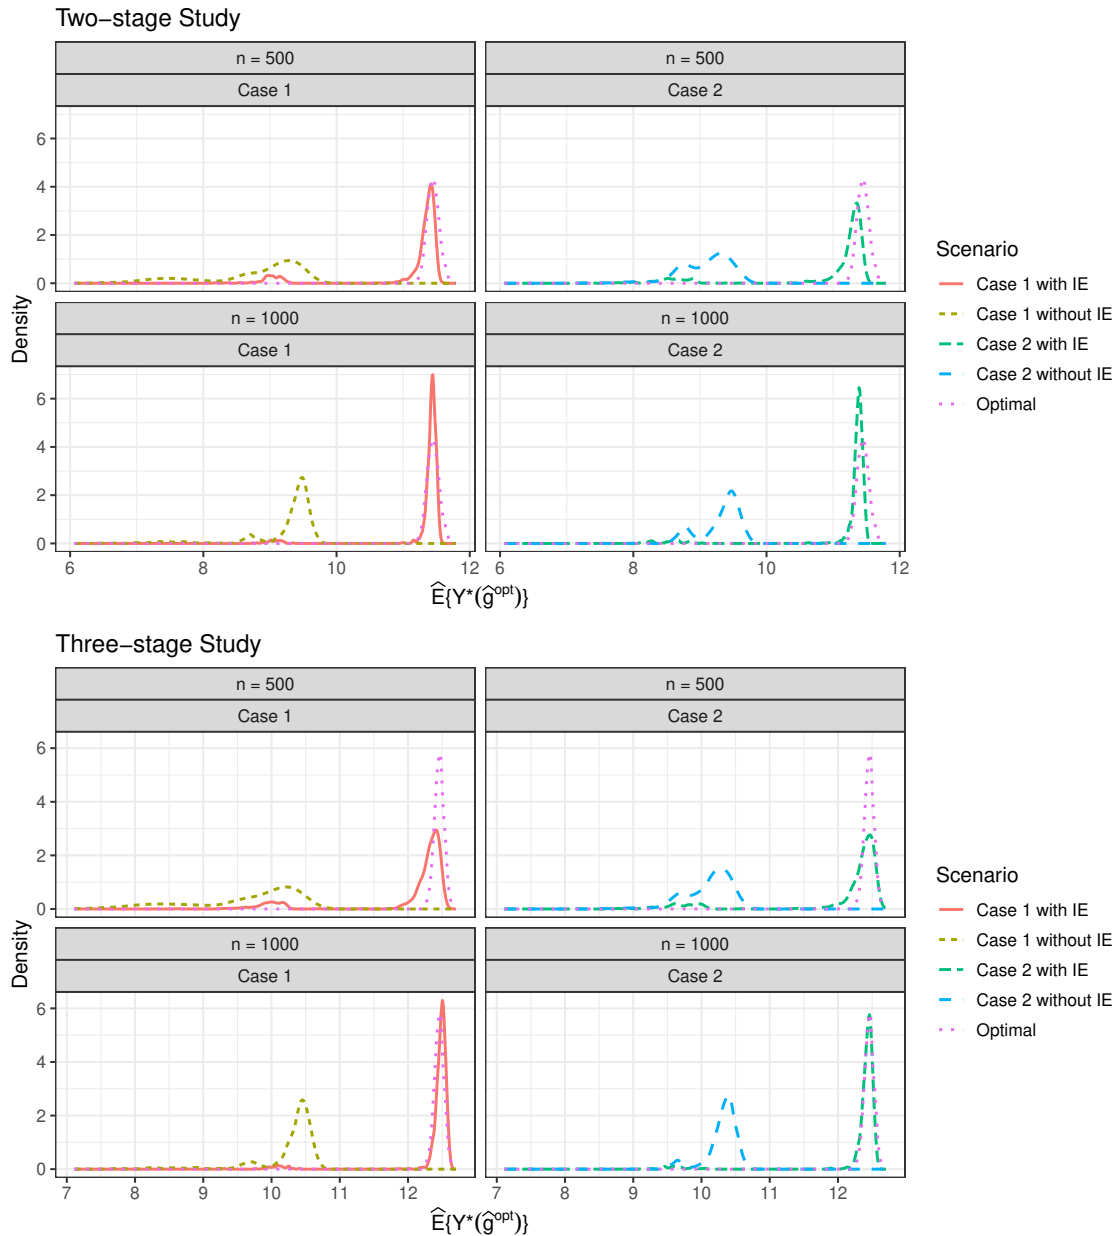


Figure 1 Empirical distribution of $\hat{E}\{Y^*(\hat{g}^{opt})\}$ from 1,000 simulation replications across different simulation scenarios. The weight variable (X_4) in structured EHR contained entry errors in Case 1 and had missing values in Case 2 but got imputed. Under both cases, the current smoking status (X_5) was not observed in structured EHR data.

Patients with the following conditions were included in the study population: (1) admitted to the ICU for at least 3 days; (2) had a first-day maximum SBP higher than 180 mmHg; and (3) had been prescribed only one type of antihypertensive agent during each stage. These inclusion criteria were selected to exclude patients with shock, significant hypotension, and those who do not require or cannot tolerate antihypertensive

therapies. Also, we limited our population to single hypertensive agent receivers to remove the interaction of background antihypertensive medications. Although the majority of patients were not admitted with hypertensive emergency as the principal diagnosis, they had severely elevated blood pressure levels and received antihypertensive treatments. Thus, we assumed that achieving tighter BP control following the intervention is a more successful outcome.

We selected the decrease in SBP as our target clinical endpoint, which is preferable with a higher value. We also assumed excessive reductions of patients blood pressure are not achievable by any single antihypertensive agent. We considered ACEI, beta-blockers, CCB, and diuretics as possible treatments for each patient in both stages, where ACEI, beta-blockers, and CCB were introduced orally, and diuretics were given intravenously (IV). Our goal is to estimate a two-stage DTR on patients blood pressure control during their first two days in ICU. Immediately after the ICU admission, the DTR will recommend the most effective antihypertensive class for the patient with severe acute arterial hypertension to use on day one based on their own clinical characteristics at ICU admission (stage 1). If the maximum SBP is still over 140 mmHg on day two, the DTR will further adjust the hypertension treatment based on patient history and the outcomes will be examined based on patients SBP on day 3 (stage 2).

