

Original Paper

# From Data to Optimal Decision Making: A Data-Driven, Probabilistic Machine Learning Approach to Decision Support for Patients With Sepsis

Athanasios Tsoukalas<sup>1</sup>, PhD; Timothy Albertson<sup>2,3</sup>, MD, PhD, MPH; Ilias Tagkopoulos<sup>1</sup>, PhD

<sup>1</sup>Department of Computer Science and Genome Center, University of California, Davis, Davis, CA, United States

<sup>2</sup>Department of Internal Medicine, University of California, Davis, Davis, CA, United States

<sup>3</sup>Northern California Veterans Administration Health Care System, Mather, CA, United States

**Corresponding Author:**

Ilias Tagkopoulos, PhD

Department of Computer Science and Genome Center

University of California, Davis

One Shields Ave

Davis, CA, CA 95616

United States

Phone: 1 5307527707

Fax: 1 5307526747

Email: [iliast@ucdavis.edu](mailto:iliast@ucdavis.edu)

## Abstract

**Background:** A tantalizing question in medical informatics is how to construct knowledge from heterogeneous datasets, and as an extension, inform clinical decisions. The emergence of large-scale data integration in electronic health records (EHR) presents tremendous opportunities. However, our ability to efficiently extract informed decision support is limited due to the complexity of the clinical states and decision process, missing data and lack of analytical tools to advice based on statistical relationships.

**Objective:** Development and assessment of a data-driven method that infers the probability distribution of the current state of patients with sepsis, likely trajectories, optimal actions related to antibiotic administration, prediction of mortality and length-of-stay.

**Methods:** We present a data-driven, probabilistic framework for clinical decision support in sepsis-related cases. We first define states, actions, observations and rewards based on clinical practice, expert knowledge and data representations in an EHR dataset of 1492 patients. We then use Partially Observable Markov Decision Process (POMDP) model to derive the optimal policy based on individual patient trajectories and we evaluate the performance of the model-derived policies in a separate test set. Policy decisions were focused on the type of antibiotic combinations to administer. Multi-class and discriminative classifiers were used to predict mortality and length of stay.

**Results:** Data-derived antibiotic administration policies led to a favorable patient outcome in 49% of the cases, versus 37% when the alternative policies were followed ( $P=1.3e-13$ ). Sensitivity analysis on the model parameters and missing data argue for a highly robust decision support tool that withstands parameter variation and data uncertainty. When the optimal policy was followed, 387 patients (25.9%) have 90% of their transitions to better states and 503 patients (33.7%) patients had 90% of their transitions to worse states ( $P=4.0e-06$ ), while in the non-policy cases, these numbers are 192 (12.9%) and 764 (51.2%) patients ( $P=4.6e-117$ ), respectively. Furthermore, the percentage of transitions within a trajectory that lead to a better or better/same state are significantly higher by following the policy than for non-policy cases (605 vs 344 patients,  $P=8.6e-25$ ). Mortality was predicted with an AUC of 0.7 and 0.82 accuracy in the general case and similar performance was obtained for the inference of the length-of-stay (AUC of 0.69 to 0.73 with accuracies from 0.69 to 0.82).

**Conclusions:** A data-driven model was able to suggest favorable actions, predict mortality and length of stay with high accuracy. This work provides a solid basis for a scalable probabilistic clinical decision support framework for sepsis treatment that can be expanded to other clinically relevant states and actions, as well as a data-driven model that can be adopted in other clinical areas with sufficient training data.

**KEYWORDS**

sepsis; clinical decision support tool; probabilistic modeling; Partially Observable Markov Decision Processes; POMDP; CDSS

## Introduction

Over the past few decades, our society has transitioned to a state where bottlenecks have shifted from a lack of data, to limitations in extracting meaningful knowledge and subsequently use that knowledge to drive decisions. This data-rich, knowledge-poor oxymoron is particularly true in computationally-driven Clinical Decision Support Systems (CDSS), where advances in automated high-throughput data acquisition and electronic health records have yet to be translated into knowledge extraction and probabilistic decision guidance. This is true even in the cases of dangerous and ubiquitous threats to human health, one of which is sepsis. Sepsis is an overwhelming systemic immune response to infection, which results in damage to the patients' own tissues and organs. This process can happen at any age, regardless of the underlying health condition and from seemingly benign incidents. Severe sepsis strikes about 18 million people annually (750,000 cases in the United States) and has a very high short-term mortality risk (28% to 50%) [1]. Severe sepsis is the leading cause of Intensive Care Unit (ICU) deaths (60-80% of ICU deaths in developing countries) and it kills more than 6 million children world-wide every year [2].

