

Optimal sepsis patient treatment using human-in-the-loop artificial intelligence

Akash Gupta^{a,*}, Michael T. Lash^{b,1}, Senthil K. Nachimuthu^c

^a California State University, 18111 Nordhoff St, Northridge, CA 91330, USA

^b University of Kansas, 1654 Naismith Dr., Lawrence, KS 66045, USA

^c 3M Health Information Systems, Inc., 575 W Murray Blvd, Murray, UT 84123, USA

ARTICLE INFO

Keywords:

Sepsis
Fluid resuscitation
Artificial intelligence
Optimization
Inverse classifier

ABSTRACT

Sepsis is one of the leading causes of death in Intensive Care Units (ICU). The strategy for treating sepsis involves the infusion of intravenous (IV) fluids and administration of antibiotics. Determining the optimal quantity of IV fluids is a challenging problem due to the complexity of a patient's physiology. In this study, we develop a data-driven optimization solution that derives the optimal quantity of IV fluids for individual patients. The proposed method minimizes the probability of severe outcomes by controlling the prescribed quantity of IV fluids and utilizes human-in-the-loop artificial intelligence. We demonstrate the performance of our model on 1122 ICU patients with sepsis diagnosis extracted from the MIMIC-III dataset. The results show that, on average, our model can reduce mortality by 22%. This study has the potential to help physicians synthesize optimal, patient-specific treatment strategies.

1. Introduction

Sepsis is a life-threatening organ dysfunction caused by a dysregulated host response to infection. Each year, at least 1.7 million adults develop sepsis (Fay et al., 2020). Patients with sepsis are at considerable risk for severe complications and death; one in three hospital deaths are due to sepsis (Rhee et al., 2019). The cornerstone of sepsis treatment is antibiotic administration and fluid resuscitation to correct hypotension. Studies have shown that the type of fluid resuscitation is correlated with mortality (Rochweg et al., 2014), and prescribing excess quantity of intravenous (IV) fluids to septic patients could be detrimental (Vincent & Gottin, 2011). Because of the significance of IV fluid in managing sepsis patients, it is critical to understand what *type* of fluid and what *amount* of fluid should be administered. In this study, we develop a prescriptive clinical model that deduces optimal, patient-specific IV fluid values for the treatment of sepsis.

The international community, Survival Sepsis Campaign (SSC), recommends early goal-directed therapy (EGDT) that underlines the importance of rapid volume resuscitation (Dellinger et al., 2013). Among septic patients, IV therapy is employed as a volume expander in the event of blood loss to keep the body tissue oxygenated. There are two types of volume expanders: *Crystalloids* and *Colloids*. Since *colloids* do not

show clear benefits over crystalloids in treating critically ill patients despite their higher cost (Perel & Roberts, 2012), we aim at finding the right type of crystalloids. The types of crystalloids considered in this study are: Dextrose (a synonym for glucose) 10% in water (D10W), Dextrose 5% (D5W), Dextrose 5% in Normal Saline (D5NS), 5% Dextrose in half normal saline (D5HNS), Dextrose 5% in Lactate Ringer (D5LR), Normal Saline (NS), Half Normal Saline (HNS), and Lactated Ringer (LR). Due to the complexity of patient physiology, there is no consensus on treatment strategy. The lack of treatment standards exacerbates the problem of antibiotics prescription for doctors who do not specialize in critical care or treating sepsis patients (Pulcini, Williams, Molinari, Davey, & Nathwani, 2011).

In teaching hospitals or in many large hospitals, junior doctors are usually responsible for initiating the immediate treatment of patients with severe sepsis for two reasons. First, on-call systems in hospitals are designed to first contact junior doctors when the condition of the patient deteriorates (often governed by the triggering of an early warning system). Second, junior doctors are often the first to attend to patients following hospitalization. SSC guidelines recommend varying the amount of fluid infusion at different disease severity levels. The study showed that 20% of all adults receiving fluid resuscitation experience complications owing to inadequate quantity of fluid resuscitation

* Corresponding author.

E-mail addresses: akash.gupta@csun.edu (A. Gupta), michael.lash@ku.edu (M.T. Lash), snachimuthu@mmm.com (S.K. Nachimuthu).

¹ Akash Gupta and Michael T. Lash equally contributed to this study.

<https://doi.org/10.1016/j.eswa.2020.114476>

Received 18 September 2020; Received in revised form 17 November 2020; Accepted 6 December 2020

Available online 17 December 2020

0957-4174/© 2020 The Author(s).

Published by Elsevier Ltd.

This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

(Courtney, Gopinath, Toward, Jain, & Rao, 2014). By the time senior doctors review the treatment, the early intervention window has passed, which is critical for sepsis treatment (Sherwin, Winters, Vilke, & Wardi, 2017; Gupta, Liu, & Shepherd, 2019). Therefore, developing a clinical tool that can help junior doctors to deduce an optimal treatment is critical.

Some of the significant barriers to clinical models being accepted in the medical community are the issues of interpretability and explainability: medical practitioners expect a model to be transparent and should provide meaningful recommendations. Part of the novelty of our proposed method is that the model considers the physician's own treatment recommendation when eliciting an optimized recommendation. Furthermore, the model can be parameterized to only deviate a small amount from the physician-provided values. Thus, our physician-in-the-loop formulation ensures that medical practitioners can trust the provided recommendations.

The method is designed to provide recommendations that are tailored to each patient's physiology. Patient physiology is characterised using clinical data available in Electronic Health Records (EHR). EHR data is collected throughout a patient's hospital stay. Fig. 1 illustrates a patient's flow through a hospital system, and also highlights the contribution of this paper. Various clinical signs can be collected as the patient progresses through the hospital system. A patient with a suspected infection enters the hospital where the primary assessment of disease criticality is performed. Depending on the severity of the disease, a patient may be recommended for hospitalization. Alternatively, a patient who has been hospitalized may develop signs of infection during the inpatient stay. Following the SSC guidelines, the physician can recommend a specific amount of IV fluids. Our model combines the physician recommendation and the patient's clinical and demographic measurements to determine the optimal quantity of the IV fluids that improves survival probability.

This study makes the following contributions:

1. Formulates a generalized optimization model to prescribe the optimal, patient-specific treatment.
2. Develops a tool to derive the optimal quantity of IV fluids for septic patients in ICUs.
3. Underlines the performance differences between the presence and the absence of human interactions in decision making using the proposed prescriptive model.
4. Employs a known feature selection procedure that integrates well with the established modeling framework.
5. Augments the theoretical literature on inverse classifiers.

Our devised human-in-the-loop method is based on an *inverse classification* framework (Lash, Lin, Street, & Robinson, 2017; Lash, Lin, Street, Robinson, & Ohlmann, 2017; Lash & Street, 2020). This framework includes two steps. First, some arbitrary classifier is trained that best predicts mortality. We employ logistic regression and neural network classifiers in our experiments. Second, the selected model is

embedded into an optimization formulation that takes patient data and physician-provided fluid resuscitation recommendations as input to derive the optimal IV fluid recommendations.

The rest of this paper is structured as follows. Section 2 reviews relevant literature, including existing IV fluid-based sepsis treatment recommendation models and applications of inverse classification to healthcare problems. Section 3 describes the proposed method in detail. Section 4 introduces the EHR data utilized in our experiments to evaluate our methodology. Section 5 presents the experimental settings and numerical results, with a discussion of the results taking place in Section 6. Finally, Section 7 summarizes the contribution of this study and briefly elaborates on future directions of research.

2. Literature review

2.1. Sepsis treatment

Due to complex pathophysiology, the treatment of sepsis imposes many challenges. As a result, the World Health Organization passed a resolution to improve the definition, diagnosis, and treatment of sepsis (Reinhart et al., 2017). The SSC, in their guidelines, advocate for initial fluid resuscitation to treat sepsis effectively (Rhodes et al., 2017). They specified that within three hours of sepsis detection, the patient should be prescribed 30 mL/kg of IV crystalloid fluid. However, the prescribed quantity should also be administered considering the patients' mass and other important vital signs (Cecconi, Evans, Levy, & Rhodes, 2018).

Due to the threat sepsis imposes on human life, many studies have shown that the fluid amount should be considered critically rather than reflexively. Chang and Holcomb (2016) study different fluid resuscitation options available for treating sepsis patients (Chang & Holcomb, 2016). The study discusses the basic physiology, science of intravenous fluid, and compares balanced and unbalanced crystalloids. Durairaj and Schmidt (2008) indicated in their review paper that excessive fluid resuscitation can lead to adverse outcomes (Durairaj & Schmidt, 2008). The authors recommended the use of a dynamic index to determine the fluid amount. This recommendation could be achieved by utilizing optimization models to prescribe patient-specific fluid amounts.

2.2. Expert systems for sepsis treatment

Some researchers have applied machine learning techniques to determine optimal sepsis treatment strategies. Komorowski, Celi, Badawi, Gordon, and Faisal (2018) proposed a model to derive an optimal treatment strategy for sepsis using reinforcement learning (Komorowski et al., 2018). The proposed model harvests the optimal treatment by comparing many treatment decisions. Raghu, Komorowski, Celi, Szolovits, and Ghassemi (2017) extended (pre-print version) the work from Komorowski et al. (2018) by employing the continuous state space solution approach to derive optimal sepsis treatment strategy (Raghu et al., 2017). The authors optimize the quantity of IV fluids and vasopressors such that the treatment outcome can be improved.

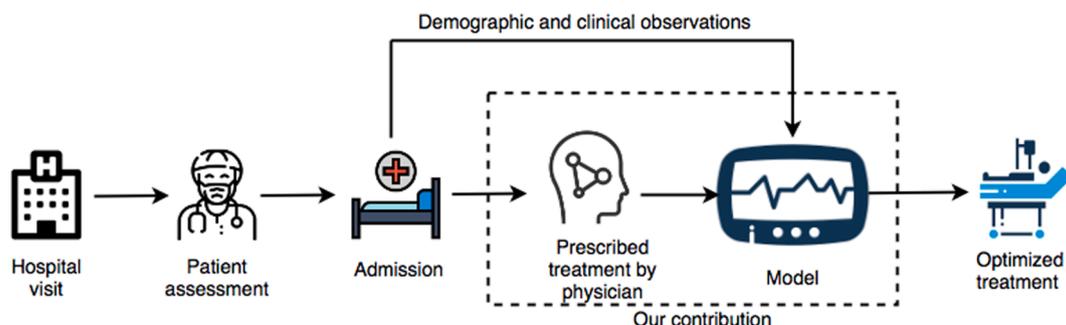


Fig. 1. Patient flow in hospital.

An optimization-driven tool could support junior doctors in determining optimal IV fluid treatments, who are usually among the first to treat patients. Courtney et al. (2014) studied the infusion of intravenous fluids to patients by junior doctors. The study found that the majority of junior doctors do not prescribe the right amount of fluid and fail to adjust for volume by mass (Courtney et al., 2014). The results of this study highlight the importance of a convenient, human (physician)-in-the-loop method that can aid junior physicians in making the right IV fluid prescription.

A few studies have shown that early fluid resuscitation is associated with reduced mortality. Lee et al. (2014) studied the correlation between aggressive, optimal fluid resuscitation administered during the early stages of treatment (Lee et al., 2014). The study found that earlier fluid resuscitation administered in adequate quantities is associated with decreased mortality. Cohen et al. (2015) suggested that, in the absence of any specific anti-sepsis treatment, the early administration of antibiotics and optimal infusion of fluid resuscitation is critical. This study designs a framework that can alert physicians to the risk of mortality during the early stages of a septic episode and assists physicians in prescribing an optimal treatment strategy using our devised inverse classification framework.

