

# Learning (predictive) risk scores in the presence of censoring due to interventions

Kirill Dyagilev<sup>1</sup> · Suchi Saria<sup>1,2</sup>

Received: 1 February 2015 / Accepted: 5 August 2015 / Published online: 20 October 2015  
© The Author(s) 2015

**Abstract** A large and diverse set of measurements are regularly collected during a patient’s hospital stay to monitor their health status. Tools for integrating these measurements into severity scores, that accurately track changes in illness severity, can improve clinicians’ ability to provide timely interventions. Existing approaches for creating such scores either (1) rely on experts to fully specify the severity score, (2) infer a score using detailed models of disease progression, or (3) train a predictive score, using supervised learning, by regressing against a surrogate marker of severity such as the presence of downstream adverse events. The first approach does not extend to diseases where an accurate score cannot be elicited from experts. The second assumes that the progression of disease can be accurately modeled, limiting its application to populations with simple, well-understood disease dynamics. The third approach, also most commonly used, often produces scores that suffer from bias due to treatment-related censoring (Paxton et al. in AMIA annual symposium proceedings, American Medical Informatics Association, p 1109, 2013). Specifically, since the downstream outcomes used for their training are observed only noisily and are influenced by treatment administration patterns, these scores do not generalize well when treatment administration patterns change. We propose a novel ranking based framework for disease severity score learning (DSSL). DSSL exploits the following key observation: while it is challenging for experts to quantify the disease severity at any given time, it is often easy to compare the disease severity at two different times. Extending existing ranking algorithms, DSSL learns a function that maps a vector of patient’s measurements to a scalar severity score subject to two

---

Editors: Byron C. Wallace and Jenna Wiens.

---

✉ Kirill Dyagilev  
kirilld@jhu.edu

Suchi Saria  
ssaria@cs.jhu.edu

<sup>1</sup> Department of Computer Science, Johns Hopkins University, 3400 N. Charles St., Baltimore, MD 21218, USA

<sup>2</sup> Department of Health Policy and Management, Johns Hopkins University, 3400 N. Charles St., Baltimore, MD 21218, USA

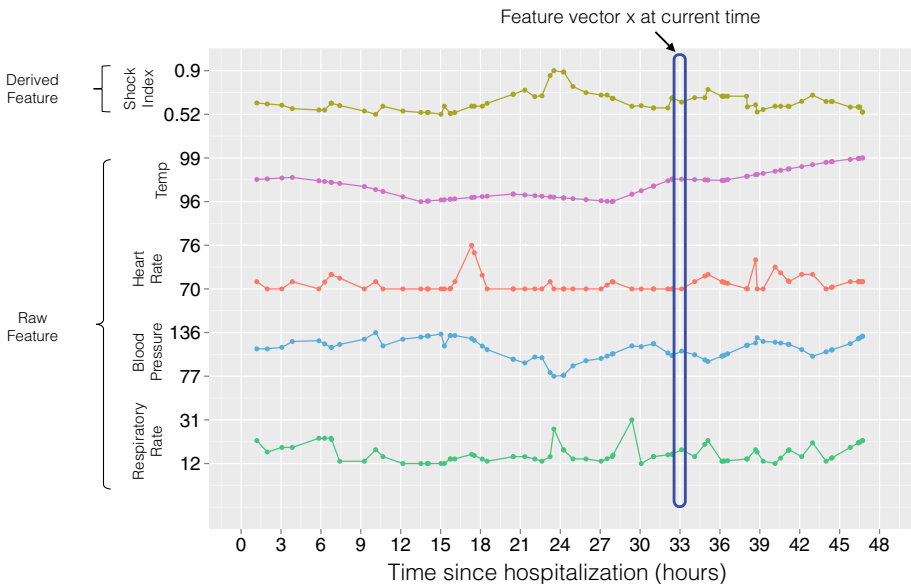
constraints. First, the resulting score should be consistent with the expert’s ranking of the disease severity state. Second, changes in score between consecutive periods should be smooth. We apply DSSL to the problem of learning a sepsis severity score using a large, real-world electronic health record dataset. The learned scores significantly outperform state-of-the-art clinical scores in ranking patient states by severity and in early detection of downstream adverse events. We also show that the learned disease severity trajectories are consistent with clinical expectations of disease evolution. Further, we simulate datasets containing different treatment administration patterns and show that DSSL shows better generalization performance to changes in treatment patterns compared to the above approaches.

**Keywords** MIMIC · Severity score · Sepsis · Gradient boosted regression trees · Ranking

### 1 Introduction

Consider the task of monitoring patients admitted to the Intensive Care Unit (ICU). Clinicians must regularly assess for changes in disease severity to plan timely interventions. Since direct observation of a patient’s disease state is rarely possible, assessing severity requires the caregiver to interpret a diverse array of markers (e.g., heart rate, respiratory rate, blood counts, and serum measurements) that measure the underlying physiologic and metabolic state. In Fig. 1, we show a subset of such data collected on a single patient in the intensive care unit over the 48-h period preceding when they experienced septic shock. Continuous assessments of whether an individual is at-risk based on this data is both time-consuming and challenging.

In this paper, we address the problem of quantifying (scoring) the latent severity of an individual’s disease at a given time. That is, we derive a mapping from the high-dimensional



**Fig. 1** Measurements over time for an example patient in an intensive care unit (ICU). In blue, we identify the feature vector  $x^p$  at time  $t = 33$  h, i.e., all available measurements for a patient  $p$  at time  $t$  (Color figure online)

observed marker data to a numeric score that tracks changes in severity of the underlying disease state over time—as health worsens, the score increases, and as the individual’s health improves, the score declines. Accurate estimation and tracking of the underlying disease severity can enable clinicians to detect critical decline such as decompensations, and acute adverse events in a timely manner. Additional benefits of accurate disease severity estimation include a means for measuring an individual’s response to therapy and stratification of patients for resource management and clinical research (Keegan et al. 2011).

*Defining a Disease Severity Score* Qualitatively, the concept of a disease severity score has been described as the total effect of disease on the body; the irreversible effect is referred to as damage, while the reversible component is referred to as activity (Medsger et al. 2003). The precise interpretation of concepts of damage and activity are typically based on the application at hand. Desirable properties of a severity scale include: (1) face and content validity i.e., the variables included are important and clinically credible, and (2) construct validity i.e., the scoring system parallels an independently ascertained severity measurement (Medsger et al. 2003).

*Prior art* Historically, severity scores have been designed in a number of different ways (Ghanem-Zoubi et al. 2011). One approach is to have clinical experts fully specify the score. Namely, using existing clinical literature, a panel of experts identifies factors that are most indicative of severity of the target disease. These factors are weighted by their relative contribution to the severity and summed together to yield the total resulting score. For example, the Acute Physiology And Chronic Health conditions score (APACHE II) (Knaus et al. 1985), which assesses the overall health state in an in-patient setting, uses factors that are most predictive of mortality. A heart rate between 110 and 139 beats per minute adds 2 points to the final score while a heart rate higher than 180 beats per minute adds 4 points. Similarly, mean arterial blood pressure between 70 and 109 mmHg adds no points while a value between 50 and 69 mmHg adds 2 points. A number of additional widely used scoring systems have been designed in this way, including the Multiple Organ Dysfunction Score (MODS) (Marshall et al. 1995), the Sequential Organ Failure Assessment (SOFA) (Vincent et al. 1996), and Medsger’s scoring system (Medsger et al. 2003).

A second approach commonly taken is to assume that the severity can be characterized in terms of another surrogate measure such as the risk of an impending adverse event or mortality. This method relies on the intuition that high severity states are more likely to be associated with adverse events and higher mortality rates. The disease severity score is then learned by regressing a mapping between observed biomarkers and elements of clinical history and the risk. For instance, the pneumonia severity index (PSI) combines 19 factors including age, vitals and laboratory test results, to calculate the probability of morbidity and mortality among patients with community acquired pneumonia (Fine et al. 1997). The relative weight of each factor in the resulting score was derived by training a logistic regression predictor of patient’s death in the following 30-h window. For simplicity of use, the relative weights were normalized so that the weight of the age would be equal to one and rounded up to the closest multiple of 10 (of 15 for temperature). Others have similarly used downstream adverse events such as the development of *Clostridium difficile* infection (Wiens et al. 2012), septic shock (Ho et al. 2012), morbidity (Saria et al. 2010b), and mortality (Pirracchio et al. 2015) as surrogate sources of supervision for training severity scores.

A third approach uses probabilistic state estimation techniques to track disease severity and progression (e.g., Mould 2012; Jackson et al. 2003; Saria et al. 2010a; Wang et al. 2014). These model disease progression as a function of the observed measurements. For example, Jackson et al. (2003) study abdominal aortic aneurysms in elderly men. They divide the progression of this disease into discrete stages of increasing severity according to successive ranges of

aortic diameter. The disease dynamics is modeled using a hidden Markov model (HMM), which allows to capture both the transition between the stages and the stage misclassification probability. The parameters of this model are estimated using maximum likelihood. Once model parameters are known, the disease severity of a patient (the unobserved state of the HMM) at a given time can be obtained by inference on the learned model.

However, all of the above-mentioned approaches for derivation of disease severity scores have their limitations. The expert-based approach captures known clinical expertise well, but does not extend to populations where the current clinical knowledge is incomplete. The progression modeling based approaches require making assumptions about the disease dynamics and are therefore only applicable to diseases where the dynamics are relatively well-understood. Finally, the third approach, also the most commonly used, often produces scores that suffer from bias due to treatment-related censoring (Paxton et al. 2013). To see why, note that for model training, supervised examples are obtained by annotating each patient's record as a positive or negative training example depending on whether they experienced the target outcome or not (Fine et al. 1997). However, a high-risk patient, if treated in a timely manner, may not experience the adverse event. If there is a group of such patients, who are consistently treated and therefore never experience the adverse event, the learning algorithm will consider their symptoms preceding their treatment as low-risk states, and give it a low severity score. This poses a problem when this severity score is moved to a different environment where treatment decisions are made based on the score alone. A caregiver may choose not to treat these high-risk state because of their low score, thereby worsening outcomes. We elaborate on this issue further with the *SyntheticFlu* example in Sect. 3.1. Accounting for the effects of treatments on the downstream outcome is one way to circumvent this issue (e.g., Henry et al. 2015). In this paper we propose an alternative framework.

*Our contribution* We propose the Disease Severity Score Learning (DSSL) framework that exploits this key observation that, while requesting experts to quantify disease severity at a given time is challenging, acquiring *clinical comparisons*—clinical assessments that order the disease severity at two different times—is often easy. These clinical comparisons, compared to labels based on downstream adverse events, are also less sensitive to treatment patterns. Further, in the majority of diseases, clinical guidelines provide rules for coarse-grained assessment of stages of a disease (see examples in AHRQ 2015). These stages can be used to augment expert-provided clinical comparisons with those that are automatically generated using these guidelines. We show how we leverage an existing guideline (Dellinger et al. 2013) in our example application.