Studies have shown that many clinical factors, including age, race, smoking status, and weight, are salient predictors of systolic blood pressure and significant risk factors for developing hypertension (Wang *et al.*, 2006). However, the MIMIC-III structured data had no information regarding patients smoking status. In addition, many patients had missing values for their bodyweight upon hospital admission. Without controlling for smoking and bodyweight, the drug effects towards SBP reduction might be biased in the counterfactual outcome model when estimating DTRs. Thus, we utilized the proposed IE method to extract smoking and bodyweight information from physician notes, discharge summaries, and general notes. We detected common patterns in the notes for smoking status by using the named boundary detection, named entity recognition, and negation annotation. For example, “X years smoking history” and “encouraged smoking cessation” indicates current smokers, “quit smoking X years ago” indicates former smokers, and “does not smoke” and “denies any smoking” suggests non-smokers. We assumed patients to be non-smokers when the smoking status was not mentioned in the notes. For bodyweight, we extracted numerical information from patterns like “weight (lb),” “wt,” and “(current): X kg.”

After adding supplemental information from clinical notes, the study population was summarized with 778 complete observations (see Table 2). The majority of patients were in their 60s or 70s when admitted to the hospital. During the first day, beta-blockers (42.4%) was the most commonly prescribed hypertension drug class, followed by diuretics (31.0%), ACEI (17.6%), and CCB (9.0%). Four hundred and forty-two patients had their blood pressure successfully controlled or stopped taking antihypertensive drugs by the end of the first day. During the second day, a larger proportion of the remaining patients had IV diuretics compared to the first day.