2.3. Recommender systems and inverse classification

Recommender systems attempt to find the items, services, products, etc. that will lead an individual to be *most satisfied*. Such methods are typically implemented in environments that contain an over-abundance of decisions, thereby leading users of the environment to a state of “information overload”. For instance, the popular movie and TV streaming website Netflix employs a recommender system to filter through the thousands of entertainment options and display those that the customer will most enjoy. Notions of “most satisfied”, “most enjoyed”, etc. are expressed in terms of some numerical measure. On Netflix, for instance, users express whether or not they enjoyed a particular TV show or movie by “liking” or “disliking” it. Recommender systems leverage a user’s past “likes” and “dislikes” to recommend not-yet-watched movies and TV shows that align with the user’s preferences. Broadly speaking, recommender system literature can be decomposed into three categories: collaborative filtering methods (Su & Khoshgof-taar, 2009), which make recommendations based on user information, content-based filtering methods (Lops, De Gemmis, & Semeraro, 2011), which make recommendations based on content information, and hybrid filtering methods (Burke, 2002), which make recommendations based on both user and content information.

In more recent years, recommender systems methodology has been incorporated into a deep learning paradigm. Deep learning methods can be either collaborative, content-based, or hybrid-based depending on the type of information used as input. These methods are best decomposed in terms of the type of neural network architecture employed. Deep recommender systems employing convolutions (Kim, Park, Oh, Lee, & Yu, 2016), recurrent units (Ko, Maystre, & Grossglauser, 2016), “vanilla” MLPs (Chen et al., 2017), and auto-encoders (Berg, Kipf, & Welling, 2017) have all been explored, as have graph-based methods (Zhang, Yao, Sun, & Tay, 2019; Berg et al., 2017).

While recommender system methodology is important work to consider provided our problem setting, it is not best suited to such a setting since an over-abundance of decisions is not the issue. A related recommendation technology is referred to as inverse classification. Inverse classification makes use of an induced classification model to find the feature value perturbations that optimize for a particular outcome of interest (measured by the induced classifier); the feature value perturbations represent the recommendations. For instance, a model might be induced to learn the mapping from patient characteristics, such as age, blood pressure, and fluid amounts, to a disease outcome of interest, such as survive/die from sepsis (as we do in this paper). Inverse classification will then use this model to find the perturbations to a new patient’s

feature values (i.e., IV fluids prescribed by the physician) that optimally minimize the probability of death due to sepsis.

Inverse classification has previously been applied to a variety of domains, including cardiovascular disease risk mitigation (Chi, Street, Robinson, & Crawford, 2012; Lash, Lin, Street, & Robinson, 2017; Lash, Lin, Street, Robinson, & Ohlmann, 2017; Yang, Street, & Robinson, 2012; Lash & Street, 2020), hiring in nurseries (Aggarwal, Chen, & Han, 2010), bankruptcy prediction and alleviation (Pendharkar, 2002), diabetes (Barbella et al., 2009), and student classroom performance (Lash, Lin, Street, & Robinson, 2017; Lash, Lin, Street, Robinson, & Ohlmann, 2017; Lash & Street, 2020). Several select works use inverse classification to explain rather than prescribe (Barbella et al., 2009; Laugel, Lesot, Marsala, Renard, & Detyniecki, 2018). Methodologically speaking, past inverse classification works can be stratified by those that employ constraints (Barbella et al., 2009; Chi, Street, Robinson, & Crawford, 2012; Lash, Lin, Street, & Robinson, 2017; Lash, Lin, Street, Robinson, & Ohlmann, 2017; Mannino & Koushik, 2000; Yang, Street, & Robinson, 2012; Lash & Street, 2020) and those that do not (Aggarwal et al., 2010; Pendharkar, 2002). Constraints encourage solutions that are real-world feasible and personable (i.e., can be tailored to each individual’s preferences) and are therefore desirable. Furthermore, some inverse classification methods are model-specific (Aggarwal et al., 2010; Barbella et al., 2009; Chi et al., 2012; Mannino & Koushik, 2000; Pendharkar, 2002; Yang et al., 2012), relying on the use of a specific predictive model, while others are model agnostic (Lash, Lin, Street, & Robinson, 2017; Lash, Lin, Street, Robinson, & Ohlmann, 2017; Lash & Street, 2020). If a particular model is able to accurately map the probability space of a particular problem, then a model-specific inverse classification method may be most appropriate. The best/most accurate model is rarely known apriori and is therefore why model-agnostic inverse classification methods are generally preferred.

2.4. Human-in-loop artificial intelligence

Several studies have shown that human interactions combined with machine learning methods can positively influence decision making (Holzinger, 2016). Holzinger et al. (2016) performed numerical experiments to show the advantages of human interactions in solving optimization problems and concluded that having a human-in-the-loop can significantly improve the solution and the computation time (Holzinger et al., 2016). Duhaime (2016) advocates integrating humans and artificial intelligence to solve healthcare problems (Duhaime, 2016). The author claims that such models are more reliable and trustworthy due to the blend of physicians’ knowledge and artificial intelligence.

Further, Amershi, Cakmak, Knox, and Kulesza (2014) present a review of the importance of human interactions in artificial intelligence-based models (Amershi et al., 2014). The study argues that human interactions provide a refinement to the artificial intelligence-obtained solution and aid in selecting alternatives. Gil et al. (2019) propose a human-guided machine learning approach (Gil et al., 2019). In the proposed method, humans interact with machines based on subject knowledge to determine a solution. In addition, the integration of human knowledge into artificial intelligence techniques produces more interpretable results, which is a key requirement for medical applications (Holzinger, Langs, Denk, Zatloukal, & Müller, 2019).

3. Proposed methodology

In this section we disclose our proposed human-in-the-loop method of eliciting optimal, patient-specific fluid resuscitation amounts to reduce the probability of death due to sepsis. The proposed prescriptive model first determines the best predictive model to estimate mortality probability. The selected model is then embedded in an optimization formulation to derive the optimal amount of IV fluid. We describe the selection of predictive model in Section 3.1 and the optimization formulation in Section 3.2. Prior to disclosing our framework we first

provide some relevant preliminary notation.

Let $\{\mathbf{x}^{(i)}, y^{(i)}\}_{i=1}^n$ denote a dataset of n instances (patient visits), where $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^p$ and $y \in \{0, 1\}$. A value of $y = 1$ indicates that a patient has died of sepsis and $y = 0$ indicates that a patient has survived. p is the size of feature vector \mathbf{x} , which represents the number of clinical variables. These include demographic measurements, such as age, vitals, such as blood pressure, lab-based measurements, such as serum creatinine, and the physician-prescribed treatment measurements. However, only specific feature categories can actually be manipulated to affect (i.e., change) the final outcome (survive or not). For instance, one cannot change his or her age, but we can change the amounts of the various IV fluids administered to the patient to improve their probability of survival. Additionally, while certain measurement categories cannot be directly manipulated, the values of the features in these categories may depend and vary according to feature values in other categories. For instance, blood pressure may be a function of both age and the administration of certain drugs.

We first ascribe notation to these different categories. Let U denote the indices in \mathbf{x} that correspond to *unchangeable* features, such as demographic information, D denote the indices corresponding to the *directly changeable* features, such as IV fluids, and I denote *indirectly changeable* features, such as vitals and lab-based measurements. Using these index sets the feature vector \mathbf{x} can be decomposed into $\mathbf{x}_U, \mathbf{x}_I, \mathbf{x}_D$ to refer to specific feature values – i.e., $\mathbf{x} = (\mathbf{x}_U, \mathbf{x}_I, \mathbf{x}_D)$. The feature sets of $\mathbf{x}_U, \mathbf{x}_I$ and \mathbf{x}_D are denoted using $\mathcal{F}_U, \mathcal{F}_I$ and \mathcal{F}_D , respectively; these notations are used predominantly in Section 4.3 to explain feature selection.

3.1. Predicting sepsis outcomes

With the above preliminaries explained, an initial model that provides probabilistic estimates of sepsis outcomes is formulated as follows:

$$\hat{y} = f(\mathbf{x}_U, \mathbf{x}_I, \mathbf{x}_D) = \frac{1}{1 + e^{-g(\mathbf{x}_U, \mathbf{x}_I, \mathbf{x}_D)}} \quad (1)$$

where $f(\cdot)$ is the logistic sigmoid function applied to another function $g(\cdot)$ that takes \mathbf{x} as input. Such a formulation allows for flexibility in defining the model used to make predictions, while still ensuring that probabilistic outputs are obtained from the model. A model that outputs probabilities is important since we wish to minimize the probability of a negative outcome directly, rather than minimizing over a set of discrete outcomes such as $\{\text{alive}, \text{notalive}\}$. When a logistic regression model (Hosmer, Lemeshow, & Sturdivant, 2013) is used, the $g(\cdot)$ function can be expressed as a linear combination of learned parameters and an instance feature vector:

$$g_{\text{logistic}}(\mathbf{x}) = [b, \boldsymbol{\theta}]^\top [1, \mathbf{x}] = b + \sum_{j=1}^p \theta_j x_j, \quad (2)$$

where b and $\boldsymbol{\theta}$ are the learned parameters. If $g(\cdot)$ is a more complex model, albeit still of the variety that can be trained to learn a probability mapping, like a neural network (Suzuki, 2013), the function can be expressed as:

$$g_{\text{NN}}(\mathbf{x}) = \mathbf{w}_k^\top h_{k-1}(\mathbf{W}_{k-1} \dots h_1(\mathbf{W}_1 \mathbf{x})), \quad (3)$$

where \mathbf{w}_k is a parameter vector for the k th (i.e., output) layer, $\mathbf{W}_j : j = 1, \dots, k-1$ are parameter matrices associated with the $k-1$ hidden layers, and $h_j : j = 1, \dots, k-1$ are each some arbitrary, non-linear activation function (e.g., Rectified Linear Unit (ReLU), sigmoid, etc.). An activation function is a general non-linear mathematical function that transforms input in a way that is meant to mimic the firing of a neuron. The ReLU activation function, for instance, is defined as:

$$h_k(\mathbf{W}_k \mathbf{x}) = \max(0, \mathbf{W}_k \mathbf{x})$$

When either g_{logistic} or g_{NN} are selected for use with f , then f and g are learned jointly through a gradient descent optimization process – e.g., gradient descent, projected gradient descent, stochastic gradient descent (in the case of g_{NN}), etc. If the selected learning algorithm does not natively provide probability estimates, however (e.g., g is an SVM model), then g must be learned separately from f , and f reduces to a specific application of Platt Scaling (Platt et al., 1999).

Defining f and g , as generally as we have, allows us to define a considerably larger hypothesis space \mathcal{H} to search across when attempting to find the optimal model to use in our inverse classification framework. The model selection process can be briefly expressed as follows:

$$f^* = \operatorname{argmax}_{f \in \mathcal{H}} \{ \text{Evaluate}(f, \mathbf{X}^{\text{val}}, c) \} \quad (4)$$

where the validation set $\mathbf{X}^{\text{val}} = \{\mathbf{x}_{\text{val}}^{(i)}\}_{i=1}^{n_{\text{val}}}$ is evaluated on trained classifier, f , based on the classification measures, c . In our experiments, classification performance is measured using accuracy and the Area Under the Receiver Characteristic Curve (AUC), explained in Section 5.2. The increased search size of \mathcal{H} is of significance since the quality of our life-saving recommendations are directly tied to the quality of the model and accuracy of the probability space mapping.