DSSL uses clinical comparisons within the same patient and across patients to train a temporally smooth disease severity score. From these clinical comparisons, DSSL learns a function that maps the patients observed feature vectors to a scalar severity score. With some abuse of terminology, we refer to this mapping function as the disease severity score (DSS) function. We present two different algorithms for learning the DSS—the first in the linear setting, and the second in the non-linear setting. In both cases, the parameters of the DSS function are found by optimizing an objective function that contains two key terms. The first term penalizes for pairs that are incorrectly ordered by their severity. The second term imposes a penalty on changes of the severity score that are driven by the temporal evolution of the disease. For example, in our application, sepsis evolves slowly and the learning objective leverages this by penalizing scores that are not smooth over those that are. We show how two commonly used ranking algorithms can be extended to our problem in a relatively straightforward manner. For the linear DSS, we extend the soft max-margin formulation by Joachims (2002) to maximize separation between ordered pairs while preserving temporal smoothness. For the non-linear DSS, the score is represented non-parametrically using a

weighted sum of regression trees. We use an optimization procedure similar to that of gradient boosted regression (Mason et al. 1999; Friedman 2001) to obtain the DSS function parameters. We show numerical results on the task of training a sepsis severity score for patients in the ICU.

Below, we highlight the main strengths of the proposed DSS learning framework:

1. Our learning algorithm provides a scalable and automatic approach to learning disease severity scores in new disease domains and populations.
2. Our learning algorithm only requires a means for obtaining *clinical comparisons*—ordered pairs comparing disease severity state at different times. This form of supervision is more natural to elicit than asking clinical experts to map the disease severity score, or encoding an accurate model of disease progression. Moreover, this supervision can often be generated automatically. Our approach allows experts to tune the quality of the score by increasing the granularity and amount of supervision given.
3. We show that our algorithm learns scores that are consistent with clinical expectations. For example, changes in the severity score over consecutive time periods are smooth and the score is higher in periods adjacent to an adverse event. Additionally, the score is sensitive to changes in disease severity state due to therapies.

## 2 The disease severity score learning (DSSL) framework

In this section, we introduce our methodology for learning DSS functions. We begin by outlining the general framework for learning a temporally smooth disease severity score in Sect. 2.1. Section 2.2 presents the soft-margin approach for learning a linear DSS function. In Sect. 2.3, we extend our methodology to non-linear DSS functions using gradient boosted regression trees.

### 2.1 Overview

We consider data that are routinely collected in a hospital setting. These include covariates such as age, gender, and clinical history (e.g., presence or absence of a clinical condition such as AIDS or Diabetes) obtained at the time of admission; time-varying measurements such as heart rate, respiratory rate, urine volume obtained throughout the length of stay; and text notes summarizing the patients evolving health status. These data are processed and transformed into tuples  $\langle x_i^p, t_i^p \rangle$  where  $\mathbf{x}_i^p \in \mathcal{R}^d$  is a  $d$ -dimensional feature vector associated with patient  $p \in P$  at time  $t_i^p$  for  $i \in \{1, \dots, T^p\}$  and  $T^p$  is the total number of tuples for patient  $p$ . A feature vector  $\mathbf{x}_i^p$  contains raw measurements (e.g., last measured heart rate or last measured white blood cell count) and features derived from one or more measurements (e.g., the mean and variance of the measured heart rate over the last 6 h or the total urine output in the last 6 h per kilogram of weight). In Fig. 1 in Sect. 1, we showed example components of a feature vector computed for a patient in the intensive care unit over a 48-h period. Let  $D$  denote the set of tuples across all patients in the study.

The problem of learning a DSS function is defined by the sets  $O$  and  $S$  of pairs of tuples from the set  $D$  of all tuples, and by the set  $G$  of permissible DSS functions. The set  $O$  contains pairs of tuples  $(\langle x_i^p, t_i^p \rangle, \langle x_j^q, t_j^q \rangle)$  that are ordered by severity based on clinical assessments. We refer to each of these paired tuples as a *clinical comparison* and the set  $O$  as the set of all available clinical comparisons. For notational simplicity, we assume that  $\mathbf{x}_i^p$  corresponds to a more severe state than  $\mathbf{x}_j^q$ . These clinical comparisons can be obtained by presenting

clinicians with data  $\mathbf{x}_i^p$  for patient  $p \in P$  at time  $t_i^p$  and data  $\mathbf{x}_j^q$  for patient  $q \in P$  at time  $t_j^q$ . For each such pair of feature vectors, the clinical expert identifies which of these correspond to a more severe health state; the expert can choose not to provide a comparison for a pair where the severity ordering is ambiguous. These pairs can also be generated in an automated fashion by leveraging existing clinical guidelines. In Sect. 3.3.1, we describe how we use an existing guideline in our application.

The set  $S$  contains pairs of tuples  $(\langle x_i^p, t_i^p \rangle, \langle x_{i+1}^p, t_{i+1}^p \rangle)$  that correspond to feature vectors that are taken from the same patient  $p$  at consecutive time steps  $t_i^p$  and  $t_{i+1}^p$ . These pairs are used to impose smoothness constraints on the learned severity scores. We thus refer to the pairs in  $S$  as the *smoothness pairs*. Finally, the set  $G$  contains a parameterized family of candidate DSS functions  $g$  that map feature vectors  $\mathbf{x}$  to a scalar severity score.

Our goal is to identify a function  $g \in G$  that quantifies the severity of the disease state represented by a feature vector  $\mathbf{x}$ . In particular, this function should correctly order any pair  $(\mathbf{x}, \mathbf{x}')$  of feature vectors by their severity, and the resulting score should be temporally smooth to mimic the natural inertia exhibited by our biological system. We use empirical risk minimization to identify such a function  $g$ . Namely, we construct an objective function  $C^g$  that maps functions  $g \in G$  to their empirical risk. The first of the two terms in  $C^g$  is

$$\sum_{(\langle x_i^p, t_i^p \rangle, \langle x_{i+1}^p, t_{i+1}^p \rangle) \in S} \left[ \frac{g(\mathbf{x}_{i+1}^p) - g(\mathbf{x}_i^p)}{t_{i+1}^p - t_i^p} \right]^2. \tag{1}$$

This term penalizes DSS functions that exhibit large changes in the severity score over short durations, hence encouraging selection of temporally smooth DSS functions. The second term in  $C^g$  penalizes  $g$  for pairs of tuples  $(\langle x_i^p, t_i^p \rangle, \langle x_j^q, t_j^q \rangle) \in O$  for which the severity ordering induced by  $g$  on vectors  $\mathbf{x}_i^p$  and  $\mathbf{x}_j^q$  is inconsistent with the ground truth clinical assessment. i.e.,  $g(\mathbf{x}_i^p) < g(\mathbf{x}_j^q)$ . We discuss the full objective comprising these two terms in greater detail in Sect. 2.2.


In the following two sections we describe the objectives and corresponding optimization algorithms for learning the linear and non-linear DSS functions in a new disease domain.

### 2.2 Learning a linear DSS

We first consider the problem of learning linear DSS functions, i.e., DSS functions of the form  $g_{\mathbf{w}}(\mathbf{x}) = \mathbf{w}^\top \mathbf{x}$ . We refer to the corresponding learning procedure as L-DSS.

We employ soft max-margin training (Joachims 2002) where we seek to maximize the distance between the pairs that are at different severity levels while keeping the distance between the consecutive pairs smooth. We briefly review the key concepts of soft max-margin ranking before we describe our extension for learning a linear DSS function given data.

*Soft max-margin ranking* Consider the toy example shown in Fig. 2. Let  $D$  contain the three feature vectors  $\{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}$  where  $\mathbf{x}_i \in \mathcal{R}^2$ , and  $O$  contain the pairs  $(\mathbf{x}_2, \mathbf{x}_1)$  and  $(\mathbf{x}_3, \mathbf{x}_2)$ , i.e., feature vectors  $\mathbf{x}_2$  and  $\mathbf{x}_3$  have higher disease severity than  $\mathbf{x}_1$  and  $\mathbf{x}_2$  respectively. Max-margin ranking seeks to find a vector  $\mathbf{w}$  such that the *margin* between pairs of different severity levels is maximized. In our example, we show parameter vectors  $\mathbf{w}_1, \mathbf{w}_2$  and  $\mathbf{w}_3$  for three candidate ranking functions in Fig. 2. For each feature vector  $\mathbf{x}$ , the assigned (severity) score for a given ranking function parameter  $\mathbf{w}_i$  is computed as the projection,  $g_{\mathbf{w}_i}(\mathbf{x})$ , of  $\mathbf{x}$  on  $\mathbf{w}_i$ . The induced ranking between two vectors  $\mathbf{x}_1$  and  $\mathbf{x}_2$  is computed based on the margin which is defined as the difference in their projections. In the example shown, the rankings



**Fig. 2** We show the projections of  $\mathbf{x}_1$ ,  $\mathbf{x}_2$  and  $\mathbf{x}_3$  on vectors  $\mathbf{w}_1$  and  $\mathbf{w}_2$  representing two candidate ranking functions; vectors drawn in red, green and blue identify projections of  $\mathbf{x}_1$ ,  $\mathbf{x}_2$  and  $\mathbf{x}_3$  respectively. Ranking is induced by the differences in the projections—for example,  $\mathbf{w}_2$  induces the ordering  $g_{\mathbf{w}_2}(\mathbf{x}_1) > g_{\mathbf{w}_2}(\mathbf{x}_2)$  because  $g_{\mathbf{w}_2}(\mathbf{x}_1) - g_{\mathbf{w}_2}(\mathbf{x}_2) > 0$

induced by both  $g_{\mathbf{w}_1}$  and  $g_{\mathbf{w}_3}$  correctly order all pairs in  $O$ , i.e.,

$$g_{\mathbf{w}_1}(\mathbf{x}_3) > g_{\mathbf{w}_1}(\mathbf{x}_2) > g_{\mathbf{w}_1}(\mathbf{x}_1) \text{ and } g_{\mathbf{w}_3}(\mathbf{x}_3) > g_{\mathbf{w}_3}(\mathbf{x}_2) > g_{\mathbf{w}_3}(\mathbf{x}_1),$$

while the rankings induced by  $\mathbf{w}_2$  do not. Furthermore,  $\mathbf{w}_3$  also induces an ordering with a larger margin between the pairs in  $O$ . Margin-maximization leads to an ordering that is more robust with respect to noise in  $\mathbf{x}$ .