All patient characteristics listed in Table 2 were considered as potential tailoring variables for the dynamic treatment regime and were included in the counterfactual outcome model. Supplemental information from clinical notes allowed us to obtain additional bodyweight observations and the whole study population’s smoking status information. Smoking status is an important predictor in the counterfactual outcome models, which provides extra precision. For both stages, we utilized random forest to model the counterfactual outcomes, as previous studies shown that the method accommodates complex individual treatment effect estimation and provides accurate predictions for counterfactual outcomes (Lu *et al.*, 2018; Su *et al.*, 2009; Foster, Taylor, and Ruberg, 2011). Removing smoking status was associated with a 1.2% and 4.5% increase in out-of-bag mean squared error for models at stage 1 and stage 2, respectively.

We apply our proposed method with information extraction with this study cohort. For stage 1 (day one in ICU), our estimate suggests that the optimal treatment was oral CCB for patients with maximum baseline SBP larger than 190 mmHg and minimal creatinine larger than 2 Mg/dL. Otherwise, stage 1 optimal treatment strategy should be ACEI. If we failed to control patients blood pressure during the first day in ICU, patients younger than 70 years of age should receive oral ACEI during the second day, while beta-blockers was the best hypertensive agent for patients older than 70 years. The estimated personalized

Table 2 Descriptive statistics of the variables among the study cohort. Note: * mean (standard deviation) for continuous variables. Otherwise listed as n (%).

Stage	1	2
Total number of patients	778	336
Treatment		
Oral ACEI	137 (17.6)	38 (11.3)
Oral beta-blockers	330 (42.4)	128 (38.1)
Oral CCB	70 (9.0)	23 (10.5)
IV diuretics	241 (31.0)	147 (43.8)
Age at Admission*	68.1 (14.8)	68.8 (13.9)
Female	380 (48.8)	175 (52.1)
Black	138 (17.7)	59 (17.6)
Current or Former Smoker	506 (65.1)	217 (64.6)
Weight (lb)*	176.5 (59.7)	173.2 (55.1)
Kidney disease	140 (18.0)	59 (17.6)
Diabetes	336 (43.2)	142 (42.3)
COPD	38 (4.9)	20 (6.0)
Chronic Hypertension	369 (47.4)	150 (44.6)
Daily Max Systolic BP*	196.6 (16.0)	181.6 (26.7)
Daily Max Diastolic BP*	100.6 (24.1)	91.9 (24.2)
Heart Rate*	80.8 (15.3)	81.3 (14.8)
Temperature (C)*	36.9 (0.6)	37.0 (0.6)
Oxygen Saturation*	97.1 (1.9)	97.0 (1.8)
Daily Maximum Hemoglobin*	11.7 (2.0)	11.6 (13.9)
Daily Minimum		
	1.8 (2.0)	1.6 (1.6)
creatinine (Mg/dL)*		

treatment decision tree is illustrated in Figure 2. In fact, this DTR aligns with the guidelines for the treatment of hypertension by the British Hypertension Society that ACEIs are the most recommended step one antihypertensive agent for younger patients (Williams et al., 2004). Younger patients often respond better to ACEI therapy, potentially due to several factors (e.g., high renin status). Additionally, high creatinine indicates possible acute or ongoing kidney function decline, and in this setting, it is not surprising that ACEI therapy might be less effective or safe for acute blood pressure-lowering given their potential to further drop glomerular filtration rate (Schoolwerth et al., 2001). Thus, our results show that patients with

creatinine higher than normal levels and baseline SBP higher than 190 mmHg might benefit more from CCB. These results can inform clinical practice with selected tailoring variables and meaningful cutoffs.

We further compared the decrease in SBP for the study population under the estimated dynamic treatment regime versus the observed treatment experiences. If the estimated treatment regime were followed, 67.8% of the patients in our study sample would have better control of SBP compared to the observed values during their first two days in ICU.