3.2. Inverse classification for optimal dose prescriptions

Using the general model defined in Eq. (1) we can obtain probabilistic estimates of death due to sepsis \hat{y} for some instance \mathbf{x} . This model will form the basis for our method of eliciting optimal, patient-specific fluid resuscitation recommendations, which we discuss in this subsection.

Our recommendation formulation is based on inverse classification, which is the process of manipulating the feature values of an instance in order to minimize the probability of an undesirable outcome. In this case, the instances are “sepsis patients” and the undesirable outcome is “death”. Therefore, our formulation will minimize the probability of death, which can be expressed:

$$\min_{\tilde{\mathbf{x}}} f(\tilde{\mathbf{x}}) \quad (5)$$

where $\tilde{\mathbf{x}}$ is a test instance whose feature values are freely manipulable by the optimization process.

The initial formulation, Eq. (5), does not include real-world feasibility considerations, however. For instance, it would make little sense to allow the formulation to manipulate the \mathbf{x}_U feature values, such as age. Moreover, the extent of changes allowed by Eq. (5) is unbounded and may produce nonsensical recommendations, such as negative fluid amounts, as a result. Finally, Eq. (5) doesn't take into account variable dependencies and interactions. As an example, blood pressure, an indirectly changeable feature, is a function of the unchangeable features \mathbf{x}_U , such as age, and the directly changeable features \mathbf{x}_D , representing IV fluids. Since we are manipulating the \mathbf{x}_D (i.e. IV fluids), we can expect blood pressure to also change as a result, and because one's blood pressure has an impact on whether one lives or dies, it is critical to capture these dependencies during optimization.

To account for the dependencies that exist between $\mathbf{x}_U, \mathbf{x}_D$ and \mathbf{x}_I , as discussed in the preceding paragraph, we propose to use a so-called *indirect feature estimator* (IFE). Let $H : \mathcal{X}^{U,D} \rightarrow \mathcal{X}^I$ denote a function that takes the U and D features as input and provides estimates for the I features, i.e.,

$$\hat{\mathbf{x}}_I = H(\mathbf{x}_U, \mathbf{x}_D). \quad (6)$$

Using $H(\cdot)$, we can account for how changes to the \mathbf{x}_D features affect the \mathbf{x}_I features. The IFE can be any differentiable regression model and is therefore also fairly flexible. To be more concrete, we explicitly require a differentiable model because our optimization methodology relies on

gradient information. In this study, we employ linear regression and neural networks. If we need to relax the assumption of differentiability then the optimization methodology would need to be adjusted accordingly (e.g., heuristic optimization). Relaxing this assumption may broaden the IFE hypothesis space at the cost of optimality guarantees.

Considering that (1) the x_U feature values are governed by x_U and x_D and (2) only the x_D feature values can be manipulated, the naive inverse classification formulation of (5) can be transformed into:

$$\min_{x_D} f(x_U, H(x_U, \tilde{x}_D), \tilde{x}_D), \tag{7}$$

where \tilde{x}_D are now the decision variables (to reflect the fact that only these variables should be changed).

While Eq. (7) is an improvement over Eq. (5), there are still necessary considerations missing from the formulation. Namely, Eq. (7) is unconstrained, which may allow the formulation to produce nonsensical recommendations, such as negative fluid resuscitation amounts. Additionally, since a key component of our method is human-in-the-loop functionality (i.e., “doctor-in-the-loop” functionality), we wish to restrict the extent of the manipulations made to the expert-specified x_D feature values. To be a bit more explicit, our rationale is that a doctor’s prescribed fluid resuscitation amounts reflect a coarse-grained recommendation. Our inverse classification method will then refine this initial, doctor-specified recommendation to provide a fine-grained, precise recommendation that does not deviate “too far” from that specified by the doctor.

Therefore, we update Eq. (7) by adding feasibility constraints to prevent nonsensical recommendations as follows:

$$\begin{aligned} \min_{x_D} f(x_U, H(x_U, \tilde{x}_D), \tilde{x}_D) \\ \text{s.t. } \|z\|_1 \leq b \\ 0 \leq \tilde{x}_D \leq 1 \end{aligned} \tag{8}$$

where $z = \tilde{x}_D - x_D$ and b is a budget term that controls the extent of recommendations allowed. Note that \tilde{x}_D are the updated feature values and x_D are the physician-provided feature values.

The b constraint is best calibrated to each physician user in an offline setting, prior to deployment. If a physician is more experienced, the b term can likely be smaller than if the physician is less experienced. A smaller b term will produce smaller recommendations since the cumulative change recommended is restricted.

The last line of Eq. (8) specifies that recommendations must be non-negative, but also less than or equal to one. This latter consideration is based on the assumption that all features have been normalized to a zero-one range and allows us to further encourage the process to produce real-world feasible recommendations. Fig. 2 illustrates the inverse classification formulation in terms of the feature value categories. Note that the human-in-the-loop (HITL) treatments denote the physician-provided IV fluid values.

The final formulation, Eq. (8), can be optimized using project

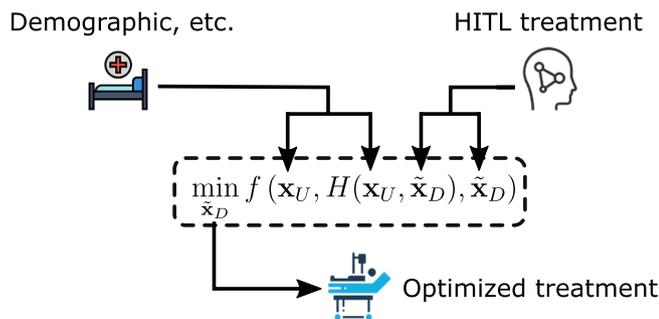


Fig. 2. A depiction of our human-in-the-loop formulation. HITL: human-in-the-loop.

gradient descent (PGD), which is an efficient gradient-based optimization method (Nesterov, 2013). Therefore, the updates to the \tilde{x}_D feature values at each iteration $t = 1, \dots, T$ of PGD can be expressed by:

$$\begin{aligned} \tilde{x}_D^{(t+1)} &= \Pi_{\Delta} \left(\tilde{x}_D^{(t)} - \eta \frac{\partial f}{\partial \tilde{x}_D} \right) \\ &= \Pi_{\Delta} \left(\tilde{x}_D^{(t)} - \eta \left[\frac{\partial f}{\partial \tilde{x}_D} + \frac{\partial f}{\partial x_I} \frac{\partial H}{\partial \tilde{x}_D} \right] \right) \\ &= \Pi_{\Delta} (\tilde{x}_D^{(t)} - \eta \nabla^{(t)} f) \end{aligned} \tag{9}$$

where $\Pi_{\Delta}(\cdot)$ is the projection operator that projects the input onto the feasible region Δ , $\tilde{x}_I^{(t)} = H(x_U, \tilde{x}_D^{(t)})$, and η is the step size. Note that the gradient depends in part upon the estimate of the x_I feature values. Further note that the projection done by $\Pi(\cdot)$ onto the feasible region Δ , dictated by the constraints in (9), can be achieved efficiently and always succeed (as long as the feasible region $\Delta \neq \emptyset$) (Lash, Lin, Street, & Robinson, 2017).

3.3. Human-in-the-loop treatment recommendation

With all of the requisite equations defined, we can now provide the full human-in-the-loop treatment recommendation process, which is disclosed in Algorithm 1. Algorithm 1 takes as input x_D , representing the values prescribed by the physician, x_U , representing the unchangeable values of the instance in question, a learned classifier f^* , a learned indirect feature estimator H^* , a budget constraint b , and two optimization parameters γ and T . The γ parameter represents the minimum relative difference in probability from one iteration to the next that must be observed in order for the algorithm to continue attempting to optimize the D feature values. We set $\gamma = 0.0001$ in our experiments. Line 9 shows how the relative difference is computed (the relative difference is represented by δ). The T parameter specifies the maximum number of iterations the algorithm is allowed to optimize the human-specified x_D feature values and ensures that, regardless of what γ value is specified, the algorithm will terminate in a finite number of iterations.

Algorithm 1. Human-in-the-Loop Treatment Recommendation HITL – TR($x_U, x_D, f^*, H^*, b, \gamma, T$)

- Require $x_D, x_U, f^*, H^*, b, \gamma, T$
- 1: $\tilde{x}_D^{(t=0)} \leftarrow x_D$ (Initialize the free variables to those specified by the physician)
- 2: $\tilde{x}_I^{(t=0)} \leftarrow H^*(x_U, \tilde{x}_D^{(t=0)})$ (Obtain an initial estimate of the indirectly changeable features according to Equation (6))
- 3: $\tilde{y}^{(t=0)} \leftarrow f^*(x_U, \tilde{x}_I^{(t=0)}, \tilde{x}_D^{(t=0)})$ (Obtain an initial outcome probability estimate using Equation (1))
- 4: $\delta^{(t=0)} \leftarrow \gamma, t \leftarrow 1$ (Initialize PGD termination criteria variables)
- 5: **while** $\delta^{(t)} > \gamma$ **AND** $t \leq T$ **do**
- 6: $\tilde{x}_D^{(t)} \leftarrow \Pi_{\Delta} (\tilde{x}_D^{(t-1)} - \eta \nabla^{(t-1)} f)$ (Apply PGD according to Equation (9) to obtain updated \tilde{x}_D values)
- 7: $\tilde{x}_I^{(t)} \leftarrow H^*(x_U, \tilde{x}_D^{(t)})$ (Obtain updated I values using Equation (6))
- 8: $\tilde{y}^{(t)} \leftarrow f^*(x_U, \tilde{x}_I^{(t)}, \tilde{x}_D^{(t)})$ (Obtain an updated outcome probability estimate using Equation (1))
- 9: $\delta^{(t)} \leftarrow (\tilde{y}^{(t-1)} - \tilde{y}^{(t)}) / \tilde{y}^{(t-1)}$ (Compute the relative difference in estimated probabilities between t and $t-1$)
- 10: $t \leftarrow t + 1$
- 11: **end while**
- 12: $\tilde{x}^* \leftarrow \tilde{x}_D^{(t)}$ (Upon termination of the while loop $\tilde{x}_D^{(t)}$ represent the optimal D values)
- 13: **ensure** \tilde{x}^*

Algorithm 1 begins by initializing the freely manipulable variables \tilde{x}_D to the x_D values specified by the physician on Line 1. Subsequently, an initial estimate for the indirectly changeable features is obtained on Line 2. Using the estimated indirect feature values from Line 2, the physician-specified x_D values, and the patient-specific unchangeable feature values x_U , we obtain an initial outcome probability estimate using f^* on Line 3.

On Line 4 we initialize δ and t , which correspond to the aforementioned termination criteria parameters γ and T , respectively. These δ and t parameters are used to control the number of optimization iterations on Line 5.

For each iteration of the optimization while loop, Lines 6 through 10 are executed. Line 6 applies PGD, specified by Eq. (9), to update the \tilde{x}_D variables. It is important to recall that \tilde{x}_D were initialized to the physician-specified x_D values. Therefore, each iteration of the optimization procedure further manipulates the original physician-specified x_D values to continue improving the probability of a good outcome. On Line 7, the indirect feature values are re-estimated using the updated \tilde{x}_D values. On Line 8 a new probability estimate is obtained using both the updated \tilde{x}_D values from Line 6 and the indirect feature values from Line 7. On Line 9 the relative difference in probability between the current iteration and the previous iteration is computed. The iteration counter t is incremented by one on Line 10. When either $t = T$ or $\delta \leq \gamma$ the algorithm terminates and $\tilde{x}_D^{(t)}$ values are set equal to \tilde{x}_D^* , representing the optimal solution, on Line 12. This solution, representing the optimal prescriptions, is then returned by the algorithm.