More formally, for each pair of feature vectors  $(\mathbf{x}_i, \mathbf{x}_j) \in O$ , we define the margin of their separation by the function  $g_{\mathbf{w}}(\cdot)$  as  $\mu_{i,j}^{\mathbf{w}} = g_{\mathbf{w}}(\mathbf{x}_i) - g_{\mathbf{w}}(\mathbf{x}_j)$ . The maximum-margin approach suggests that we can improve generalization and robustness of the learned separator by selecting  $\mathbf{w}$  that maximizes the number of tuples that are ordered correctly (i.e.,  $\mu_{i,j}^{\mathbf{w}} > 0$ ) while simultaneously maximizing the minimal normalized margin  $\mu_{i,j}^{\mathbf{w}} / \|\mathbf{w}\|$ . Using the standard soft max-margin framework, the SVMRank algorithm (Joachims 2002) approximates the above-mentioned problem as the following convex optimization program:

$$\min_{\mathbf{w}, \zeta_O^{i,j}} \left[ \frac{1}{2} \|\mathbf{w}\|^2 + \frac{\lambda_O}{|O|} \sum_{(\mathbf{x}_i, \mathbf{x}_j) \in O} \zeta_O^{i,j} \right] \tag{2}$$

subject to the following **ordering constraints** :

$$\forall (\mathbf{x}_i, \mathbf{x}_j) \in O : g_{\mathbf{w}}(\mathbf{x}_i) - g_{\mathbf{w}}(\mathbf{x}_j) \geq 1 - \zeta_O^{i,j} \text{ and } \zeta_O^{i,j} \geq 0$$

*The L-DSS objective and optimization algorithm* We now describe our algorithm for learning linear DSS functions. We return to our original setting where we are given sets  $O$  and  $S$  which contain feature vectors belong to more than one patients at varying times.

We augment the soft-max margin objective with the additional term, shown in Eq. (1), that encourages temporal smoothness. We state the full L-DSS objective below.

$$\min_{\mathbf{w}, \zeta_O^{ij}} \left[ \frac{1}{2} \|\mathbf{w}\|^2 + \frac{\lambda_O}{|O|} \sum_{((x_i^p, t_i^p), (x_j^q, t_j^q)) \in O} \zeta_O^{(p,i),(q,j)} \right. \tag{3}$$

$$\left. + \frac{\lambda_S}{|S|} \sum_{((x_i^p, t_i^p), (x_{i+1}^p, t_{i+1}^p)) \in S} \left[ \frac{g_{\mathbf{w}}(\mathbf{x}_{i+1}^p) - g_{\mathbf{w}}(\mathbf{x}_i^p)}{t_{i+1}^p - t_i^p} \right]^2 \right]$$

subject to the following **ordering constraints** :

$$\forall ((x_i^p, t_i^p), (x_j^q, t_j^q)) \in O : g_{\mathbf{w}}(\mathbf{x}_i^p) - g_{\mathbf{w}}(\mathbf{x}_j^q) \geq 1 - \zeta_O^{(p,i),(q,j)} \quad \text{and}$$

$$\forall ((x_i^p, t_i^p), (x_j^q, t_j^q)) \in O : \zeta_O^{(p,i),(q,j)} \geq 0$$

Here, the coefficients  $\lambda_O$  and  $\lambda_S$  control the relative degree of emphasis on the smoothness versus the margin-maximization component of the objective. For a given setting of  $\lambda_O$ , different choices of  $\lambda_S$  yield trajectories with differing levels of smoothness. An appropriate choice of  $\lambda_S$  could be determined by the clinical user based on the rate of change in severity that is to be expected in that domain. For example, in sepsis, changes in severity do not occur within minutes while in many cardiac conditions, rapid changes in severity can occur. Alternately, this parameter can be set using cross-validation to optimize performance for a particular application of DSS.

In Eq. (3), for every value of  $\mathbf{w}$ , the optimal values of  $\zeta_O^{(p,i),(q,j)}$  are given by

$$\zeta_O^{(p,i),(q,j)} = \max\{0, 1 - (g_{\mathbf{w}}(\mathbf{x}_i^p) - g_{\mathbf{w}}(\mathbf{x}_j^q))\}. \tag{4}$$

Substituting Eq. (4) and  $g_{\mathbf{w}}(\mathbf{x}) = \mathbf{w}^T \mathbf{x}$  in Eq. (3), we obtain the following unconstrained convex optimization formulation:

$$\min_{\mathbf{w}} \frac{1}{2} \|\mathbf{w}\|^2 + \frac{\lambda_O}{|O|} \sum_{((x_i^p, t_i^p), (x_j^q, t_j^q)) \in O} \max\{0, 1 - \mathbf{w}^T(\mathbf{x}_i^p - \mathbf{x}_j^q)\} \tag{5}$$

$$+ \frac{\lambda_S}{|S|} \sum_{((x_i^p, t_i^p), (x_{i+1}^p, t_{i+1}^p)) \in S} \left[ \frac{\mathbf{w}^T(\mathbf{x}_{i+1}^p - \mathbf{x}_i^p)}{t_{i+1}^p - t_i^p} \right]^2$$

Instead of solving the dual formulation as in [Joachims \(2002\)](#), following the reasons of efficiency and accuracy discussed by [Chapelle and Keerthi \(2010\)](#), we solve the primal form of this optimization program as follows.

The terms of the form  $\max\{0, a\}$ , also called the hinge loss, are not differentiable at  $a = 0$ . We approximate these terms with the Huber loss  $L_h$  for  $0 < h < 1$  given by

$$L_h(a) = \begin{cases} 0 & , \text{ if } a < -h \\ \frac{(a+h)^2}{4h} & , \text{ if } |a| \leq h \\ a & , \text{ if } a > h \end{cases}$$



This approximation yields the following unconstrained, convex, twice-differentiable optimization problem:

$$\begin{aligned}
 \min_{\mathbf{w}} \quad & \frac{1}{2} \|\mathbf{w}\|^2 \\
 & + \frac{\lambda_O}{|O|} \sum_{((x_i^p, t_i^p), (x_j^q, t_j^q)) \in O} L_h(1 - \mathbf{w}^\top (\mathbf{x}_i^p - \mathbf{x}_j^q)) \\
 & + \frac{\lambda_S}{|S|} \sum_{((x_i^p, t_i^p), (x_{i+1}^p, t_{i+1}^p)) \in S} \left[ \frac{\mathbf{w}^\top (\mathbf{x}_{i+1}^p - \mathbf{x}_i^p)}{t_{i+1}^p - t_i^p} \right]^2
 \end{aligned} \tag{6}$$

L-DSS Objective :

We solve this optimization program using the Newton–Raphson algorithm. We show experiments using the L-DSS learner in Sect. 3.

### 2.3 Learning a non-linear DSS

In many disease domains, assuming a linear mapping between the measurements and the latent disease severity may be too restrictive. For example, ranges for measurements values that are considered to be *normal* (or from a low-severity state) are often age dependent or clinical history dependent. Consider an individual with a pre-existing kidney condition; he or she is likely to have a worse baseline creatinine level (a test that measures kidney function) compared to an individual with fully-functioning kidneys. Thus, when measuring changes in severity related to the kidney, these individuals are likely to manifest a disease differently. See the guideline by [Dellinger et al. \(2013\)](#) for other examples.

To learn non-linear DSS functions, we represent  $g$  as a weighted sum of regression trees. Alternate choices for learning non-linear DSS functions exist including extending the soft-margin formulation presented for learning L-DSS via use of the “kernel-trick” ([Kuo et al. 2014](#)). We chose to extend boosted regression trees as this is one of the most widely used algorithms for ranking (e.g., see [Mohan et al. 2011](#)).

Our hypothesis class  $G$  includes all linear combinations of shallow regression trees, i.e., functions of the form  $g(\mathbf{x}) = \sum_{k=1}^K \alpha_k f_k(\mathbf{x})$ , where  $f_k$  for  $k = 1, \dots, K$  are shallow (limited-depth) regression trees and  $K$  is finite. In our experiments,  $K$  is set to 5. Similar to the objective for L-DSS in Eq. (6), we construct the NL-DSS objective to identify  $g \in G$  that maximizes the dual criteria of ordering accuracy and temporal smoothness as:

$$\begin{aligned}
 C^g(g) = \frac{1}{|O|} \sum_{((x_i^p, t_i^p), (x_j^q, t_j^q)) \in O} L_h(1 - (g(\mathbf{x}_i^p) - g(\mathbf{x}_j^q))) \\
 + \frac{\lambda_S}{|S|} \sum_{((x_i^p, t_i^p), (x_{i+1}^p, t_{i+1}^p)) \in S} \left[ \frac{g(\mathbf{x}_{i+1}^p) - g(\mathbf{x}_i^p)}{t_{i+1}^p - t_i^p} \right]^2
 \end{aligned} \tag{7}$$

NL-DSS Objective :

Note that since the soft max-margin formulation is not defined for a non-linear classifier we drop the term  $\|\mathbf{w}\|^2 / 2$ . Thus, without loss of generality,  $\lambda_O$  can be replaced by 1. Now, the relative emphasis on the smoothing versus the ordering components are changed by varying  $\lambda_S$ .

We optimize the NL-DSS objective using the gradient boosted regression trees (GBRT) learning algorithm (Friedman 2001; Mason et al. 1999; Burges 2010). Gradient boosting methods grow  $g$  incrementally, in a greedy fashion, by adding a weak learner—in this case, a regression tree—at each iteration. A tree that most closely approximates the gradient of  $C^g$  evaluated at  $g$  obtained in the previous iteration is added (Friedman 2001).

The per-iteration computational complexity of this approach is equivalent to the computational complexity of building a single regression tree, which is  $|T| \log |T|$  (Hothorn et al. 2006), where  $|T|$  is the number of unique tuples in the set  $\mathcal{O} \cup \mathcal{S}$  of tuple pairs.

### 3 Experiments

We now describe the evaluation of the proposed DSSL framework. Before discussing numerical results on a real-world dataset, in Sect. 3.1, we use a simple toy example to illustrate the behavior of DSS related to the following questions. First, when DSS is transported between environments with different degrees of interventional confounds, how is the performance of DSS impacted compared to the performance of a supervised learning algorithm that uses downstream-events as labels? Second, when clinical comparisons are generated by implementing automated coarse-grading rules, does the learned score simply learn the rule itself? In Sect. 3.2, we provide background on our application: we introduce sepsis, the dataset used, and the guideline used for generating the clinical comparisons needed to train the DSS scores. Next, in Sect. 3.3, we provide an overview of the experiments and the experimental setup followed by a detailed discussion of the numerical results on the sepsis data in Sect. 3.3.3.

#### 3.1 Learning DSS for *SyntheticFlu*

For these experiments, we create a simple toy disease called *SyntheticFlu* as follows. We quantify severity as a function of the patient’s temperature and white blood cell counts (WBC)—as the temperature or the WBC increases, risk of mortality increases. We assume that the disease manifests in two ways: with 50% probability, patients are sampled from a model where the temperature tends to *deteriorate* while the WBC remains *normal*, and for the other fraction of the population, their WBC tends to deteriorate while the temperature remains normal. Each of these measurements assume one of 10 states (e.g., the temperature ranges from 99 to 108 °F). In the absence of treatment, for a measurement that is deteriorating, it retains its value  $T$  in the following timestep with probability 0.3, increases to  $T + 1$  with probability 0.5, and decreases to  $T - 1$  with probability 0.2. For a measurement that is assumed to stay normal, the corresponding transition probabilities are 0.7, 0.1 and 0.2. States 1–2 are defined to be “benign” (e.g., temperatures of 99 and 100 °F) where an individual can be discharged with probability 0.5 (i.e. their sampled trajectory ends). States 6–9 are defined to be “severe” (e.g., temperatures between 104 to 107 °F) where an individual may receive treatment with probability  $\rho_T$ . Administration of treatment (e.g., antibiotics) transitions the individual to one of the benign states. Finally, an individual dies when she or he reaches state 10.