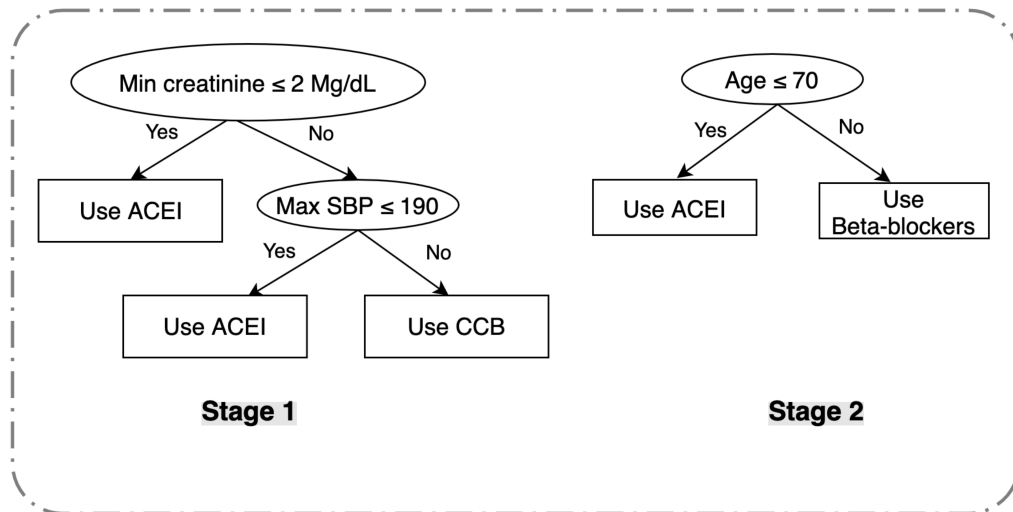


Figure 2 Estimated optimal dynamic treatment regime for blood pressure management among critically ill patients with severe acute hypertension. The optimal DTR was estimated using T-RL with extra information extracted from IE. Stage 1 indicates the time from the first day to the second in ICU and stage 2 indicates the time from the second day to the third day in ICU.

5 Discussion

In this paper, we proposed a joint estimation approach pairing tree-based reinforcement learning with information extraction from free-text clinical observations to estimate the optimal DTR. Our approach can effectively alleviate data quality problems in the structured EHR data, including missing values, erroneous entries, and unobserved risk factors. The reported simulation and clinical-data experiments show that T-RL techniques significantly benefit from the use of IE. This strategy enables clinical decision support for larger study populations, provides more accurate counterfactual outcome modeling, and supports a wider pool of candidate tailoring variables.

The proposed algorithm fills missing values in structured data using clinical text information before the analysis. However, we might still have missingness after the filling completes when information is unavailable in either structured or unstructured data. Depending on the study goals and potential variables involved, we might encounter different missing mechanisms. After merging structured and unstructured data, if we have missing completely at random or missing at random, our method provides valid estimation and inference with imputation. For cases when we have missing not at random (MNAR), without any other prior constraints, it is very difficult to use data-driven approaches to account for systematic differences between missing and observed values. In certain situations with MNAR, analytical bias from imputed data might even exceed the potential bias estimated using the complete observations. We refer readers to the pattern-mixture modeling and sensitivity analysis literature for potential solutions (Little, 1993; National Research Council and others, 2010; Rubin, 2004; Permutt, 2016).

Measurement error is another crucial issue in many empirical applications using EHR data. It may arise as a result of data entry errors, inconsistent documentation of health conditions among physicians, no documentation for out-of-facility services, and inability to capture information in unstructured fields (Chan et al., 2018). In this paper, we have discussed the structured data entry error scenario where clinical text might offer an extra validation source. In our method, we treat structured EHR data as the primary information source. When data is missing or potentially containing errors in the primary source, we seek additional information from clinical text. However, we might have measurement errors in both observed structured variables and clinical text when one of the two data sources is not available. In these situations, we cannot confirm the results by comparing and contrasting two data sources. Hence, the DTR estimation might be biased when the errors are differential. In the existing literature for handling measurement error, popular methods assume additive error in observed variables and suggest modeling the error-prone information using latent variables (Carroll, Gail, and Lubin, 1993; Carroll, 1998; Carroll et al., 2006). While this is beyond the scope of the current study, it is a valuable future research topic.