Additionally, to improve the readability of Algorithm 1 we have provided a corresponding flow chart presented in Fig. 3. Note that the line numbers provided in Fig. 3 correspond to the lines in Algorithm 1.

4. Data preparation

4.1. Data

This study extracted EHR data from MIMIC III (Medical Information Mart for Intensive Care III) (Johnson et al., 2016) to demonstrate the performance of the proposed methodology. These clinical data are freely-available de-identified electronics health records consisting of over sixty thousand patients (61,532) hospitalized in critical care units (ICU) at the Beth Israel Deaconess Medical Center between 2001 and 2012. We selected the adult patient visits with a diagnosis of sepsis for use in our study.

4.2. Data pre-processing

Several necessary pre-processing steps were taken to curate our final dataset. Fig. 4 illustrates these data preparation steps. The data, initially available in multiple tables, are merged to consolidate information about patients, clinical observations and treatments. A particular challenge when pre-processing medical data is the plethora of terms conveying the same meaning (synonyms), which must be accounted for when deriving a suitable dataset for this study. For example, respiratory rate is recorded as *Respiratory Rate*, *Breath rate*, *Res. rate* etc..

Ultimately, the derivation of the dataset used in our study is performed using a clinical expert's opinions and by consulting previous studies related to sepsis (Komorowski et al., 2018; Gupta, Liu, & Crick, 2020). For each patient visit, the clinical variables are recorded longitudinally (i.e., are measured across time). We then aggregate the longitudinal data by computing the mean of each clinical variable for each visit. This aggregation approach was also performed for IV fluids. Any missing data is imputed using multiple imputation by using the chain equation approach (Buuren & Groothuis-Oudshoorn, 2010).

Pre-processing the data results in a total of 1122 patient visits with a diagnosis of sepsis. For each patient visit, we extract 30 clinical features. The statistical summary (minimum, first quarter, second quarter or median, mean, third quarter and maximum) of each feature is included in Table 1. The mean and the median age of patients is 68 and 79 years, respectively. Among all patient visits, the outcome for 244 (22%) visits is expired. The data include a large proportion of patient visits with severe outcomes because the study focuses on ICU patients.

Besides demographics, vitals, and lab-based data, we also extracted the physician-prescribed dosage of IV fluids. Table 2 lists the names of IV

fluids, as recorded in the dataset. As mentioned earlier, a key medical data challenge is the use of synonymous terms; many IV fluids are recorded with different names, but refer to the same fluid. We translate these synonymous fluid names into a standard convention following a clinical expert's suggestions. This study incorporates nine different types of IV fluids. Additives/solutes (e.g. potassium chloride) that were administered via an IV infusion were parsed to obtain the type of the underlying IV fluid. We assume that the absence of a specific type of IV fluid for a patient indicates that the specific type of IV fluid was not prescribed.

4.3. Feature selection

Oftentimes a model with superior predictive performance can be produced by including/excluding certain features. The process of discovering the group of features that produces this superior model is called *feature selection*. Because the inverse classification procedure relies on an underlying model to make life-saving fluid resuscitation recommendations, it is of paramount importance that the most accurate model be learned. Therefore, we propose to employ a feature selection methodology, *Classifier Subset Evaluation*-based (CSE) disclosed by Algorithm 2, to further improve model performance and thereby the accuracy of the probability space mapping captured by the learned model. The adopted CSE method is a specific implementation of the more general *ClassifierSubsetEval* method found in Weka (Frank, Hall, & Witten, 2016).

Algorithm 2 takes as input a training set $\{\mathbf{x}^{(i)}, \mathbf{y}^{(i)}\}_{i=1}^n$, the full feature set \mathcal{F} , an arbitrary, trained classifier f , a classifier performance measure c , such as accuracy or AUC, and s , which specifies the number of consecutive non-improving iterations allowed before termination. Upon execution, the algorithm initializes \mathcal{F}_{opt} , which will hold the selected features, to an empty set, k , which counts the number of iterations, is initialized to zero, the termination criteria *term* is initialized to false, $c^{(k)}$ is the performance of the classifier trained at the k th iteration with $c^{(0)}$ initialized to 0, δ^2 , which measures classifier improvement in terms of c by adding the feature at the k th iteration (i.e., $c^{(k)} - c^{(k-1)}$), is initialized to infinity, and s' , which represents the number of successive, non-improving iterations currently observed, is set to 0. From here the algorithm begins iterating until *term* is set equal to true, the conditions for which are expressed on lines 14 and 15 and will be explained momentarily.

Algorithm

2. Classifier Subset Evaluation $\text{ClassSubEval}(\{\mathbf{x}^{(i)}, \mathbf{y}^{(i)}\}_{i=1}^n, \mathcal{F}, f, c, s)$.

```

Require  $\{\mathbf{x}^{(i)}, \mathbf{y}^{(i)}\}_{i=1}^n, \mathcal{F}, f, c, s$ 
1: Initialize  $\mathcal{F}_{opt} \leftarrow \{\}, k \leftarrow 0, \delta \leftarrow \infty, \text{term} \leftarrow \text{False}, c^{(0)} \leftarrow 0, s' \leftarrow 0$ 
2: while term == False do
3:    $k \leftarrow k + 1, f \sim \mathcal{N}(\mathcal{F}), \mathcal{F} \leftarrow \mathcal{F} \setminus f$ 
4:    $\mathcal{F}_{opt} \leftarrow \mathcal{F}_{opt} \cup f$ 
5:    $c^{(k)} \leftarrow \text{train}^c(\{\mathbf{x}^{(i)}, \mathbf{y}^{(i)}\}_{i=1}^n, f, \mathcal{F}_{opt})$ 
6:    $\delta \leftarrow c^{(k)} - c^{(k-1)}$ 
7:   if  $\delta \leq 0$  and  $|\mathcal{F}_{opt}| \neq 1$  then
8:      $\mathcal{F}_{opt} \leftarrow \mathcal{F}_{opt} \setminus f$ 
9:      $\mathcal{F} \leftarrow \mathcal{F} \cup f$ 
10:     $s' \leftarrow s' + 1$ 
11:   else
12:      $s' \leftarrow 0$ 
13:   end if
14:   if  $s' == s$  or  $\mathcal{F} == \emptyset$  then
15:     term  $\leftarrow$  True
16:   end if
17: end while
Ensure  $\mathcal{F}_{opt}$ 

```

² Note that we have overloaded the definition of δ , which was originally used as part of the stopping criteria in Algorithm 1.

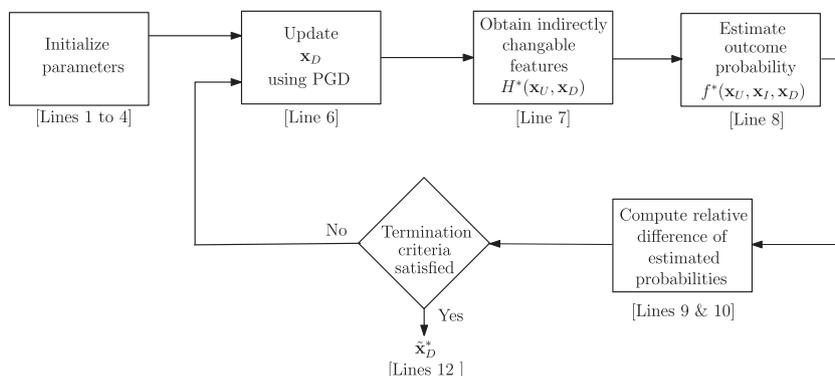


Fig. 3. Pictorial representation of the proposed human-in-the-loop treatment recommendation methodology. The line numbers correspond Algorithm 1.

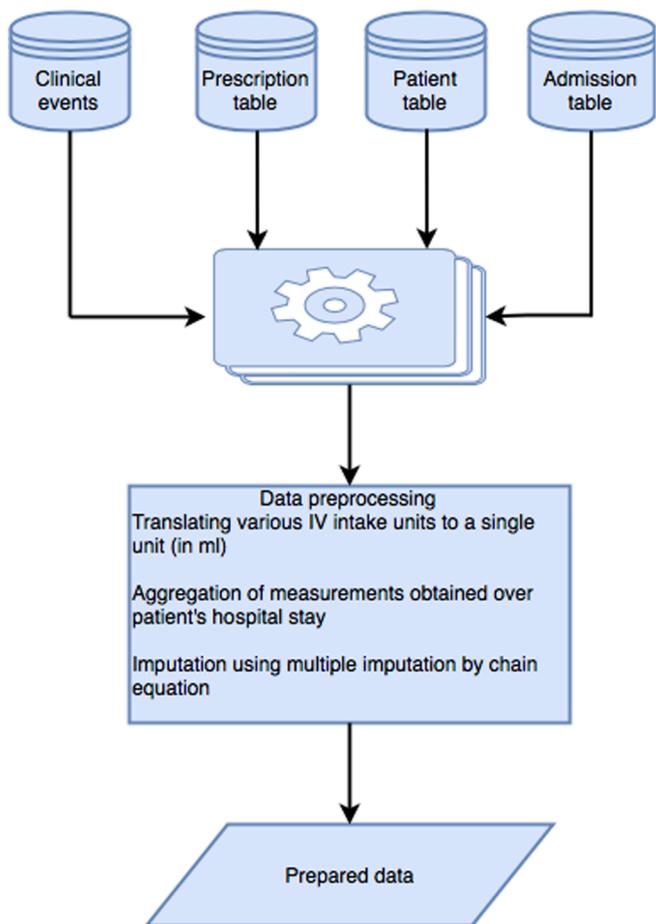


Fig. 4. Data preparation steps.

At each iteration, k is incremented by one, a feature $f \sim \mathcal{U}(\mathcal{F})$ is randomly selected from \mathcal{F} , and \mathcal{F} is updated to exclude the selected feature ($\mathcal{F} \leftarrow \mathcal{F} \setminus f$) (line 3). Next (line 4), the selected feature is added to \mathcal{F}_{opt} where, subsequently, a model f is trained using only the \mathcal{F}_{opt} features (line 5) and evaluated in terms of metric c to obtain $c^{(i)}$; δ is then computed on line 6. On line 7, δ is evaluated to see if adding f has improved predictive performance (this occurs when $\delta > 0$). If f does not improve the model ($\delta = 0$), or worsens model performance ($\delta < 0$), then the addition of f to \mathcal{F}_{opt} is undone: f is re-added to \mathcal{F} (lines 8–9), and s' is incremented (line 10); otherwise (i.e., $\delta > 0$) s' is set to zero (lines 11–12). Line 14 specifies the termination criteria: if s concurrent

iterations have failed to produce an improvement or if all features have been added to \mathcal{F}_{opt} . Utilizing the outlined CSE-based feature selection method, we improve the predictive performance of our model and thereby the accuracy of the probability space mapping.

5. Results

Previously, we elaborated on our proposed methodology and the selected clinical dataset. In this section we illustrate the performance of our model on this clinical dataset. The experimental results are stratified into five segments. First, we experiment with our CFS-based feature selection method while searching for the best predictive model. The best model and the selected subset of features are used in subsequent experiments. Second, we conduct experiments to find the best IFE. Third, using our optimal predictive model and optimal IFE, we derive optimized treatments relative to budget constraint b . Fourth, we perform a robustness check of our model to evaluate performance should physician recommendations not be available. Finally, we present the average changes recommended by our method relative to several selected budget constraint values b .