*Evaluating transportability* The first question we investigate is regarding the transportability of the different scores. Namely, when DSS is moved between different treatment regimes, how is the performance of DSS impacted compared to the performance of a supervised learning algorithm that uses downstream-events as labels. Since treatments affect the prevalence of adverse outcomes, we show that the risk scores learned via the latter approach are highly


**Table 1** Analysis of transportability of the L-DSS, NL-DSS and logistic regression based severity scores between different treatment regimes.  $\rho_T^{\text{train}}$ ,  $\rho_{\text{WBC}}^{\text{train}}$ ,  $\rho_T^{\text{test}}$ , and  $\rho_{\text{WBC}}^{\text{test}}$  denote the probability of treatment for temperature and WBC in the train and test regimes respectively

Scenario	$\rho_T^{\text{train}}$	$\rho_{\text{WBC}}^{\text{train}}$	$\rho_T^{\text{test}}$	$\rho_{\text{WBC}}^{\text{test}}$	Logistic regression	L-DSS	NL-DSS
#1	0	0	0	0	0.974	0.973	0.974
#2	0.1	0	0.1	0	0.978	0.990	0.991
#3	0.1	0	0	0	0.963	0.974	0.981
#4	0.3	0	0	0	0.769	0.973	0.981
#5	0.3	0	0	0.3	0.510	0.978	0.996

sensitive to treatment patterns and therefore, in our example setting, generalize poorly compared to DSS. We sample 1000 patients each for the train and test sets. Data are sampled from different treatments regimes as shown in Table 1. For example, in scenario 1, no treatments are prescribed in either the train or test regimes. In scenario 5, in the train regime, treatments are prescribed with 0.3 probability only for treating high temperature but not for high WBC. However, in the test regime, treatments are prescribed only for high WBC but not for high temperature. Training L-DSS and NL-DSS requires generating clinical comparisons. To do so, we randomly sample pairs from the observed data. For a pair  $(\mathbf{x}_i^p, \mathbf{x}_j^q)$ , we consider  $\mathbf{x}_i^p$  to represent a more severe state if one of the measurements is at least 2 units higher in  $\mathbf{x}_i^p$  than in  $\mathbf{x}_j^q$  (e.g., temperature of 103 °F is more severe than 101 °F) and the other measurement is at least as high in  $\mathbf{x}_i^p$  as in  $\mathbf{x}_j^q$ . Many such pairs are sampled for the train and test set. We use a standard protocol for training logistic regression (LR) (Paxton et al. 2013). Train and test samples are generated from the patient trajectories via a sliding window approach with the outcome defined as whether or not the patient died within 10 timesteps in the future. On the test set, we consider a patient to be correctly identified as at-risk patient if his severity score was greater than a certain threshold value at any point of patient’s hospital stay. We measure performance of the obtained scores using the area under the curve (AUC) obtained on the task of predicting per-patient mortality (Paxton et al. 2013).

**Results:** In scenario 1 and scenario 2 where the treatment patterns are the same across the train and test regimes, all three scores perform equally well. However, as the treatment patterns begin to diverge to an increasing degree as seen in scenarios 4 and 5, LR’s performance drops while the DSS performance does not change. It is worth noting that such discrepancies between the train and test regimes can occur within the same hospital when comparing treatment practice before and after deploying a predictive model. Specifically, while clinicians continue to treat high-risk states, the resulting LR score learned from this data will underestimate risks for the treated states. Once the decision support system is deployed and clinicians begin to rely on the predictive tool, they may erroneously undertreat high-risk patients, thereby worsening outcomes.

*Relationship of learned scores to the coarse grades* As mentioned, for many diseases, clinical guidelines provide rules for coarse-grained assessment of severity stages of a disease. These guidelines can be used for automated generation of clinical comparison pairs with unambiguous severity ordering. This raises a natural question of whether a DSS learned from such clinical comparisons simply recovers these clinical guidelines and thus yields no generalization beyond the coarse grading. To evaluate this hypothesis, we extend the setup



**Fig. 3** Experiments on synthetic data: relationship of learned scores to the coarse grades. L-DSS and NL-DSS scores for varying values of  $\lambda_S$  for a patient with WBC in state 2 and temperature ranging from 99 to 107 °F. The scores are normalized so that temperature of 99 °F is given the score of 0 and the temperature of 107 °F is given the score of 1

described above. Not sure if the previous sentence gives too strong of a claim. After all, we show a single example where this does not happen. Specifically, we augment the feature vectors to include the coarse grades which are in turn derived from the temperature and WBC measurements. We derive the coarse severity grade from the observed feature vectors as follows. We assign a feature vector  $\mathbf{x}_i^p$  a severity of 0 if both the corresponding WBC and temperature measurements are in states 1 or 2 (e.g., temperature below 101 °F), a severity of 1 if exactly one of the measurements is in state 3 or higher (e.g., temperature above or equal to 101 °F), and a severity of 2 if both measurements are in states 3 or higher. We consider all three combinations of coarse graded severity pairs and randomly sample 6000 clinical comparisons and a similar number of smoothness pairs. We use  $\lambda_O = 100$  for the L-DSS and sweep values  $\lambda_S$  between 0.1 and 1000. We show results using data sampled from regime 1 though the conclusions do not depend on the treatment pattern.

Since perfect ordering accuracy can be achieved using only the coarse grading component of the feature vector, one might expect that the learned scores will rely on the coarse grading feature alone. In fact, this is the case when  $\lambda_S$  is small ( $=0.1$ ). However, such DSS score will exhibit abrupt changes between consecutive time points, e.g., when the temperature or WBC progresses from state 2 to state 3 or vice versa. As  $\lambda_S$  increases, the smoothness term in the DSS objective encourages the learning of temporally smooth scores which rely increasingly on the WBC and temperature features alone. Thus, in this scenario, the smoothness constraint allows the L-DSS and NL-DSS scores to generalize beyond coarse grades. In Fig. 3, we depict the severity score assigned to a feature vector in which WBC is in state 1 and temperature varies from 99 to 107 °F for varying values of  $\lambda_S$ .

### 3.2 Sepsis, MIMIC-II and the surviving sepsis campaign guideline

In the following experiments on the real-world clinical data, our goal is to learn a score that assesses the severity of sepsis.

*Sepsis* Sepsis is a whole-body inflammatory response to infection; it is a leading cause of death in the inpatient setting, with especially high mortality among patients who develop septic shock, a major sepsis-related adverse event. Both sepsis and septic shock are known to be associated with high morbidity, longer hospital stay and increased health care cost (Kumar et al. 2011). Often, the risk of sepsis-related adverse outcomes can be reduced by early treatment (Sebat et al. 2007), thus a scoring system that allows precise tracking of changes in sepsis-related disease severity is of great importance.

*Dataset* We use MIMIC-II, a publicly available dataset containing electronic health record data from patients admitted to the ICUs at the Beth Israel Deaconess Medical Center from 2001 to 2008 (Saeed et al. 2002, 2011). We only include adults (>15 years old) in our study ( $N = 16,234$ ). We compute 45 different features that derived from vital sign measurements, clinical history variables, and laboratory test results. We provide the complete list of features in Appendix 6. We impute missing data using linear interpolation. Other examples of imputation methods used for ICU monitoring datasets include model-based (e.g., Ho et al. 2012) and Last Observation Carried Forward (LOCF) (Hug 2009). These methods require making numerous domain-specific assumptions. For example, in LOCF, each measurement is carried forward for a finite time window which length is determined based on the typical sampling frequency of this measurement. Since this choice is orthogonal to the focus of our paper, we implement linear interpolation as it is the simplest of the above methods.

*Guideline for grading sepsis severity* In order to train a severity score in the DSSL framework, we need a set  $O$  of pairs of feature vectors ordered by their severity. We create these pairs automatically by leveraging the coarse severity grading of sepsis established in the Surviving Sepsis Campaign Guideline (SSCG) (Dellinger et al. 2013). The SSCG provides rules for identifying when an individual is in each of the four stages of severity: septic shock, severe sepsis, SIRS and none. For each of these stages, the guideline defines criteria using (1) a combination of thresholds for individual measurements, and (2) presence of specific diagnosis codes or diagnoses noted in their clinical notes. For example, the SIRS criterion is met when values of at least two out of these five features are out of their normal range. White blood cell count per microliter, for example, is considered to be out of the normal range if its value is either below 4.0 or above 12.0. Heart rate is considered to be out of the normal range if it is above 90 beats-per-minute. The stage of severe sepsis is reached when physician suspects that patient developed an infection, the SIRS criterion is met, and at least one organ system is showing signs of failure. Finally, septic shock is defined as severe sepsis with observed hypotension despite significant fluid resuscitation. In Table 2, we specify the criteria used for grading each of the stages. It is also worth noting that the guideline does not provide a grade at all times: on average, application of these criteria grades less than 40% of data entries of a patient. We use the time points with available SSSG grading to generate clinical comparison as described in Sect. 3.3.1. We also note that the SSSG grades use features beyond those used in the severity score. For example, to assess the severity grade, we use keyword search in transcripts written at the time of discharge to determine whether the patient had developed an infection. Similarly, features such as whether or not the patient received sufficient fluids are used for grading but excluded when learning the severity scores as these are caregiver driven and only indirectly measure severity.

### 3.3 Learning DSS for sepsis in ICU patients

We begin with an overview of the experiments. In our first experiment, we assess the quality of the trained L-DSS and NL-DSS scores by their performance on the task of distinguishing between the different severity stages of sepsis. This is done by calculating severity ordering of held out pairs of feature vectors and measuring their concordance with the ground truth provided by the SSCG guideline. We show that our scores significantly outperform the three main scoring systems that are widely used in ICUs.