The proposed method is a promising remedy for several problems regarding estimating optimal DTR using EHR data. However, the improvement of DTR estimation largely depends on the quality of unstructured clinical notes. The benefit of information extraction may be limited when there is little additional informative content embedded in the clinical free-text. Moreover, since optimal information extraction is difficult to attain, some of the extracted information may potentially introduce bias in the final treatment regime. Therefore, cautiously validating the derived information from the original text is recommended. For the rule-based information extraction techniques, scalability can be a potential issue when the number of variables to extract is large. Future studies might consider more advanced or scalable methods such as deep learning for extensive information extraction in a large text corpus when estimating optimal DTRs.

Acknowledgements

This study is supported in part by grants from NSF (1916425, 1734853, 1636840, 1416953) and NIH (UL1TR002240, R01CA233487, R01MH121079, R01MH126137, T32GM141746).

References

- Bartal, M (2001). “Health effects of tobacco use and exposure”. In: *Monaldi archives for chest disease* 56 (6), pp. 545–554.
- Carroll, Raymond J (1998). “Measurement error in epidemiologic studies”. In: *Encyclopedia of biostatistics* 3, pp. 2491–2519.
- Carroll, Raymond J, Mitchell H Gail, and Jay H Lubin (1993). “Case-control studies with errors in covariates”. In: *Journal of the American Statistical Association* 88.421, pp. 185–199.
- Carroll, Raymond J et al. (2006). *Measurement error in nonlinear models: a modern perspective*. Chapman and Hall/CRC.
- Chakraborty, Bibhas and Susan A. Murphy (2014). “Dynamic Treatment Regimes”. In: *Annual Review of Statistics and Its Application* 1 (1), pp. 447–464. ISSN: 2326-8298. DOI: 10.1146/annurev-statistics-022513-115553. URL: <http://www.annualreviews.org/doi/10.1146/annurev-statistics-022513-115553>.
- Chan, Pui Ying et al. (2018). “Peer Reviewed: Using Calibration to Reduce Measurement Error in Prevalence Estimates Based on Electronic Health Records”. In: *Preventing chronic disease* 15.
- Doll, Richard (1998). “Uncovering the effects of smoking: historical perspective”. In: *Statistical methods in medical research* 7 (2), pp. 87–117.
- Foster, Jared C, Jeremy MG Taylor, and Stephen J Ruberg (2011). “Subgroup identification from randomized clinical trial data”. In: *Statistics in medicine* 30.24, pp. 2867–2880.
- Hernn, Miguel A, Babette Brumback, and James M Robins (2001). “Marginal structural models to estimate the joint causal effect of nonrandomized treatments”. In: *Journal of the American Statistical Association* 96 (454), pp. 440–448.
- Johnson, Alistair E W et al. (2016). “MIMIC-III, a freely accessible critical care database”. In: *Scientific data* 3, p. 160035.
- Karttunen, Lauri et al. (1996). “Regular expressions for language engineering”. In: *Natural Language Engineering* 2 (4), pp. 305–328.
- Laan, Mark J Van Der and Daniel Rubin (2006). “Targeted maximum likelihood learning”. In: *The international journal of biostatistics* 2 (1).
- Laber, E B and Y Q Zhao (2015). “Tree-based methods for individualized treatment regimes”. In: *Biometrika* 102 (3), pp. 501–514.
- Little, Roderick JA (1993). “Pattern-mixture models for multivariate incomplete data”. In: *Journal of the American Statistical Association* 88.421, pp. 125–134.
- Lu, Min et al. (2018). “Estimating individual treatment effect in observational data using random forest methods”. In: *Journal of Computational and Graphical Statistics* 27.1, pp. 209–219.
- Mahmud, Azra and John Feely (2007). “Choice of First Antihypertensive: Simple as ABCD?.” in: *American Journal of Hypertension* 20 (8), pp. 923–927. ISSN: 0895-7061. DOI: 10.1016/j.amjhyper.2007.03.011. URL: <https://doi.org/10.1016/j.amjhyper.2007.03.011>.
- Marik, Paul E and Joseph Varon (2007). “Hypertensive crises: challenges and management”. In: *Chest* 131 (6), pp. 1949–1962.
- Mays, James A and Patrick C Mathias (2019). “Measuring the rate of manual transcription error in outpatient point-of-care testing”. In: *Journal of the American Medical Informatics Association* 26 (3), pp. 269–272.
- Murphy, Susan A et al. (2001). “Marginal mean models for dynamic regimes”. In: *Journal of the American Statistical Association* 96 (456), pp. 1410–1423.
- National Research Council and others (2010). “Principles and methods of sensitivity analyses”. In: *The Prevention and Treatment of Missing Data in Clinical Trials*. National Academies Press (US).
- Permutt, Thomas (2016). “Sensitivity analysis for missing data in regulatory submissions”. In: *Statistics in medicine* 35.17, pp. 2876–2879.