All results were obtained by first randomly partitioning our dataset into training, validation, and test sets. Since only approximately 20% of the dataset instances (patient visits) belong to the positive class, we ensured that equal proportions of positive instances were allocated to each set (i.e., 20% of the instances in each training, validation, testing set are positive). Dataset sizes were selected to be 80%, 10%, and 10% for training, validation, and test sets, respectively. All f and H models were trained using the training set. The best type of each model (f^*, H^*) was selected based on validation set performance. The testing set was reserved for evaluating our recommendations and was not used in constructing or selecting predictive models. Finally, all features are normalized to a range of $[0, 1]$ using min–max scaling.

5.1. Variable selection

As discussed in Section 3.1, we need to search across \mathcal{H} and select the best predictive model (f^*) to estimate the probability of mortality (Eq. 4). Our first model-building step is to eliminate clinical features that do not contribute to the prediction, and may even detract from predictive performance. Therefore, we establish a variable selection procedure (outlined in a preceding section) that integrates well with our problem setting.

Table 3 lists both independent (or predictors) and dependent (or response) features. The independent variables are divided into three categories: x_D, x_I and x_U . In our dataset, the amount and the type of the prescribed IV fluid are under direct control of the physician; hence, such variables are in the category of directly changeable features (x_D). The vitals (e.g., heart rate, blood pressure) and lab-based measurements (e.g., creatinine, white blood cell count) can not be manipulated directly,

Table 1
Statistical summary of clinical features (Q1: first quarter, Q2: second quarter, Q3: third quarter).

Clinical Variables (units)	Minimum	Q1	Median or Q2	Mean	Q3	Maximum
Base Excess (mEq/L)	-31.0	-5.0	-2.3	-2.7	0.0	16.5
Blood CO ₂ (mEq/L)	4.5	19.8	22.6	22.8	25.8	44.0
Blood Hemoglobin (g/dL)	6.0	9.1	9.9	10.1	11.0	19.5
Blood Urea Nitrogen (mg/dL)	1.0	17.0	29	36.01	48.2	212.5
Body Temperature (°F)	47.4	97.5	98.1	98.1	98.8	107.2
Diastolic Blood Pressure (mmHg)	18.0	50.8	56.4	57.1	63.0	90.3
Glasgow Coma Scale Score	3.0	10.6	13.9	12.5	15.0	15.0
Heart Rate (/min)	46.6	78.4	88.2	89.0	98.9	137.3
Hematocrit (%)	19.7	27.8	29.9	30.7	32.9	61.6
Lactate (mg/dL)	0.6	1.5	2.0	2.5	2.9	18.3
O ₂ Flow (L/min)	0.3	2.2	3.4	5.1	6.3	100.0
PaCO ₂ (mmHg)	19.0	33.6	38.5	39.8	43.8	121.0
PaO ₂ (mmHg)	27.0	83.0	104.1	108.1	127.5	350.0
pH	2.4	7.3	7.4	7.4	7.4	53.5
PO ₂ (mmHg)	26.0	73.3	100.0	103.9	126.7	467.0
PT (s)	11.6	13.8	14.9	16.6	17.6	55.2
PTT (s)	14.9	24.0	27.9	31.7	35.4	128.3
Platelet count (×1000/mm ³)	16.8	145.7	215.3	227.5	288.0	985.0
Respiratory Rate (/min)	10.7	17.7	20.3	20.5	22.9	38.1
Serum Creatinine (mg/dL)	0.2	0.8	1.2	1.9	2.0	141.9
Serum Chloride (mEq/L)	84.0	102.6	106.2	106.1	109.6	137.6
Serum Glucose (mg/dL)	30.3	107.9	126.7	135.4	150.8	447.7
Serum Magnesium (mEq/L)	1.1	1.8	2.0	2.0	2.1	18.3
Serum Potassium (mEq/L)	2.7	3.7	4.0	4.1	4.3	7.3
Serum Sodium (mEq/L)	118.3	136.9	139.2	139.3	141.8	163.1
Systolic Blood Pressure (mmHg)	0.0	102.2	109.9	111.7	120.7	210.1
WBC Count (×1000/mm ³)	0.5	8.3	11.7	13.2	15.8	97.1
Age (years)	19.0	56.3	68.0	66.4	79.0	89.0
Median Weight (kg)	0.0	63.8	76.8	80.3	90.8	233.9

but can be manipulated *indirectly* through manipulation of the administered IV fluids (the x_D features). Therefore, all patient vitals and lab-based measurements fall under the category of indirectly changeable features (x_I). The dataset also includes variables such as age, gender etc. These patient attributes can not be altered. Hence, demographic features fall under the category of unchangeable features (x_U).

Table 2
Available types of IV fluids in the dataset.

IV name (in dataset)	Standardized IV fluid name	Standardized acronym
D10W	Dextrose 10% in Water	D10W
D5 1/2NS	D5 1/2NS	D5HNS
Potassium Chl 20 mEq/ 1000 mL D5 1/2 NS	D5 1/2NS	D5HNS
Potassium Chl 40 mEq/ 1000 mL D5 1/2 NS	D5 1/2NS	D5HNS
D5LR	Dextrose 5% in Lactated Ringer	D5LR
D5NS	Dextrose 5% in Normal Saline	D5NS
Iso-Osmotic Dextrose D5W	Dextrose 5%	D5W
5% Dextrose	Dextrose 5%	D5W
Dextrose 5%	Dextrose 5%	D5W
D5W (EXCEL BAG)	Dextrose 5%	D5W
5% Dextrose (EXCEL BAG)	Dextrose 5%	D5W
Potassium Chl 40 mEq/ 1000 mL D5W	Dextrose 5%	D5W
Amino Acids 4.25% W/ Dextrose 5%		DNS
1/2 NS	Half normal saline	HNS
0.45% Sodium Chloride LR	Half normal saline Lactated Ringers	HNS LR
Lactated Ringers	Lactated Ringers	LR
NS	Normal Saline	NS
SW	Normal Saline	NS
0.9% Sodium Chloride	Normal Saline	NS
0.9% Sodium Chloride (Mini Bag Plus)	Normal Saline	NS
NS (Mini Bag Plus)	Normal Saline	NS
Iso-Osmotic Sodium Chloride	Normal Saline	NS
Isotonic Sodium Chloride	Normal saline	NS
NS (Glass Bottle)	Normal Saline	NS
Potassium Chl 40 mEq/ 1000 mL NS	Normal Saline	NS
Potassium Chl 20 mEq/ 1000 mL NS	Normal Saline	NS

Table 3
Clinical features in the dataset.

Variable category	Notation	Complete set of variables	Selected variables
Independent	x_D	Amount of fluid resuscitation (D10W, D5HNS, D5LR, D5NS, D5W, DNS, HNS, LR and NS)	D5HNS, D5LR, D5W, LR, NS
	x_I	Base excess, blood CO ₂ , blood hemoglobin, blood urea nitrogen, body temperature, diastolic blood pressure, Glasgow Coma Scale, heart rate, hematocrit, lactate, O ₂ flow, PaCO ₂ , PaO ₂ , pH, PO ₂ , PT, PTT, platelet, respiratory rate, serum creatinine, serum chloride, serum glucose, serum magnesium, serum potassium, serum sodium, systolic blood pressure, WBC (27 variables)	Base excess, blood CO ₂ , blood hemoglobin, blood urea nitrogen, diastolic blood pressure, Glasgow Coma Scale, heart rate, hematocrit, lactate, PaCO ₂ , PaO ₂ , PO ₂ , PTT, platelet, respiratory rate, serum creatinine, serum glucose, serum sodium, systolic blood pressure, WBC (20 variables)
	x_U	Age, gender, weight	Age, gender
Dependent	y	Discharge type (binary)	Discharge type

Feature selection was performed using CSE, discussed in Section 4.3. CSE can be applied to each $f \in \mathcal{F}$ in an “online” fashion, or as a pre-processing step, where a single f is selected and used to find the subset

of features that will be used when searching for $f^* \in \mathcal{F}$. We adopt the latter, pre-processing strategy to reduce experiment compute time. Our selected f was a neural network model trained for 150 epochs with a single hidden layer containing three hidden nodes. This particular parameterization of f was selected because it is representative of the parameterizations explored during the model tuning phase, the results of which are discussed in the next subsection.

Table 3 shows the variables before and after employing the variable selection procedure. The variable selection results show that *D5HNS*, *D5LR*, *D5W*, *LR* and *NS* are the only fluids that affect the probability of mortality significantly. We also observe that the size of the indirect variable set has been reduced from 27 to 20 features. *Discharge type* is considered the outcome, or dependent feature. In the next subsection, we show the performance of the predictive model on the data with the complete set and on the data with the selected features.

5.2. Predictive model tuning

In Section 5.1, CFS is utilized to find the best subset of features. In this section, we employ a grid search to find the model f^* that has the best predictive performance and therefore the best probability space mapping. We apply the grid search to two datasets: a dataset containing the full set of features and a dataset containing only those features that were selected using CFS. This will allow us to choose not only the best model, but to assess whether CFS is in fact able to produce a superior model. We limit our study to logistic regression ($g_{\text{logistic}}(\mathbf{x})$) and neural network variants ($g_{\text{NN}}(\mathbf{x})$). After determining the optimal f^* and dataset, a grid search is performed to find the optimal IFE function H^* . We limit this grid search to multivariate linear regression and variants of neural networks.

Accuracy and AUC are employed to assess the performance of our classification models $f \in \mathcal{F}$. Accuracy is defined as the ratio obtained by dividing the number of correctly predicted instances by the total number of instances. As a metric, however, accuracy is susceptible to class imbalance, which is present in our dataset (i.e., only 20% of instances

are positive). Therefore, we also adopt the AUC metric, which is insensitive to class imbalance. AUC plots the true positive rate (TPR) against the false positive rate (FPR), thus creating a curve. The area underneath this curve (called the receiver operating characteristic curve) is the AUC (area under the curve). When a model predicts *only* the majority class (e.g., the model always predicts the negative class), the AUC is 0.50 and is why AUC is considered insensitive to class imbalance. An AUC of 1.0 represents completely perfect predictions.

Table 4 presents the results of our grid search in terms of accuracy and AUC, obtained on both the training and validation sets, on the two datasets (original and feature selected) discussed earlier in this section. Note that “Tr” stands for “training set” and “Val” for “validation set”. We varied the number of training epochs from 100 to 250 for all models. We also varied the number of hidden nodes in our neural network from 3 to 10. We concentrated our grid search on these specific values after experimentally noticing that performance degraded or did not improve when smaller or larger epoch/hidden node values were parameterized. Note that logistic regression can be viewed as a neural network with no hidden nodes and layers. The Adam optimizer was used for training all models with an initial learning rate set to 0.01, the exponential decay rate for the first moment estimates to .90 (the default), and the exponential decay rate for the second moment estimates to 0.999 (the default) (Kingma & Ba, 2015). Furthermore, we set the batch size of all models to 64 and did not employ dropout or weight regularization since the models were relatively small in size (i.e., contain few nodes and only a single hidden layer). All of our experiments were coded using Python 3.6 and employed the Tensorflow 2.0.0 package for training our machine learning models. We have made our code publicly available at <https://github.com/michael-lash/HITL-Sepsis-OO>.

Upon executing the grid search we found the best performing model f^* is obtained using a neural network with three hidden nodes (AUC: 0.8792) on the “feature selection” dataset. Therefore we adopt this model and dataset for the remainder of our experiments.

Next, we execute a grid search to find the optimal IFE function H^* . Recall that the IFE takes as input x_I and x_D and provides estimates for x_I and will be used during the recommendation procedure. Fig. 5 illustrates

Table 4

Grid search results for both the original and “feature selection” datasets on a randomly held out validation set. HN: Hidden nodes, Tr Acc: Accuracy on training data, Tr AUC: AUC on training dataset, Val Acc: Accuracy on validation dataset, Val AUC: AUC on validation dataset.