The success of our scores in distinguishing between sepsis stages is encouraging, but expected, since our scores are explicitly trained for severity ordering. In the following experiments, we evaluate whether the learned scores also generalize well to measuring fine-grained changes in severity. Towards this, we first examine whether DSS is sensitive to changes in

**Table 2** Definition of the Surviving Sepsis Campaign Guideline (SSCG) for sepsis severity grading

Stage	Criteria
SIRS	At least two out of the following four conditions hold: <ol style="list-style-type: none"> <li>1. Heart rate is <math>&gt;90</math> beats-per-minute and was measured in the last 2 h</li> <li>2. Temperature is either <math>&gt;38</math> or <math>&lt;36^{\circ}\text{C}</math> and was measured in the last 8 h</li> <li>3. Respiratory rate is <math>&gt;20</math> beats-per-minute and was measured in the last 2 h, or arterial partial pressure of <math>\text{CO}_2</math> is <math>&lt;32</math> mmHg and was measured in the last 8 h</li> <li>4. White blood cell count in thousands per microliter is either <math>&gt;12.0</math> or <math>&lt;4.0</math> and was measured in the last 8 h</li> </ol>
Severe Sepsis	Patient's clinical record contain words sepsis or septic or ICD-9 code for infection SIRS criteria holds, and at least one of the following nine criteria holds: <ol style="list-style-type: none"> <li>1. Systolic blood pressure is <math>&lt;90</math> mmHg and was measured in the last 2 h</li> <li>2. Blood lactate measurement is <math>&gt;2.0</math> micromol per liter and was taken in the last 2 h</li> <li>3. Urine output over the past 2 h is <math>&lt;0.5</math> milliliter per kg</li> <li>4. Patient has no chronic renal insufficiency and her or his creatinine measurement is <math>&gt;2.0</math> milligrams per deciliter, and was taken in the last 8 h</li> <li>5. Patient has no chronic liver disease, her or his bilirubin measurement is <math>&gt;2.0</math> mg per deciliter and was taken in the last 8 h</li> <li>6. Platelet count is <math>&lt;100,000</math> per microliter and was measured in the last 8 h</li> <li>7. International normalized ratio (INR) is <math>&gt;1.5</math> and was measured in the last 8 h</li> <li>8. Patient experienced pneumonia during their hospital stay as indicated by ICD-9 codes, measurements of both partial arterial pressure of oxygen (<math>\text{PaO}_2</math>) and of fraction of inspired oxygen (<math>\text{FiO}_2</math>) were taken in the last 8 h, and it holds that <math>\text{PaO}_2/\text{FiO}_2 &lt; 200</math></li> <li>9. Patient experienced acute lung infection unrelated to pneumonia during their hospital stay as indicated by ICD-9 codes, the measurements of <math>\text{PaO}_2</math> and <math>\text{FiO}_2</math> were taken in the last 8 h, and it holds that <math>\text{PaO}_2/\text{FiO}_2 &lt; 250</math></li> </ol>
Septic Shock	The following two conditions hold: <ol style="list-style-type: none"> <li>1. The patient has severe sepsis</li> <li>2. The patient experiences hypotension (i.e., systolic blood pressure <math>&lt;90</math> mmHg) for at least last 30 min</li> </ol>
None	The following two conditions hold: <ol style="list-style-type: none"> <li>1. Heart rate was measured in the last 2 h, temperature was measured in the last 8 h, respiratory rate was measured in the last 2 h, arterial partial pressure of <math>\text{CO}_2</math> was measured in the last 8 h, and white blood cell count was measured in the last 8 h</li> <li>2. At most one of the SIRS conditions holds</li> </ol>

severity state leading up to adverse events. We consider septic shock, an adverse event of sepsis, and measure whether the learned severity scores increase leading up to septic shock as one would expect. Indeed, we show that the L-DSS and NL-DSS scores show a significant upward trend in the time period leading up to the adverse event.

Next, we evaluate whether scores trained by L-DSS and NL-DSS are sensitive to changes in severity state due to therapy. Specifically, we compare the trend of the disease severity score before and after *fluid bolus*, a therapy used to relieve hypotension in septic patient (Dellinger et al. 2013). We show that the learned scores show significant change in their trend around the time of treatment administration, thus indicating sensitivity to treatment responses. For instance, a DSS score that trends upward over the time period leading up to

administration of fluid bolus, is likely to trend down during the time period after treatment administration or to trend up at a slower pace.

Motivated by the results showing sensitivity to impending adverse events, in the last experiment, we measure the performance of the learned severity scores for early detection of septic shock. We train a simple classifier using the DSS and its trend features to predict risk of septic shock onset in the next 48 h. We show that this predictor significantly outperforms routinely used clinical scoring systems.

The rest of this section proceeds as follows. In Sect. 3.3.1 we describe our experimental setup and the procedure for the automated generation of clinical comparison pairs. Section 3.3.2 presents the baseline methods. Finally, in Sect. 3.3.3 we present the numerical results and analysis.

### 3.3.1 Experimental setup and automatic generation of the clinical comparison pairs

We begin by randomly dividing the 16, 234 patients in our dataset into training (60%) and testing sets (40%). Within the training set, we assign two thirds of the patients to the development set and the remaining third to the validation set. For each of the development, validation and testing sets of patients we generate a separate set  $O$  of clinical comparison pairs. We consider six combinations of possible pairs of different sepsis stages, i.e., none-SIRS, none-severe, none-shock, SIRS-severe, SIRS-shock, severe-shock. For each combination of stages, we randomly select an equal number of feature vectors  $(\mathbf{x}_i^p, \mathbf{x}_j^q)$  sampled at time points  $(t_i^p, t_j^q)$  such that  $\mathbf{x}_i^p$  corresponds to a more severe state and  $\mathbf{x}_j^q$  corresponds to a less severe state in this combination. For the development and testing sets we sample 2000 clinical comparisons for each combinations of sepsis severity stages resulting in total of 12000 clinical comparisons for each set. For the validation set that contains only half of the number of patients in the development and testing sets, we sample 1000 clinical comparisons for each combination of sepsis severity stages resulting in total of 6000 clinical comparisons.


### 3.3.2 Baselines: routinely used clinical severity scores in the ICU

We compare the performance of the learned disease severity scores to three widely used ICU-based severity scoring systems (Keegan et al. 2011). The first two scores are based on the Sequential Organ Failure Assessment or the SOFA score (Vincent et al. 1996) which was originally designed to assess sepsis-related organ damage severity. The SOFA method scores severity at the per-organ level. Two variants of the SOFA that are commonly used are: (1) total SOFA computed as the sum of SOFA scores of all organ systems, and (2) worst SOFA represented as the highest value of SOFA score among of all organ systems. We also compare the performance of our score to that of the Acute Physiology and Chronic Health Evaluation or APACHE II (Knaus et al. 1985), which is a widely used scoring system for assessing general (not necessarily sepsis-related) disease severity in hospitalized individuals.

### 3.3.3 Performance evaluation of scores trained using the L-DSS and NL-DSS algorithms

In this section, we present numerical results of evaluation of the learned L-DSS and NL-DSS scores.

*Selection of free parameters* The L-DSS and NL-DSS algorithms contain free parameters  $\lambda_O$  and  $\lambda_S$  that remain to be specified. Let us first consider the L-DSS algorithm. With  $\lambda_S$



**Fig. 4** Experiment 1. Severity ordering accuracy and the shock prediction AUC on the validation set for different values of  $\lambda_S$ . We mark the values of  $\lambda_S$  selected for further evaluations, i.e.,  $\lambda_S^{soa}$  and  $\lambda_S^{auc}$ , with vertical lines

set to 0, we sweep the values of  $\lambda_O$  from 1 to  $10^{15}$  and select the value of  $\lambda_O$  that maximizes accuracy of ordering held out pairs on the validation set. That is, we count the fraction of ordering pairs in the set  $O$  that are concordant with the ordering prescribed by the ground truth comparisons. We refer to this quantity as the severity ordering accuracy or SOA. In the evaluations below,  $\lambda_O$  is thus set to  $10^7$ .

For a given  $\lambda_O$ , different choices of  $\lambda_S$  yield trajectories with differing levels of smoothness. An appropriate choice of  $\lambda_S$  can be made by sweeping through a wide range of values and selecting the value that optimizes performance for a use case that the end user has in mind. For example, if the primary application of the score is for the early detection of individuals at risk for septic shock, the  $\lambda_S$  that maximizes prediction accuracy on the validation set is selected. Alternately, when no assumptions are given,  $\lambda_S$  can be selected to maximize smoothness without hurting ordering accuracy on the validations set. We present results using these two approaches for setting  $\lambda_S$ ; we refer to these settings as  $\lambda_S^{auc}$  and  $\lambda_S^{soa}$  respectively. In Fig. 4, we show performance on the validation set and mark the selected setting of  $\lambda_S$  for each of the scores. Thus,  $\lambda_S^{auc}$  and  $\lambda_S^{soa}$  were set to be  $1.62 \cdot 10^8$  and  $1.13 \cdot 10^5$  for L-DSS and 2000 and 100 for NL-DSS. While we do not experiment with this approach, it is worth noting that yet another means for selecting  $\lambda_S$  is based on an expert’s knowledge of the degree of short-term to long-term variability expected within that disease domain. For example, in slowly evolving diseases, the value of  $\lambda_S$  that yields a small ratio of short-term to long-term variability may be preferable.

*Experiment 1: Distinguishing between the severity stages of sepsis* We begin by evaluating whether L-DSS and NL-DSS can distinguish and correctly order the different stages of sepsis severity. We compare their severity ordering accuracy to that of routinely used clinical



**Table 3** Experiment 1. Severity ordering accuracy (SOA) for different methods. The 95 % confidence interval on SOA is obtained using the bootstrap algorithm and is contained in the  $\pm 0.002$  band around SOA

	Method	$\lambda_S = \lambda_S^{\text{soa}}$	$\lambda_S = \lambda_S^{\text{auc}}$
Proposed	L-DSS	0.860	0.844
Scores	NL-DSS	0.946	0.938
Routine	APACHE II	0.68	
Clinical	Total SOFA	0.63	
Scores	Worst SOFA	0.63	

scores—APACHE-II, Total SOFA and Worst SOFA (Keegan et al. 2011). The results of this evaluation are presented in Table 3.

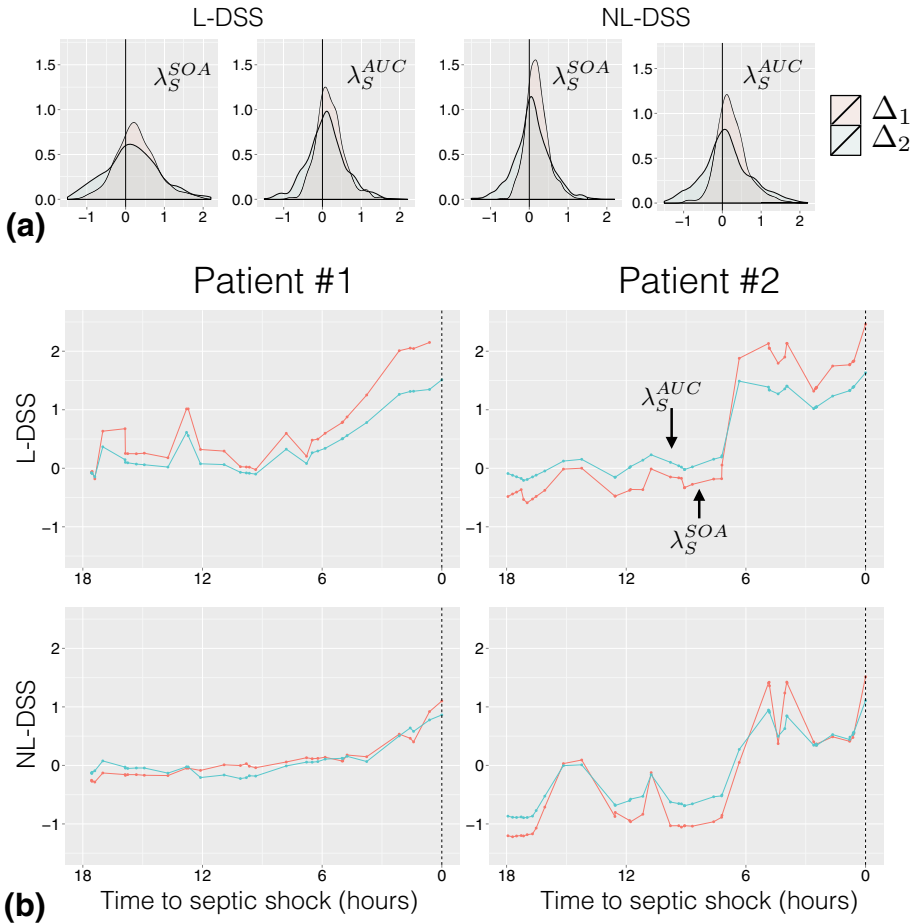
We observe that L-DSS and NL-DSS significantly outperform APACHE II, Total SOFA and Worst SOFA scores for all considered values of  $\lambda_S$ . The performance achieved by L-DSS and NL-DSS is significant from a clinical standpoint as it orders severity states more accurately than the three clinical scores, all of which are widely used to assess severity of ICU patients. In particular, the SOFA score was designed specifically to measure sepsis related severity.