- Robins, James M (1994). “Correcting for non-compliance in randomized trials using structural nested mean models”. In: *Communications in Statistics-Theory and methods* 23 (8), pp. 2379–2412.
- (1997). *Causal inference from complex longitudinal data*. Springer, pp. 69–117.
 - (2004). “Optimal structural nested models for optimal sequential decisions”. In: pp. 189–326.
- Rubin, Donald B (2004). *Multiple imputation for nonresponse in surveys*. Vol. 81. John Wiley & Sons.
- Salgado, Diamantino Ribeiro, Eliezer Silva, and Jean-Louis Vincent (2013). “Control of hypertension in the critically ill: a pathophysiological approach”. In: *Annals of intensive care* 3 (1), p. 17.
- Savova, Guergana K et al. (2010). “Mayo clinical Text Analysis and Knowledge Extraction System (cTAKES): architecture, component evaluation and applications”. In: *Journal of the American Medical Informatics Association* 17 (5), pp. 507–513.
- Schoolwerth, Anton C et al. (2001). “Renal considerations in angiotensin converting enzyme inhibitor therapy: a statement for healthcare professionals from the Council on the Kidney in Cardiovascular Disease and the Council for High Blood Pressure Research of the American Heart Association”. In: *Circulation* 104 (16), pp. 1985–1991.
- Shafi, Tariq (2004). “Hypertensive urgencies and emergencies.” In: *Ethnicity & Disease* 14 (4), S232.
- Stekhoven, Daniel J and Peter Bühlmann (2012). “MissForestnon-parametric missing value imputation for mixed-type data”. In: *Bioinformatics* 28.1, pp. 112–118.
- Styron, Joseph F et al. (2009). “Initial emergency department systolic blood pressure predicts left ventricular systolic function in acute decompensated heart failure”. In: *Congestive Heart Failure* 15 (1), pp. 9–13.
- Su, Xiaogang et al. (2009). “Subgroup analysis via recursive partitioning.” In: *Journal of Machine Learning Research* 10.2.
- Szczech, Lynda A et al. (2010). “Acute kidney injury and cardiovascular outcomes in acute severe hypertension”. In: *Circulation* 121 (20), p. 2183.
- Tao, Yebin and Lu Wang (2017). “Adaptive contrast weighted learning for multi-stage multi-treatment decision-making”. In: *Biometrics* 73 (1), pp. 145–155.
- Tao, Yebin, Lu Wang, and Daniel Almirall (2018). “Tree-based reinforcement learning for estimating optimal dynamic treatment regimes”. In: *The annals of applied statistics* 12 (3), p. 1914.
- Wang, Lu et al. (2012). “Evaluation of viable dynamic treatment regimes in a sequentially randomized trial of advanced prostate cancer”. In: *Journal of the American Statistical Association* 107 (498), pp. 493–508.
- Wang, Wenyu et al. (2006). “A longitudinal study of hypertension risk factors and their relation to cardiovascular disease: the Strong Heart Study”. In: *Hypertension* 47 (3), pp. 403–409.
- Wang, Yanshan et al. (2018). “Clinical information extraction applications: a literature review”. In: *Journal of biomedical informatics* 77, pp. 34–49.
- Williams, Bryan et al. (2004). “British Hypertension Society guidelines for hypertension management 2004 (BHS-IV): summary”. In: *Bmj* 328 (7440), pp. 634–640.