Models	Parameters		All Features				Selected Features			
	HN	Epochs	Tr Acc	Tr AUC	Val Acc	Val AUC	Tr Acc	Tr AUC	Val Acc	Val AUC
Logistic		100	0.8185	0.8988	0.7768	0.8523	0.8029	0.8820	0.7768	0.8257
		150	0.8151	0.9001	0.7589	0.8510	0.8163	0.8953	0.7946	0.8437
		200	0.8285	0.9101	0.7946	0.8695	0.8285	0.9058	0.8304	0.8615
		250	0.8229	0.9111	0.7946	0.8665	0.8263	0.9101	0.8036	0.8654
NN	3	100	0.8296	0.9137	0.8036	0.8691	0.8263	0.9061	0.7857	0.8679
	3	150	0.7996	0.8942	0.7589	0.8437	0.8096	0.8975	0.7679	0.8601
	3	200	0.8408	0.9224	0.7857	0.8735	0.8318	0.9156	0.7946	0.8777
	3	250	0.8374	0.9235	0.7768	0.8690	0.8374	0.9205	0.8125	0.8792
	5	100	0.8285	0.9057	0.7857	0.8624	0.8007	0.8854	0.7679	0.8427
	5	150	0.8151	0.8906	0.7679	0.8367	0.8051	0.8830	0.7768	0.8481
	5	200	0.8452	0.9228	0.8036	0.8633	0.8352	0.9172	0.8214	0.8774
	5	250	0.8530	0.9228	0.8036	0.8565	0.8352	0.9159	0.8125	0.8736
	10	100	0.8218	0.9083	0.7857	0.8632	0.8196	0.8976	0.7679	0.8593
	10	150	0.8140	0.8909	0.7768	0.8461	0.8040	0.8898	0.7768	0.8601
	10	200	0.8463	0.9285	0.7946	0.8724	0.8474	0.9234	0.8125	0.8760
	10	250	0.8419	0.9268	0.7768	0.8709	0.8441	0.9242	0.8125	0.8728

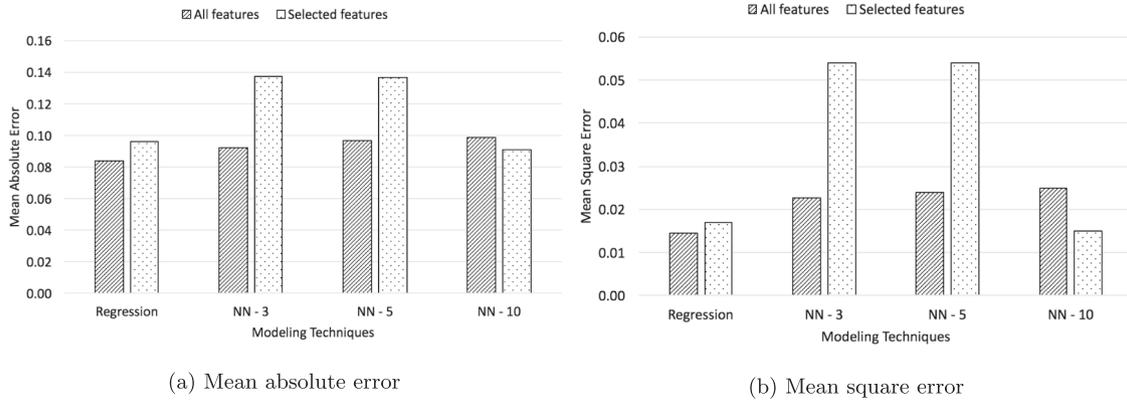


Fig. 5. Model performance of indirect feature estimators on both the “original” and “feature selection” datasets on a randomly held out validation set. NN-3, NN-5 and NN-10 are the neural network with 3, 5 and 10 hidden nodes, respectively.

the performance of the model in the validation dataset with 250 epochs. We explored epochs ranging from 100 to 350 and found that 250 epochs produced the best results. Figs. 5a and 5b show Mean Absolute Error (MAE) and Mean Square Error (MSE), respectively. We utilized MSE as the performance measure to select the best model. The results show that a neural network with ten hidden nodes produces the best model with an MSE of 0.015.

Table 5 lists the selected models that will be used in subsequent experiments.

5.3. Human-in-the-loop recommendations: probability improvement

Using the optimal $f^*(\cdot)$ and $H^*(\cdot)$, discovered in the preceding subsection, we apply our human-in-the-loop inverse classification formulation, discussed in Section 3.2, to the test set. The experiments are performed by varying the budget b from 0.1 to 1 in increments of 0.1. In these experiments the IV fluid values, x_D , specified by each physician are cumulatively allowed to be changed by only an amount b , according to our formulation in Section 3.2. Therefore, larger values of b will allow larger changes to be made to the physician’s recommendation.

Our method derives optimal, patient-specific IV fluid recommendations along with an estimated risk (probability) of mortality (*objective function*) at the varying budget levels mentioned. Fig. 6 shows the average results across these different budget levels. Fig. 6a shows the average (μ) probability of mortality (i.e., death due to sepsis; y-axis) at each budget level (x-axis), as indicated by the blue line; the gray shading indicates one quarter of one standard deviation above and below the average ($\mu \pm \frac{1}{4}\sigma$). As expected, with an increase in b , the probability of mortality is further reduced (on average). The average probability of mortality with no adjustment in x_D is 0.46, while the average probability of mortality with x_D allowed to be adjusted to the maximum (i.e., $b = 1.0$) is 0.37.

Fig. 6b shows the relative probability improvement (y-axis) at each budget setting (x-axis), indicated by the red line; the gray shading indicates one quarter of one standard deviation above and below the average. With an increase in b , the probability of mortality is further reduced (on average). We can observe that the average relative improvement in mortality, provided a limited budget $b = 1.0$, is about 22%, which is significantly better than the 1.8–3.6% improvement found in Raghu et al. (2017). Therefore, the proposed model can

Table 5 Selected modeling technique (star indicates the best selected model).

Functions	Selected Model
f^*	Neural network with three hidden nodes and sigmoid output activation
H^*	Neural network with ten hidden nodes with no output activation

significantly improve the chances of survival by adjusting the infusion of IV fluids to the optimal value.

5.4. Human-in-the-loop recommendations: robustness

Fig. 7 illustrates the results obtained using human-in-the-loop initialization and random initialization. We refer to the scenarios where the physician’s recommendations are incorporated as *human-in-the-loop* initialization, while the scenarios with small, randomly specified input are referred to as *random* initialization. These small, random input values were drawn uniformly from a range of [0, 0.1], reflecting a “cautious” random initialization. We use the term *initialization* because the input x_D values represent the starting place for the optimization procedure. Therefore, the specified x_D values likely have an impact on the recommendation and, consequently, the amount of probability improvement that can be extended to each patient (test instance). Similar to the results presented in the preceding subsection, we show both actual and relative probability improvement at varying budget levels.

Fig. 7a show the average probability improvement results. The x-axis represents the budget and y-axis represents average probability of mortality. The green line represents human-in-the-loop initialization and the purple line represents random initialization. Clearly, integrating a “human into the loop” produces reduced mortality results, as compared to the random initialization result. The results demonstrate the importance of the “human-in-the-loop” component of our formulation. Nevertheless, the results also show that poor (i.e., random) initializations can still be turned into recommendations that provide comparable benefits (in terms of probability improvement).

Similarly, Fig. 7b shows relative probability improvement across varying budget levels using both types of initialization. Again, the green line represents human-in-the-loop initialization and the purple line represents random initialization. Here, we make the same observations that we do for Fig. 7a. However, we can also see that as the budget is increased, the results obtained from random initialization tends closer to those obtained from physician initialization. This observation further shows that, provided a sufficiently large budget, random initialization can come close to providing the same probabilistic improvement.

5.5. Human-in-the-loop recommendations: average recommendations

In this section, we examine the recommendations produced by our method. Fig. 8 shows the average recommendation results obtained from applying our method. Each of the Figs. 8d represent a different budget level (0.3, 0.5, 0.7 and 0.9). The x-axis shows each IV fluid. The y-axis represents the average recommended change to a physicians input. It is impractical to present individual recommendations and, as

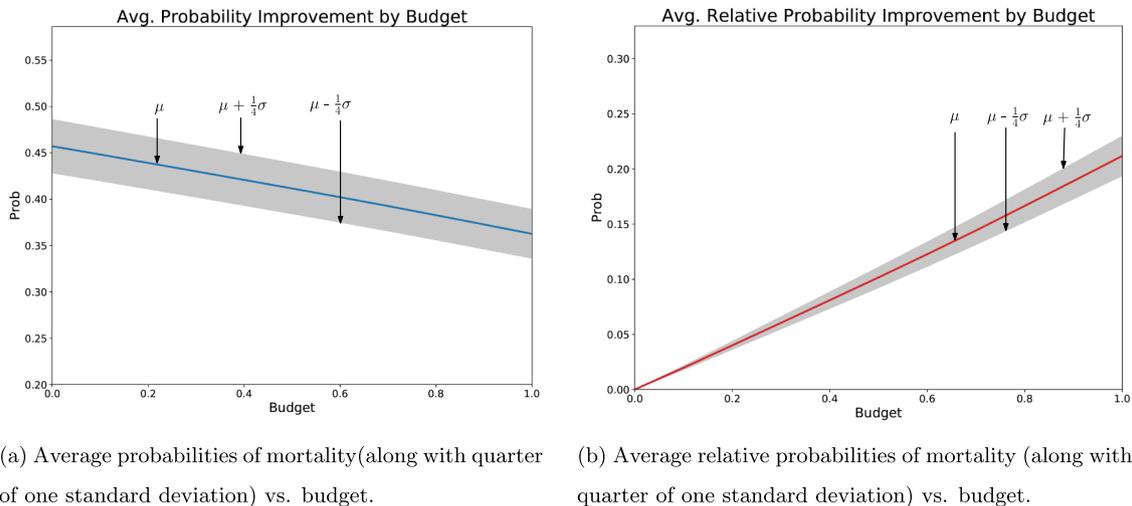


Fig. 6. The average change in the probability of mortality over varying budgets.

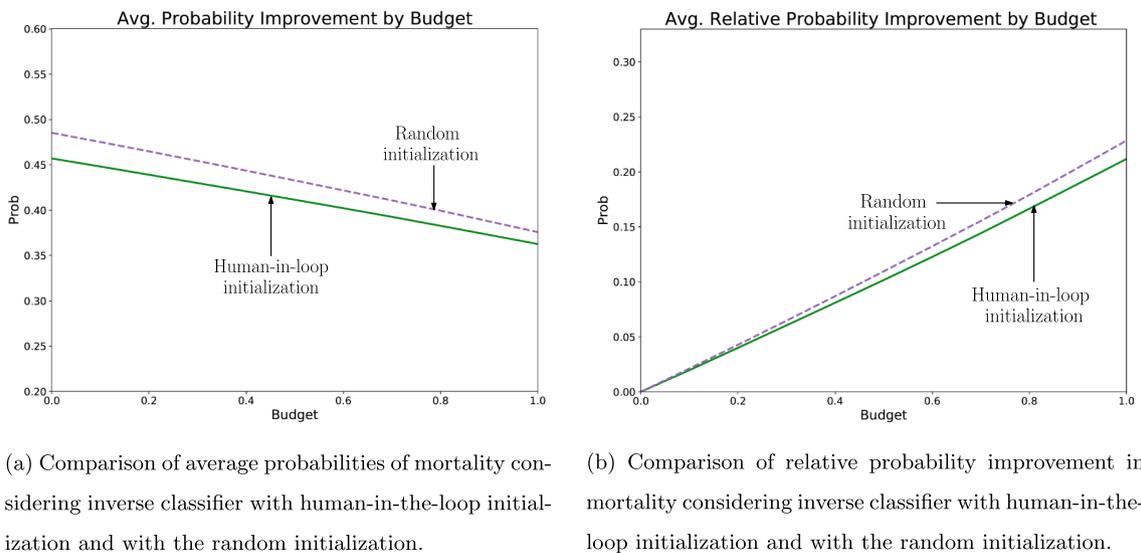


Fig. 7. Human-in-the-loop vs. random initialization results.

such, we present average recommendations by budget value.