While the above-mentioned result is promising, it remains to be seen whether the obtained scores are also sensitive enough to capture changes in severity status that extend beyond the coarse grading that the different stage definitions provide. Towards this, we next evaluate whether the learned scores exhibit the following desirable characteristics: (1) Are they sensitive to changes in severity leading up to septic shock, an adverse event of sepsis? and, (2) Are they sensitive to post-therapy changes in severity?

*Experiment 2: Are the learned DSS sensitive to changes in severity leading up adverse events?* To address the question of whether the learned scores are sensitive enough to capture changes in severity that can occur leading up to an adverse event, we examine the L-DSS and NL-DSS behavior in the 18-h duration leading up to septic shock.

We consider all patients with septic shock in our test set with at least 18 h of data prior to septic shock onset ( $N = 587$ ). On these patients, we define three time intervals of interest: (1) 6 h prior to the onset of septic shock; (2) 6–12 h prior to the onset of the septic shock; (3) 12–18 h prior to the onset of septic shock. We denote the average values of the learned scores in these intervals by  $\bar{s}_{0-6}$ ,  $\bar{s}_{6-12}$ , and  $\bar{s}_{12-18}$ , respectively.

We calculate values of  $\Delta_1 = \bar{s}_{0-6} - \bar{s}_{6-12}$  and  $\Delta_2 = (\bar{s}_{0-6} - \bar{s}_{6-12}) - (\bar{s}_{6-12} - \bar{s}_{12-18})$  for each patient. In Fig. 5 (a) we show the full probability density of  $\Delta_1$  and  $\Delta_2$ . The value of  $\Delta_1$  is positive in at least 70 % of the cases for all four considered scores. The value of  $\Delta_2$  is positive in at least 57 % of the cases. Using the standard one-tailed t-test, we assess the  $p$  value (denoted by  $p_{\text{trend-up}}$  in Table 4) for whether the recorded  $\Delta_1$  can be observed by chance under the null hypothesis that  $\Delta_1$  are drawn from a zero mean distribution. Similarly, we assess the  $p$  value (denoted by  $p_{\text{rate acceleration}}$  in Table 4) for whether the recorded  $\Delta_2$  can be observed by chance under the null hypothesis that  $\Delta_2$  are drawn from a zero mean distribution. Across all values of  $\lambda_S$ , for both the L-DSS and the NL-DSS, the obtained  $p$  values rule out the null hypothesis, that is, the learned scores leading up to septic shock show significant upward trend and acceleration. Using the bootstrap, we estimate the median  $p$  value for a range of sample sizes and significance is achieved (i.e., the median  $p$  value for that sample size is below 0.01) with as few as 30 samples for  $\Delta_1$  and 420 for  $\Delta_2$ . As an



**Fig. 5** Experiment 2. Sensitivity of learned DSS to changes in severity leading up to adverse events. **a** Probability density of  $\Delta_1$  and  $\Delta_2$ ; **b** DSS trajectories over the 18 h period leading up to septic shock for two example patients

example, in Fig. 5(b), we show the L-DSS and NL-DSS trajectories for two patients for the period leading up to septic shock.

*Experiment 3: Are the learned DSS sensitive to post-therapy changes in severity?* We now evaluate whether the scores trained by the L-DSS and NL-DSS methods are sensitive to changes in severity state due to administration of fluid bolus—a treatment used for septic shock (Dellinger et al. 2013). Towards this, we use the self-controlled case series method. We compare trends exhibited by DSS values over the 5-h intervals prior to and post the administration of fluid bolus. We refer to the trends over these intervals as  $\Delta_{\text{prior}}$  and  $\Delta_{\text{post}}$ . The value of  $\Delta_{\text{prior}}$  is computed as the difference between the value of the DSS at the time of treatment administration and the mean value of DSS over the 5-h interval prior to treatment administration. Similarly, the value of  $\Delta_{\text{post}}$  is calculated as the difference between the mean value of DSS over the 5-h interval after treatment administration and the value of DSS at the moment of treatment administration. If the patient is responsive to fluid therapy,

**Table 4** Experiment 2.  $p$  value  $p_{\text{trend-up}}$  for the observed  $\Delta_1$ ;  $p$  value  $p_{\text{rate acceleration}}$  for the observed  $\Delta_2$ ; the fraction of positive  $\Delta_1$  with 95 % confidence interval, and the fraction of positive  $\Delta_2$  with 95 % confidence interval

Method	$\lambda_S = \lambda_S^{\text{soa}}$	$\lambda_S = \lambda_S^{\text{auc}}$
	$P_{\text{trend-up}}$	
L-DSS	$10^{-36}$	$10^{-39}$
NL-DSS	$10^{-39}$	$10^{-46}$
	$P_{\text{rate increase}}$	
L-DSS	$1.3 \times 10^{-5}$	$4 \times 10^{-4}$
NL-DSS	$1.3 \times 10^{-3}$	$2.6 \times 10^{-3}$
	Fraction of positive $\Delta_1$ (95 % confidence interval)	
L-DSS	0.71 (0.68–0.75)	0.70 (0.67–0.74)
NL-DSS	0.74 (0.70–0.77)	0.75 (0.71–0.78)
	Fraction of positive $\Delta_2$ (95 % confidence interval)	
L-DSS	0.58 (0.54–0.62)	0.57 (0.53–0.61)
NL-DSS	0.59 (0.56–0.63)	0.58 (0.55–0.63)

Calculations of all values are based on 587 examples

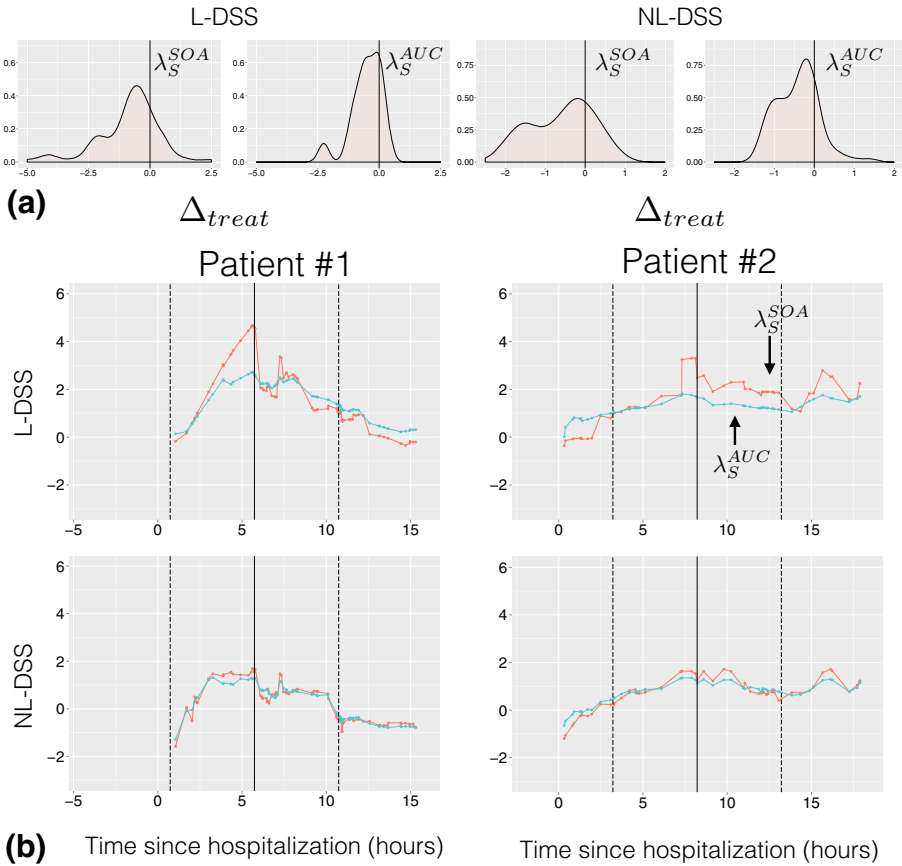
then  $\Delta_{\text{treat}} = \Delta_{\text{post}} - \Delta_{\text{prior}} < 0$ , that is, if the DSS was trending up prior to treatment administration, we expect this trend to be attenuated or even reversed by the treatment.

We identify cases of fluid administration events related to sepsis using the following criteria: (1) the patient is experiencing SIRS, severe sepsis or septic shock at the time of treatment administration, and (2) the patient is hypotensive (has systolic blood pressure below 100 mmHg), a commonly used criteria for prescribing fluids in sepsis. To avoid confounding due to multiple administration of fluids, we restrict our attention to treatment administrations that were not preceded or followed by another fluid bolus administration within a 5-h window. This yielded a total of 81 fluid bolus administration events.

In Fig. 6 (a) we plot the distribution of  $\Delta_{\text{treat}} = \Delta_{\text{post}} - \Delta_{\text{prior}}$ . Overall, the change of trend  $\Delta_{\text{treat}}$  is negative in at least 75 % of recorded values of  $\Delta_{\text{treat}}$ . Employing the one-tailed t-test, we obtain the  $p$  value  $p_{\text{treatment response}}$  (shown in Table 5) for whether the observed values of  $\Delta_{\text{treat}} = \Delta_{\text{post}} - \Delta_{\text{prior}}$  can be observed by chance under the null hypothesis that  $\Delta_{\text{treat}}$  are drawn from a zero mean distribution. For our sample size of 81 cases, across all values of  $\lambda_S$  for both the L-DSS and NL-DSS, the obtained  $p$  values rule out the null hypothesis in favor of the stated hypothesis, that is, DSS shows significant response to therapy. Moreover, using the bootstrap to estimate the median  $p$  value for a range of samples sizes, we observe that significance is achieved with as few as 20 samples. In Fig. 6 (b), we show the L-DSS and NL-DSS trajectories for two example patients around the time of fluid bolus administration.

*Predictive score for septic shock using DSS:* The high ordering performance in experiment 1 and the significant upward trend observed in experiment 2 suggest that a score such as DSS maybe useful for detecting individuals at risk for septic shock. Thus, we conclude this section by showing that the value of the severity score and its temporal trajectory can be used in prediction tasks, specifically, in the task of early detection of septic shock.