The deviation from a value of zero represents a recommended change. A positive (negative) value indicates that the physicians recommended IV fluid value should be increased (decreased). The results are further stratified by “predicted positive”, shown in red, and “predicted negative” shown in blue. An instance was predicted as belonging to the positive class, representing “expiration due to sepsis”, if their predicted probability of mortality was greater than 50%, and negative otherwise.

6. Discussion

In the previous section (Section 5), we showed the efficacy of our proposed method through a variety of experiments. Since the proposed method uses two levels of estimation – H to estimate the indirectly changeable features (x_i) and f to estimate the outcome (probability of mortality) – we first conducted experiments to find the optimal f^* and H^* . We found that, in both cases, some variation of neural network produced the best predictive performance (see Table 4 for the best performing f and Fig. 5 for the best performing H). One possible reason for such performance is that neural networks can capture non-linear

relationships among the covariates. Thus, if nonlinear interactions among features are truly indicative of the associated outcome (mortality in this case), neural networks tend to produce better models than their linear counterparts (logistic regression), which is what we observe here. Furthermore, our experiments to find f^* also showed that our feature selection dataset (obtained using the CSE-based method disclosed by Algorithm 2) produced a superior model. We believe the reason for this is the same as our reasoning behind introducing the CSE-based method used to obtain the feature selection dataset in the first place: features that are irrelevant to the prediction can introduce spurious noise into the model, thus detracting from predictive performance. Intuitively, when the “right” set of predictive features are used to construct the predictive model, predictive performance is improved.

After obtaining f^* and H^* , our experiments concentrated on evaluating our proposed human-in-the-loop treatment optimization methodology. The proposed model takes a physician’s initial treatment prescription as input to determine the optimal, patient-specific IV fluid dosing. Rudd et al. (2018) highlight the major challenges to the global burden of sepsis and underline the importance of developing data-driven and patient-specific treatment strategies (Rudd et al., 2018). Our study aims at eliminating this barrier and providing personalized treatment

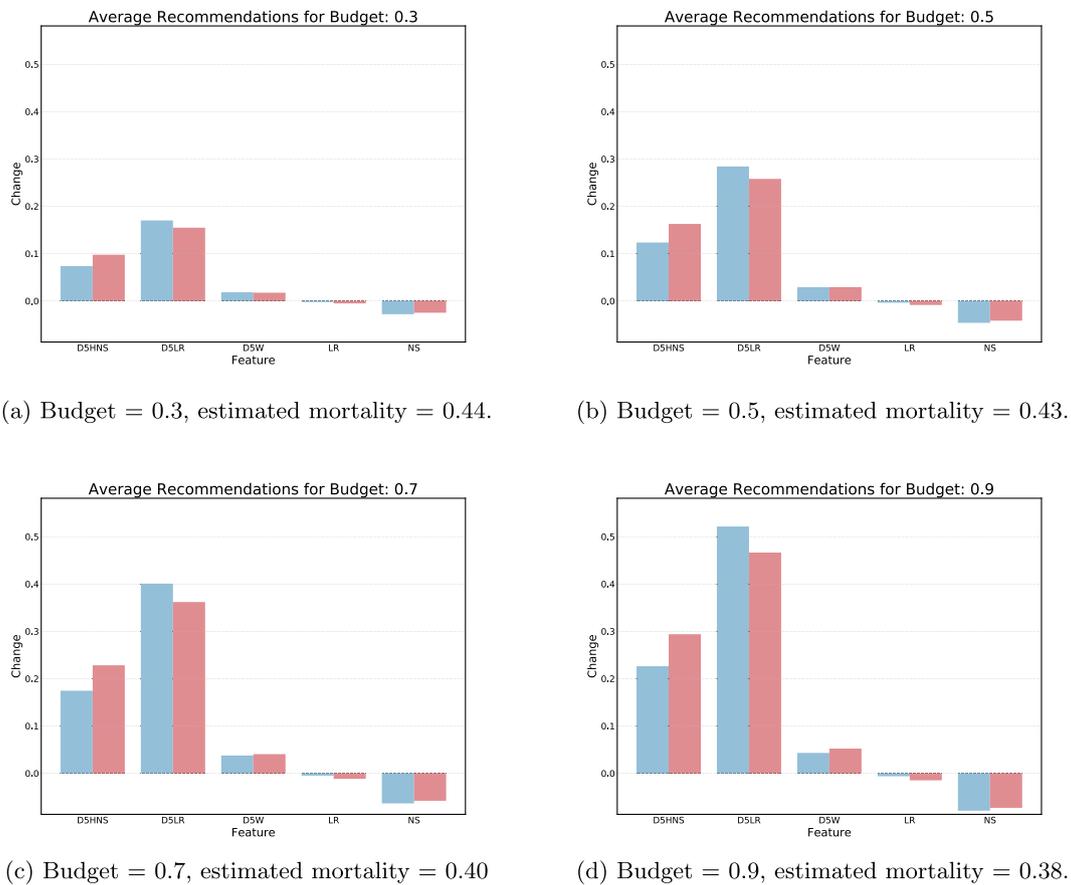


Fig. 8. Average recommendations at budget levels 0.3, 0.5, 0.7, and 0.9, stratified by predicted positive test instances, shown in red, and predicted negative test instances, shown in blue.

recommendations. To demonstrate this, we applied our method to our dataset and showed the benefit in terms of average probability improvement (Fig. 6a) and relative average probability improvement (Fig. 6b) in Section 5.3. We find that our model improves the relative probability of mortality by 23% on average, depending on the budget. By contrast, Raghu et al. (2017) are only able to provide a 3.6% relative improvement in mortality (Raghu et al., 2017).

Following the promising results obtained in Section 5.3, we wanted to investigate the robustness of our model in the absence of any physician (i.e., intelligent) recommendation. Therefore, we compare the results obtained using physician recommendations to those obtained using random inputs (Section 5.4). For each test instance, we randomly initialized each of the features to values in the range of $[0, 0.1]$. This range of values was selected to reflect a *cautious* initialization (i.e., small, rather than large values). The results show that our model can be used in a standalone fashion (i.e., without initial physician recommendations) to generate life-saving recommendations (Fig. 7). The results show that physician input clearly improves the obtained result, however, particularly at more modest budget levels. Nevertheless, these results underscore the robustness of our method and demonstrate that faulty initial recommendations can be overcome. Robustness is a desirable facet since junior doctors with limited experience are prone to error (Courtney et al., 2014), yet are usually the first to be called by hospital administration when patient conditions deteriorate (Callaghan, Kinsman, Cooper, & Radomski, 2017). The model can provide additional support to junior doctors to make informative decisions in the early stage of treatment.

In Section 5.5, we present the average recommendations made by our human-in-the-loop method. While many scoring criteria (Ferreira, Bota, Bross, Mélot, & Vincent, 2001; Gupta, Liu, Shepherd, & Paiva,

2018; Vincent et al., 1996) and tools (Gupta et al., 2020; Henry, Hager, Pronovost, & Saria, 2015) exist to assess the risk of mortality of septic patients, there are limited studies that comprehensively assess risk, account for physician input, and provide treatment recommendations. Our study and proposed method provides all three of these benefits.

Upon inspecting Figs. 8d, we can clearly see that an increased budget allows greater changes to be made to a physicians suggested IV fluid values. For example, with budget of 0.3, the average suggested change for D5LR for positive and negative predicted cases is 0.15 and 0.17, respectively. At a budget of 0.9 these values are 0.47 and 0.51, respectively. These figures also provide the following insights:

1. On average, our model suggests increasing the intake of D5LR, D5HNS and D5W, while recommending that LR and NS be decreased. These observations are in alignment with the literature (Huang et al., 1995) that suggests that resuscitation using LR is associated with increased renal failure.
2. Resuscitation using 5% dextrose in lactate ringer is encouraged among septic patients in ICUs to improve the probability of survival.

7. Conclusions

This study proposes a clinical prescriptive model with human-in-the-loop functionality that recommends optimal, individual-specific amounts of IV fluids for the treatment of septic patients in ICUs. The proposed methodology combines constrained optimization and machine learning techniques to arrive at optimal solutions. A key novelty of the proposed clinical model is utilization of a physician's input to derive optimal solutions. The efficacy of the method is demonstrated using a real world medical dataset. We further validated the robustness of the

proposed approach to show that our method benefits from the human-in-the-loop component, but is also robust to poor input, which is a crucial consideration for new physicians. The results showed, under the limited budget, the optimal solution can improve the average relative probability of survival by 22%. The proposed method can potentially be embedded in an existing electronic health record system to make life-saving IV fluid recommendations. This model can also be used for training junior physicians to synthesize the appropriate treatment strategy, and prevent user error after the inclusion of additional clinical variables and prospective validation. An important limitation of the model is the non-inclusion of vasopressors and antibiotics, which are two important classes of drugs used to treat sepsis.

There are numerous directions for future work. First, we intend to extend this work to capture the temporal dynamics, typically found in clinical data, through the use of longitudinal machine learning approaches. A longitudinal approach will allow us to capture the sequential treatment decision-making process employed by a physician and subsequently prescribe a course of treatment taking such decisions into account. Second, while the assumption of classifier differentiability is not definitively prohibitive, inclusion and exploration of optimization methods that operate on non-differentiable classifications functions, such as random forests, will be beneficial. Finally, derivation and exploration of a data-driven cost function (i.e., learned cost function) will provide added benefit to our formulation, as currently such costs must be specified by the user.

CRedit authorship contribution statement

Akash Gupta: Conceptualization, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing - original draft, Visualization, Project administration. **Michael T. Lash:** Methodology, Software, Validation, Formal analysis, Investigation, Resources, Writing - original draft, Visualization. **Senthil K. Nachimuthu:** Conceptualization, Validation, Writing - review & editing, Supervision.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

The authors would like to thank Dr. Rebekah Child, Ph.D., RN, Associate Professor, California State University – Northridge for initially discussing the medical importance of the problem.

References

- Aggarwal, C. C., Chen, C., & Han, J. (2010). The inverse classification problem. *Journal of Computer Science and Technology*, 25(May), 458–468. <https://doi.org/10.1007/s11390-010-9337-x>
- Amershi, S., Cakmak, M., Knox, W. B., & Kulesza, T. (2014). Power to the people: The role of humans in interactive machine learning. *Ai Magazine*, 35(4), 105–120.
- Barbella, D., Benzaïd, S., Christensen, J., Jackson, B., Qin, X. V., & Muscant, D. (2009). Understanding support vector machine classifications via a recommender system-like approach. In *Proceedings of the international conference on data mining* (pp. 305–311).
- Berg, R. V. D., Kipf, T. N. & Welling, M. (2017). Graph convolutional matrix completion. arXiv preprint arXiv:1706.02263.
- Burke, R. (2002). Hybrid recommender systems: Survey and experiments. *User Modeling and User-Adapted Interaction*, 12(4), 331–370.
- Buuren, S. V., & Groothuis-Oudshoorn, K. (2010). mice: Multivariate imputation by chained equations in R. *Journal of Statistical Software*, 1–68.
- Callaghan, A., Kinsman, L., Cooper, S., & Radomski, N. (2017). The factors that influence junior doctors' capacity to recognise, respond and manage patient deterioration in an acute ward setting: an integrative review. *Australian Critical Care*, 30(4), 197–209.
- Cecconi, M., Evans, L., Levy, M., & Rhodes, A. (2018). Sepsis and septic shock. *The Lancet*, 392(10141), 75–87.
- Chang, R., & Holcomb, J. B. (2016). Choice of fluid therapy in the initial management of sepsis, severe sepsis, and septic shock. *Shock (Augusta, Ga.)*, 46(1), 17.

- Chen, C., Zhao, P., Li, L., Zhou, J., Li, X., & Qiu, M. (2017). Locally connected deep learning framework for industrial-scale recommender systems. In *Proceedings of the 26th international conference on world wide web companion* (pp. 769–770).
- Chi, C. L., Street, W. N., Robinson, J. G., & Crawford, M. A. (2012). Individualized patient-centered lifestyle recommendations: An expert system for communicating patient specific cardiovascular risk information and prioritizing lifestyle options. *Journal of Biomedical Informatics*, 45(6), 1164–1174.
- Courtney, M., Gopinath, B., Toward, M., Jain, R., & Rao, M. (2014). Are adequate fluid challenges prescribed for severe sepsis? *International Journal of Health Care Quality Assurance*.
- Dellinger, R. P., Levy, M. M., Rhodes, A., Annane, D., Gerlach, H., Opal, S. M., ... et al. (2013). Surviving sepsis campaign: international guidelines for management of severe sepsis and septic shock, 2012. *Intensive Care Medicine*, 39(2), 165–228.
- Duhaime, E. P. (2016). Combining human and artificial intelligence for analyzing health data. In 2016 aaai spring symposium series.
- Duraiaraj, L., & Schmidt, G. A. (2008). Fluid therapy in resuscitated sepsis: Less is more. *Chest*, 133(1), 252–263.
- Fay, K., Sapiano, M. R., Gokhale, R., Dantes, R., Thompson, N., Katz, D. E., ... et al. (2020). Assessment of health care exposures and outcomes in adult patients with sepsis and septic shock. *JAMA Network Open* 3(7) e206004–e206004.
- Ferreira, F. L., Bota, D. P., Bross, A., Mélot, C., & Vincent, J.-L. (2001). Serial evaluation of the sofa score to predict outcome in critically ill patients. *JAMA*, 286(14), 1754–1758.
- Frank, E., Hall, M., & Witten, I. (2016). *Data mining: Practical machine learning tools and techniques*. The WEKA Workbench.
- Gil, Y., Honaker, J., Gupta, S., Ma, Y., D'Orazio, V., Garijo, D., & Jahanshad, N. (2019). Towards human-guided machine learning. In *Proceedings of the 24th international conference on intelligent user interfaces* (pp. 614–624).
- Gupta, A., Liu, T., & Crick, C. (2020). Utilizing time series data embedded in electronic health records to develop continuous mortality risk prediction models using hidden markov models: A sepsis case study. *Statistical Methods in Medical Research*, 0962280220929045.
- Gupta, A., Liu, T., & Shepherd, S. (2019). Clinical decision support system to assess the risk of sepsis using tree augmented bayesian networks and electronic medical record data. *Health Informatics Journal*, 1460458219852872.
- Gupta, A., Liu, T., Shepherd, S., & Paiva, W. (2018). Using statistical and machine learning methods to evaluate the prognostic accuracy of sirs and qsofa. *Healthcare Informatics Research*, 24(2), 139–147.
- Henry, K. E., Hager, D. N., Pronovost, P. J. & Saria, S. (2015). A targeted real-time early warning score (trewscore) for septic shock. *Science Translational Medicine* 7(299) 299ra122–299ra122.
- Holzinger, A. (2016). Interactive machine learning for health informatics: when do we need the human-in-the-loop? *Brain Informatics*, 3(2), 119–131.
- Holzinger, A., Lings, G., Denk, H., Zatloukal, K., & Müller, H. (2019). Causability and explainability of artificial intelligence in medicine. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 9(4). e1312.
- Holzinger, A., Plass, M., Holzinger, K., Crisgan, G. C., Pintea, C.-M. & Palade, V. (2016). Towards interactive machine learning (iml): applying ant colony algorithms to solve the traveling salesman problem with the human-in-the-loop approach. In *International conference on availability, reliability, and security* (pp. 81–95).
- Hosmer, D. W., Jr, Lemeshow, S., & Sturdivant, R. X. (2013). *Applied logistic regression* (vol. 398). John Wiley & Sons.
- Huang, P. P., Stucky, F. S., Dimick, A. R., Treat, R. C., Bessey, P. Q., & Rue, L. W. (1995). Hypertonic sodium resuscitation is associated with renal failure and death. *Annals of Surgery*, 221(5), 543.
- Johnson, A. E., Pollard, T. J., Shen, L., Li-wei, H. L., Feng, M., Ghassemi, M., & Mark, R. G. (2016). Mimic-iii, a freely accessible critical care database. *Scientific Data*, 3, 160035.
- Kim, D., Park, C., Oh, J., Lee, S., & Yu, H. (2016). Convolutional matrix factorization for document context-aware recommendation. In *Proceedings of the 10th ACM conference on recommender systems* (pp. 233–240).
- Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. In *3rd international conference on learning representations (iclr 2015)*. Retrieved from <http://arxiv.org/abs/1412.6980>.
- Ko, Y. J., Maystre, L. & Grossglauser, M. (2016). Collaborative recurrent neural networks for dynamic recommender systems. In *Journal of machine learning research: Workshop and conference proceedings* (Vol. 63).
- Komorowski, M., Celi, L. A., Badawi, O., Gordon, A. C., & Faisal, A. A. (2018). The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nature Medicine*, 24(11), 1716–1720.
- Lash, M. T., Lin, Q., Street, W. N., Robinson, J. G., & Ohlmann, J. (2017). Generalized inverse classification. In *Proceedings of the 2017 SIAM international conference on data mining* (pp. 162–170). SIAM.
- Lash, M. T., Lin, Q., Street, W. N., & Robinson, J. G. (2017). A budget-constrained inverse classification framework for smooth classifiers. In *Data Mining Workshops (ICDMW), 2017 IEEE International Conference on* (pp. 1184–1193). IEEE.
- Lash, M. T., & Street, W. N. (2020). Personalized Cardiovascular Disease Risk Mitigation via Longitudinal Inverse Classification. *Bioinformatics and Biomedicine Workshops, 2020 IEEE International Conference on*, 2610–2617.
- Laugel, T., Lesot, M.-J., Marsala, C., Renard, X. & Detyniecki, M. (2018). Comparison-based inverse classification for interpretability in machine learning. In *International conference on information processing and management of uncertainty in knowledge-based systems* (pp. 100–111).
- Lee, S. J., Ramar, K., Park, J. G., Gajic, O., Li, G., & Kashyap, R. (2014). Increased fluid administration in the first three hours of sepsis resuscitation is associated with reduced mortality: A retrospective cohort study. *Chest*, 146(4), 908–915.

- Lops, P., De Gemmis, M., & Semeraro, G. (2011). Content-based recommender systems: State of the art and trends. In *Recommender systems handbook* (pp. 73–105). Springer.
- Mannino, M. V., & Koushik, M. V. (2000). The cost minimizing inverse classification problem: A algorithm approach. *Decision Support Systems*, 29(3), 283–300. [https://doi.org/10.1016/S0167-9236\(00\)00077-4](https://doi.org/10.1016/S0167-9236(00)00077-4)
- Nesterov, Y. (2013). Gradient methods for minimizing composite objective function. *Mathematical Programming, Series B*, 140, 125–161.
- Pendharkar, P. C. (2002). A potential use of data envelopment analysis for the inverse classification problem. *Omega*, 30(3), 243–248.
- Perel, P., & Roberts, I. (2012). Colloids versus crystalloids for fluid resuscitation in critically ill patients. *Cochrane Database of Systematic Reviews*, 6.
- Platt, J., et al. (1999). Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in Large Margin Classifiers*, 10(3), 61–74.
- Pulcini, C., Williams, F., Molinari, N., Davey, P., & Nathwani, D. (2011). Junior doctors' knowledge and perceptions of antibiotic resistance and prescribing: a survey in France and Scotland. *Clinical Microbiology and Infection*, 17(1), 80–87.
- Raghu, A., Komorowski, M., Celi, L. A., Szolovits, P., & Ghassemi, M. (2017). Continuous state-space models for optimal sepsis treatment—a deep reinforcement learning approach. arXiv preprint arXiv:1705.08422.
- Reinhart, K., Daniels, R., Kissoon, N., Machado, F. R., Schachter, R. D., & Finfer, S. (2017). Recognizing sepsis as a global health priority—a who resolution. *New England Journal of Medicine*, 377(5), 414–417.
- Rhee, C., Wang, R., Zhang, Z., Fram, D., Kadri, S. S., Klompas, M., et al. (2019). Epidemiology of hospital-onset versus community-onset sepsis in us hospitals and association with mortality: A retrospective analysis using electronic clinical data. *Critical Care Medicine*, 47(9), 1169–1176.
- Rhodes, A., Evans, L. E., Alhazzani, W., Levy, M. M., Antonelli, M., Ferrer, R., et al. (2017). Surviving sepsis campaign: international guidelines for management of sepsis and septic shock: 2016. *Intensive Care Medicine*, 43(3), 304–377.
- Rochwerg, B., Alhazzani, W., Sindi, A., Heels-Ansdell, D., Thabane, L., Fox-Robichaud, A., et al. (2014). Fluid resuscitation in sepsis: A systematic review and network meta-analysis. *Annals of Internal Medicine*, 161(5), 347–355.
- Rudd, K. E., Kissoon, N., Limmathurotsakul, D., Bory, S., Mutahunga, B., Seymour, C. W., & West, T. E. (2018). The global burden of sepsis: Barriers and potential solutions. *Critical Care*, 22(1), 1–11.
- Sherwin, R., Winters, M. E., Vilke, G. M., & Wardi, G. (2017). Does early and appropriate antibiotic administration improve mortality in emergency department patients with severe sepsis or septic shock? *The Journal of Emergency Medicine*, 53(4), 588–595.
- Su, X., & Khoshgoftaar, T. M. (2009). A survey of collaborative filtering techniques. *Advances in Artificial Intelligence*, 2009.
- Suzuki, K. (2013). *Artificial neural networks: Architectures and applications*. BoD—Books on Demand.
- Vincent, J. L., & Gottin, L. (2011). Type of fluid in severe sepsis and septic shock. *Minerva Anestesiologica*, 77(12), 1190–1196.
- Vincent, J. L., Moreno, R., Takala, J., Willatts, S., De Mendonça, A., Bruining, H., ...Thijs, L. G. (1996). The sofa (sepsis-related organ failure assessment) score to describe organ dysfunction/failure. Springer-Verlag.
- Yang, C., Street, W. N., & Robinson, J. G. (2012). 10-year CVD risk prediction and minimization via inverse classification. In *Proceedings of the 2nd acm sighth symposium on international health informatics – ihi '12* (pp. 603–610).
- Zhang, S., Yao, L., Sun, A., & Tay, Y. (2019). Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys (CSUR)*, 52(1), 5.