We begin by describing the trend features that we derive from sequences of instantaneous values of severity scores. For a patient  $p$ , we let  $s_i^p$  for  $i = 1, \dots, T_p$  be the value of the assigned severity scores at time  $t_i^p$ . At every time point  $t_i^p$ , we augment the score value  $s_i^p$  with the seven derived trend features that were inspired by Wiens et al. (2012) and adapted to



**Fig. 6** Experiment 3. Sensitivity of DSS to post-treatment changes in severity. **a** Probability density of  $\Delta_{treat}$ ; **b** DSS trajectories before and after administration of fluid bolus for two example patients

**Table 5** Experiment 3. Statistical significance of DSS response to fluid bolus treatment.  $p$  value  $p_{treatment\ response}$  for the observed  $\Delta_{treat}$ ; the fraction of negative  $\Delta_{treat}$  with 95% confidence interval

Method	$\lambda_S = \lambda_S^{soa}$	$\lambda_S = \lambda_S^{auc}$
	$p_{treatment\ response}$	
L-DSS	$5 \times 10^{-7}$	$2 \times 10^{-10}$
NL-DSS	$5.1 \times 10^{-8}$	$4.8 \times 10^{-10}$
	Fraction of negative $\Delta_{treat}$ (95% confidence interval)	
L-DSS	0.84 (0.77–0.93)	0.77 (0.68–0.85)
NL-DSS	0.75 (0.67–0.85)	0.84 (0.77–0.93)

Calculations of all values are based on 81 examples

the specifics of the MIMIC II dataset. We present the complete list of these features in Table 6. In this table, Feature 1 is the average value of the score since admission. We note that the feature vectors are not sampled at a fixed rate, i.e., the length of the time interval between

**Table 6** Trend features derived from the trajectory of the severity score

Feature	Description	Expression
1	Average score since admission	$(1/t_i^p) \sum_{j=1}^i (t_j^p - t_{j-1}^p) s_j^p$
2	Linear weighted score	$\sum_{j=1}^i (t_j^p \cdot s_j^p) / (\sum_{j=1}^i t_j^p)$
3	Quadratic weighted score	$\sum_{j=1}^i ((t_j^p)^2 s_j^p) / (\sum_{j=1}^i (t_j^p)^2)$
4	Average score change rate	$(s_i^p - s_1^p) / (t_i^p - t_1^p)$
5	Linearly weighted average score change rate	$\sum_{j=1}^i (s_j^p - s_{j-1}^p) \cdot \frac{t_j^p - t_{j-1}^p}{t_i^p - t_1^p}$
6	Average absolute score change rate	$\sum_{j=1}^i  s_j^p - s_{j-1}^p  \cdot \frac{t_j^p - t_{j-1}^p}{t_i^p - t_1^p}$
7	Linearly weighted absolute score change rate	$\sum_{j=1}^i  s_j^p - s_{j-1}^p  \cdot \frac{t_j^p - t_{j-1}^p}{t_i^p - t_1^p}$

two consecutive feature vectors of a patient need not be fixed. We thus weigh every value of the severity score with the length of the time interval between the current feature vector and the previous one. Features 2 and 3 are versions of Feature 1 where more recent scores have a higher relative weight. Feature 4 captures the average rate of change in severity score since admission. Feature 5 is a version of Feature 4 in which more recent rates of score change are given higher relative weight. Finally, Features 6 and 7 capture the variability of the score.

To train a predictive score for the task of early identification of sepsis, we use logistic regression to learn a mapping from a patient’s feature vectors to the probability of occurrence of septic shock in the following 48 h. This approach was inspired by (Clermont et al. 2001; Minne et al. 2008; Ho et al. 2012). Specifically, we use logistic regression regularized with an elastic net penalty (Zou and Hastie 2005). As positive training examples, we consider feature vectors  $\mathbf{x}_i^p$  taken less than 48-h prior to an adverse event. As negative training examples, we take feature vectors from patients that do not experience the adverse event during their hospital stay and do not receive any treatment of fluid-bolus. The choice of leaving out individuals who receive fluid-bolus but do not experience septic shock is owing to the fact that their outcome is censored due to treatment (Paxton et al. 2013). We refer to this procedure as *LR-Shock*. Employing LR-Shock with severity score and its trend features as its input, we obtain a predictor of the onset of septic shock in the next 48 h. We refer to predictors based on the L-DSS and NL-DSS scores and their derived features as LR-Shock+L-DSS+Derived and LR-Shock+NL-DSS+Derived, respectively. As a baseline for comparison, we train LR-Shock+ $\mathbf{x}$ , which is a predictor trained by LR-Shock with feature vectors  $\mathbf{x}$  as its input.

The performance of these predictors is measured in terms of per-patient prediction accuracy that is determined in the following way. Consider an arbitrary value of a threshold  $\tau$ . We say that a patient was correctly identified to have septic shock if he or she had a severity score higher than  $\tau$  at least once prior to the onset of septic shock. We say that a patient was falsely identified to have septic shock if he did not experience septic shock during his hospital stay, but his or her severity score rose above  $\tau$  at some time point. By considering a series of values of  $\tau$ , we can obtain receiver operating curve that corresponds to our predictor and the associated area under the curve (AUC). In our experiment, we use the AUC on the validation set of patients to find the optimal values of free parameters of the LR-Shock method. We

**Table 7** Per-patient accuracy of early detection of sepsis in terms of the corresponding AUC. The 95 % confidence interval on the AUC is obtained using the bootstrap method and is given in parentheses

	Features	$\lambda_S = \lambda_S^{\text{soa}}$	$\lambda_S = \lambda_S^{\text{auc}}$
Predictors based on proposed scores	L-DSS	0.836 (0.824–0.849)	0.853 (0.841–0.865)
	NL-DSS	0.859 (0.846–0.872)	0.878 (0.866–0.890)
Predictor based on feature vectors alone	L-DSS + Derived	0.856 (0.844–0.868)	0.857 (0.845–0.869)
	NL-DSS + Derived	0.861 (0.849–0.874)	0.874 (0.862–0.886)
Routine	APACHE II	0.620 (0.600–0.641)	
Clinical	Total SOFA	0.602 (0.582–0.622)	
Scores	Worst SOFA	0.601 (0.581–0.621)	

report the performance of our classifiers in terms of AUC on the test set (see Table 7). This table also contains the predictive performance of the L-DSS and NL-DSS scores.

The learned scores are significantly more accurate than the APACHE II and the SOFA based scores. However, more interestingly, we note that DSS performance is comparable to LR-Shock+x which was trained to optimize predictive performance unlike the DSS scores. Thus, our learning objective addresses the bias introduced due to treatment related confounding without hurting predictive performance in this application.

## 4 Related work on ranking

Our work is closely related to the body of literature on pairwise methods for ranking in the field of Information Retrieval (IR). In IR, the ranking problem is typically formulated as the task of sorting retrieved documents by their relevance to the query. Pairwise ranking approaches (e.g., [Joachims \(2002\)](#); [Zheng et al. \(2008\)](#)) aim to learn a ranking function that orders pairs of documents in concordance to their relevance. These ranking functions are derived using a variety of machine learning techniques. [Joachims \(2002\)](#) proposed a max-margin approach to learning a ranking, dubbed SVMRank. A computationally efficient method for training of the SVMRank as a primal-form optimization problem was later proposed by [Chapelle and Keerthi \(2010\)](#). [Burges et al. \(2005\)](#) proposed RankNet, a neural network based algorithm for ranking that tunes its parameters to minimize a simple explicit probabilistic cost function that captures the task of pairwise ranking. In later work, [Burges et al. \(2006\)](#) and [Burges \(2010\)](#) observed that a ranking function can be learned from gradients of the ranking objective function alone, without explicit specification of the whole cost function. They proposed two new ranking algorithms that build on this observation, one based on neural networks and one based on gradient boosting trees. [Zheng et al. \(2008\)](#) proposed a general boosting framework for learning ranking functions for a wide family of cost functions. Additional approaches include FRank ([Tsai et al. 2007](#)), nested ranker ([Matveeva et al. 2006](#)), and multiple hyperplane ranker ([Qin et al. 2007](#)). In some disease domains, as is the case in ours, one might be able to obtain rank data rather than pair-wise comparisons. In these cases, ordinal regression based approaches have been developed for ranking ([Herbrich et al. 2000](#); [Chu and Keerthi 2007](#)). However, acquiring ranked samples from the clinician is often not practical.

Moreover, by relying on pairwise comparisons, our framework opens up the possibility of exploring new forms of supervision that can be automatically generated. For example, two time slices with the same severity grade may be ordered based on their time to an adverse event in the case when no interventions have been administered between these time slices.

## 5 Discussion and future work

This paper proposes DSSL, a novel ranking-based framework for scalable and automated learning of disease severity scores in new disease domains and populations. DSSL only requires a means for obtaining clinical comparisons—ordered pairs comparing disease severity state at different times. We argue that this form of supervision is more natural to elicit than asking clinical experts to map the disease severity score, or to encode an accurate model of disease progression. Moreover, supervision of this type can also be obtained in an automated way by leveraging existing clinical guidelines.

We test DSSL by applying it to a large, real-world electronic health record dataset and to synthetic clinical records. Using synthetic clinical records, we show that scores learned using DSSL are less sensitive to changes in treatment administration patterns between the train and test environments compared to the regression based approach that is currently used. Using a large real-world dataset of ICU clinical records, we show that the scores learned using DSSL are significantly more accurate, both for severity assessment and early adverse-event detection, compared to widely used clinical severity scores. Further, these scores have face validity—their behavior aligns with what is expected clinically. They trend upwards leading up to an adverse event, and show decline post-treatment.

DSSL has a number of other advantages. It allows experts to automatically tune the quality of the score by increasing the granularity and amount of supervision given. Additionally, the quality of the learned scores can be improved by incorporating additional constraints related to disease progression. For example, expected clinical response to therapy can be directly incorporated as a constraint within the optimization objective.

One limitation of our current work is the heavy reliance on the availability of a large number of clinical comparison pairs<sup>1</sup> In domains where existing clinical guidelines cannot be leveraged, it is unrealistic to obtain thousands of clinical comparison pairs from experts. In these domains, the use of active learning may help mitigate this limitation. Supervision in the form of additional constraints related to disease progression may also prove helpful. Another aspect that deserves further exploration is how the proposed scores can be made interpretable in practice. While NL-DSS yields high performance, the score is constructed using a bag of regression trees. Thus, is it not obvious how one might make the score interpretable at the point of care. In practice, scores are often deployed with a specific use case in mind. For example, if the score were to be used for early detection, the precision-recall curve can be used to identify suitable thresholds for taking action based on the DSS score. Another approach might be to identify and to display which factors led to the increase or decrease in the score value. Finally, using simulated data, our experiments show that the scores learned using DSSL are less dependent on the practice patterns of the regime where the model was developed. While promising, further analysis is needed to understand susceptibility of DSSL to treatment

---

<sup>1</sup> In other experiments (Dyagilev and Saria 2015), we have shown that the performance of our framework degrades gracefully as the number of clinical comparisons is reduced. In particular, on the task of severity assessment, scores trained on approximately 150 clinical comparisons are as accurate as APACHE and SOFA.

administration patterns. This includes developing metrics that characterize susceptibility of prediction algorithms to changes in provider practice patterns.

In summary, electronic tools that can integrate the diverse and the large set of measurements collected clinically to produce an accurate, real-time severity score can enable clinicians to provide more timely interventions. Further, these scores should be robust to changes in clinical practice patterns as the mere introduction of a decision-support tool can change clinician behavior. This paper introduces a new ranking-based formulation for the problem of learning (predictive) disease severity scores. By leveraging clinical comparisons, a form of supervision that is less susceptible to clinician practice patterns, DSSL provides a promising alternative to existing methods.

**Acknowledgments** The authors would like to thank Katharine Henry for her help with data preprocessing. This work was supported by ACH grant #117762.

## 6 Appendix: Features used for DSS learning

In this section we provide the complete list of all features provided to L-DSS and NL-DSS methods for DSS learning in experiments in Sect. 3.3.3. These can be divided into three categories: clinical information, measurements of vital signals, and results of laboratory analysis.

*Clinical information* Age of the patient; whether patient has a pacemaker; whether patient was diagnosed with AIDS; whether patient received treatment that compromised his immune system; patient's current weight and his weight on admission; presence of ICD-9 codes for diabetes, dialysis, chronic renal insufficiency, heart failure, or chronic liver disease; whether patient is currently in the cardiac surgery recovery unit; presence or absence of hematologic malignancy; jaundice; whether a patient was mechanically ventilated; presence of metastatic carcinoma.

*Measurements of vital signals* Glasgow coma score; heart rate; Riker Sedation-Agitation Scale; temperature; respiratory rate; systolic blood pressure; shock index defined as the ratio of heart rate to systolic blood pressure; peripheral capillary oxygen saturation;

*Results of laboratory analysis* Blood urea nitrogen levels (BUN); hematocrit; international normalized ratio (INR); white blood cell count (WBC); blood pH level as measured by an arterial line; partial pressure of arterial oxygen ( $PaO_2$ ); fraction of inspired oxygen ( $FiO_2$ ); ratio of  $PaO_2$  to  $FiO_2$ ; partial pressure of  $CO_2$ ; blood lactate measurements; bilirubin; creatinine; potassium and sodium levels; platelet count; hemoglobin; total urine output over the past 2 h per kg of weight; partial thromboplastin time; arterial  $CO_2$  levels; levels of aspartate aminotransferase.

## References

- AHRQ. (2015). *Guideline syntheses*. <http://www.guideline.gov/syntheses/index.aspx>.
- Burges, C., Shaked, T., Renshaw, E., Lazier, A., Deeds, M., Hamilton, N., & Hullender, G. (2005). Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on machine learning*, ACM, (pp. 89–96).
- Burges, C. J. (2010). *From ranknet to lambdarank to lambdamart: An overview*. Technical report, Microsoft Research.
- Burges, C. J., Ragno, R., & Le, Q. V. (2006). Learning to rank with nonsmooth cost functions. In *Advances in neural information processing systems*, (pp. 193–200).
- Chapelle, O., & Keerthi, S. S. (2010). Efficient algorithms for ranking with SVMs. *Information Retrieval*, 13(3), 201–215.



- Chu, W., & Keerthi, S. S. (2007). Support vector ordinal regression. *Neural Computation*, 19(3), 792–815.
- Clermont, G., Angus, D. C., DiRusso, S. M., Griffin, M., & Linde-Zwirble, W. T. (2001). Predicting hospital mortality for patients in the intensive care unit: A comparison of artificial neural networks with logistic regression models. *Critical Care Medicine*, 29(2), 291–296.
- 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.
- Dyagilev, K., & Saria, S. (2015). Learning a severity score for sepsis: A novel approach based on clinical comparisons. In *AMIA Annual symposium proceedings, American Medical Informatics Association*
- Fine, M. J., Auble, T. E., Yealy, D. M., Hanusa, B. H., Weissfeld, L. A., Singer, D. E., et al. (1997). A prediction rule to identify low-risk patients with community-acquired pneumonia. *New England Journal of Medicine*, 336(4), 243–250.
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29(5), 1189–1232.
- Ghanem-Zoubi, N. O., Vardi, M., Laor, A., Weber, G., & Bitterman, H. (2011). Assessment of disease-severity scoring systems for patients with sepsis in general internal medicine departments. *Critical Care Medicine*, 15(2), R95.
- Henry, K. E., Hager, D. N., Provonost, P. J., & Saria, S. (2015). A targeted real-time early warning score (TREWScore) for septic shock. *Science Translational Medicine*, 7, 299ra122.
- Herbrich, R., Graepel, T., & Obermayer, K. (2000). Large margin rank boundaries for ordinal regression. In: *Advances in Large Margin Classifiers*, (pp. 115–132). Cambridge: The MIT Press.
- Ho, J. C., Lee, C. H., & Ghosh, J. (2012). Imputation-enhanced prediction of septic shock in ICU patients. In *Proceedings of the ACM SIGKDD workshop on health informatics (HI-KDD12)*.
- Hothorn, T., Hornik, K., & Zeileis, A. (2006). Unbiased recursive partitioning: A conditional inference framework. *Journal of Computational and Graphical Statistics*, 15(3), 651–674.
- Hug, C. (2009). *Detecting hazardous intensive care patient episodes using real-time mortality models*. PhD thesis.
- Jackson, C. H., Sharples, L. D., Thompson, S. G., Duffy, S. W., & Couto, E. (2003). Multistate Markov models for disease progression with classification error. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 52(2), 193–209.
- Joachims, T. (2002). Optimizing search engines using clickthrough data. In *Proceedings of the eighth ACM SIGKDD international conference on knowledge discovery and data mining*, ACM, (pp. 133–142).
- Keegan, M. T., Gajic, O., & Afessa, B. (2011). Severity of illness scoring systems in the intensive care unit. *Critical Care Medicine*, 39(1), 163–169.
- Knaus, W. A., Draper, E. A., Wagner, D. P., & Zimmerman, J. E. (1985). APACHE II: A severity of disease classification system. *Critical Care Medicine*, 13(10), 818–829.
- Kumar, G., Kumar, N., Taneja, A., Kaleekal, T., Tarima, S., McGinley, E., et al. (2011). Nationwide trends of severe sepsis in the 21st century (2000–2007). *CHEST Journal*, 140(5), 1223–1231.
- Kuo, T. M., Lee, C. P., & Lin, C. J. (2014). Large-scale kernel RankSVM. In *Proceedings of the 2014 SIAM international conference on data mining*, SIAM.
- Marshall, J. C., Cook, D. J., Christou, N. V., Bernard, G. R., Sprung, C. L., & Sibbald, W. J. (1995). Multiple organ dysfunction score: A reliable descriptor of a complex clinical outcome. *Critical Care Medicine*, 23(10), 1638–1652.
- Mason, L., Baxter, J., Bartlett, P., & Frean, M. (1999). Boosting algorithms as gradient descent in function space. *Advances in Neural Information Processing Systems*, 12, 512–518.
- Matveeva, I., Burges, C., Burkard, T., Laucius, A., & Wong, L. (2006). High accuracy retrieval with multiple nested ranker. In *Proceedings of the 29th annual international ACM SIGIR conference on research and development in information retrieval* (pp. 437–444), ACM.
- Medsgger, T., Bombardieri, S., Czirjak, L., Scorza, R., Rossa, A., & Bencivelli, W. (2003). Assessment of disease severity and prognosis. *Clinical and Experimental Rheumatology*, 21(3; SUPP29), S42–S46.
- Minne, L., Abu-Hanna, A., de Jonge, E., et al. (2008). Evaluation of SOFA-based models for predicting mortality in the ICU: A systematic review. *Critical Care Medicine*, 12(6), R161.
- Mohan, A., Chen, Z., & Weinberger, K. Q. (2011). Web-search ranking with initialized gradient boosted regression trees. In *Yahoo! learning to rank challenge, Citeseer*, (pp. 77–89).
- Mould, D. (2012). Models for disease progression: New approaches and uses. *Clinical Pharmacology & Therapeutics*, 92(1), 125–131.
- Paxton, C., Niculescu-Mizil, A., & Saria, S. (2013). Developing predictive models using electronic medical records: Challenges and pitfalls. In *AMIA annual symposium proceedings, American Medical Informatics Association*, vol. 2013, p. 1109.

- Pirracchio, R., Petersen, M. L., Carone, M., Rigon, M. R., Chevret, S., & van der Laan, M. J. (2015). Mortality prediction in intensive care units with the super ICU learner algorithm (SICULA): A population-based study. *The Lancet Respiratory Medicine*, 3(1), 42–52.
- Qin, T., Zhang, X. D., Wang, D. S., Liu, T. Y., Lai, W., & Li, H. (2007). Ranking with multiple hyperplanes. In *Proceedings of the 30th annual international ACM SIGIR conference on research and development in information retrieval*, ACM (pp. 279–286).
- Saeed, M., Lieu, C., Raber, G., & Mark, R. (2002). MIMIC II: A massive temporal ICU patient database to support research in intelligent patient monitoring. In *Computers in Cardiology*, 2002, IEEE, (pp. 641–644).
- Saeed, M., Villarroel, M., Reisner, A. T., Clifford, G., Lehman, L. W., Moody, G., et al. (2011). Multiparameter Intelligent Monitoring in Intensive Care II (MIMIC-II): A public-access intensive care unit database. *Critical Care Medicine*, 39(5), 952.
- Saria, S., Koller, D., & Penn, A. (2010a). Learning individual and population level traits from clinical temporal data. In *Predictive models in personalized medicine workshop, neural information processing systems*.
- Saria, S., Rajani, A. K., Gould, J., Koller, D., & Penn, A. A. (2010b). Integration of early physiological responses predicts later illness severity in preterm infants. *Science Translational Medicine*, 2(48), 48ra65–48ra65.
- Sebat, F., Musthafa, A. A., Johnson, D., Kramer, A. A., Shoffner, D., Eliason, M., et al. (2007). Effect of a rapid response system for patients in shock on time to treatment and mortality during 5 years. *Critical Care Medicine*, 35(11), 2568–2575.
- Tsai, M. F., Liu, T. Y., Qin, T., Chen, H. H., & Ma, W. Y. (2007). Frank: A ranking method with fidelity loss. In *Proceedings of the 30th annual international ACM SIGIR conference on research and development in information retrieval*, ACM (pp. 383–390).
- Vincent, J. L., Moreno, R., Takala, J., Willatts, S., De Mendonça, A., Bruining, H., et al. (1996). The SOFA (Sepsis-related Organ Failure Assessment) score to describe organ dysfunction/failure. *Intensive Care Medicine*, 22(7), 707–710.
- Wang, X., Sontag, D., & Wang, F. (2014). Unsupervised learning of disease progression models. In *Proceedings of the twentieth ACM SIGKDD international conference on knowledge discovery and data mining*, ACM (pp. 85–94).
- Wiens, J., Horvitz, E., & Guttag, J. V. (2012). Patient risk stratification for hospital-associated c. diff as a time-series classification task. *Advances in Neural Information Processing Systems*, 25, 467–475.
- Zheng, Z., Zha, H., Zhang, T., Chappelle, O., Chen, K., & Sun, G. (2008). A general boosting method and its application to learning ranking functions for web search. In *Advances in Neural Information Processing Systems*, vol. 20, pp. 1697–1704.
- Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2), 301–320.