Surprisingly, while sepsis is one of the most common diseases (more deaths than prostate cancer, breast cancer, and HIV/AIDS combined [3]), it has the lowest state-funding rates for research. This is in contrast to its severity, occurrence in our society (sepsis hospitalizations have more than doubled over the last 10 years [2]) and money spent to battle it (US \$14.6 billion in 2008, an increase by an average of 11.9% each year). The diagnosis of sepsis is often delayed because it is difficult to differentiate from other high-risk conditions and this delay can lead to the rapid deterioration of the patient. One potentially transformative approach to this problem would be to exploit the vast amount of information that is hidden in the Electronic Health Records (EHR) of patients to derive a CDSS.

Adoption of EHR's by health care systems was predicted to vastly improve the efficiency and quality of patient care [4]. Unfortunately, despite explosive EHR adoption, and enormous associated capital expenditures, these gains have yet to be realized [5,6]. One reason for this failure is that our capacity to utilize complex, large-scale data to generate knowledge and inform clinical decisions remains limited. For example, while CDSS have existed for decades, they are mostly limited to alert systems and (data-oblivious) agent-based suggestions that rely on hard-coded criteria. Although in certain systems patient cases are used for probabilistic training, these efforts focus on feature correlations and final clinical outcomes [7-10] rather than actionable policy (see [11,12] for a review). Our previous work on the associations among EHR observations for lactic acid prediction work falls also in this category (11).

Some of the most powerful methods for modeling decision making in clinical decision support are those that treat the

learning problem as a Markov Decision Process (MDPs) [13]. A MDP is a discrete-time stochastic control process, where the next state depends only on the current state and the action that the decision maker performs, while it is conditionally independent of all other states and actions. An extension to MDPs are the Partially Observable MDPs, (POMDPs), where the states themselves are hidden and only observations are available. In that case, a belief regarding the current system state is formed based on the observations and their state-based likelihoods [14]. There are many methods to solve MDPs/POMDPs, including dynamic programming, linear programming and reinforcement learning [13-17]. When the problem becomes intractable, reinforcing learning methods are preferred as they do not require knowledge of the underlying MDP model.

Surprisingly, although the use of MDP methods in clinical settings is well established, there are only a few notable examples where POMDP has been explored for disease-specific decision support with probabilistic outcomes. MDP has been used for decision support in determining liver acceptance for liver transplants [18-20], HIV therapy initiation [21], breast cancer screening [22], treatment of hepatitis C [23], statin therapy timing [24], among others [25]. However, in most cases the pathology is complex, the underlying state of the patient is latent and we can only observe emitted signals (observations) with some uncertainty. A notable example of POMDP adoption in medical decision support is that of Hauskrecht and Fraser [26,27] who modeled Ischemic Heart Disease with a POMDP model that included both treatment and test actions. This work used an impressive hierarchical model for state variables and performed reasonably well in predicting optimal policies in that given scenario. Similarly, Turgay et al [28] used a POMDP model to support personalized mammography screening decisions in a model that used six states, three of which were fully observable. Kreke et al created a 2-state POMDP model for pneumonia-related sepsis patients that included only a Cytokine test as an investigative action and incorporated the Sequential Organ Failure Assessment (SOFA) score for the MDP state space definition [29,30]. In all these cases, both the state model and policy were course-grained, the parameter space limited and the training sets were a few dozen patients.

Here, we use a point-based POMDP solver together with a dataset of 1492 patients that is complete with time-stamped blood tests, vitals, and other relevant records for sepsis. We model patient trajectories and treatment by defining data-driven states, observations, probabilistic beliefs, actions, and rewards. We then evaluate the potential of this method to inform on optimal administration of antibiotic combinations, defined as antibiotic "policy". In addition, we assess the predictive ability of applying machine learning methods to predict patient mortality and length of stay, in order to drive clinical decision support.

## Methods

### Data Collection

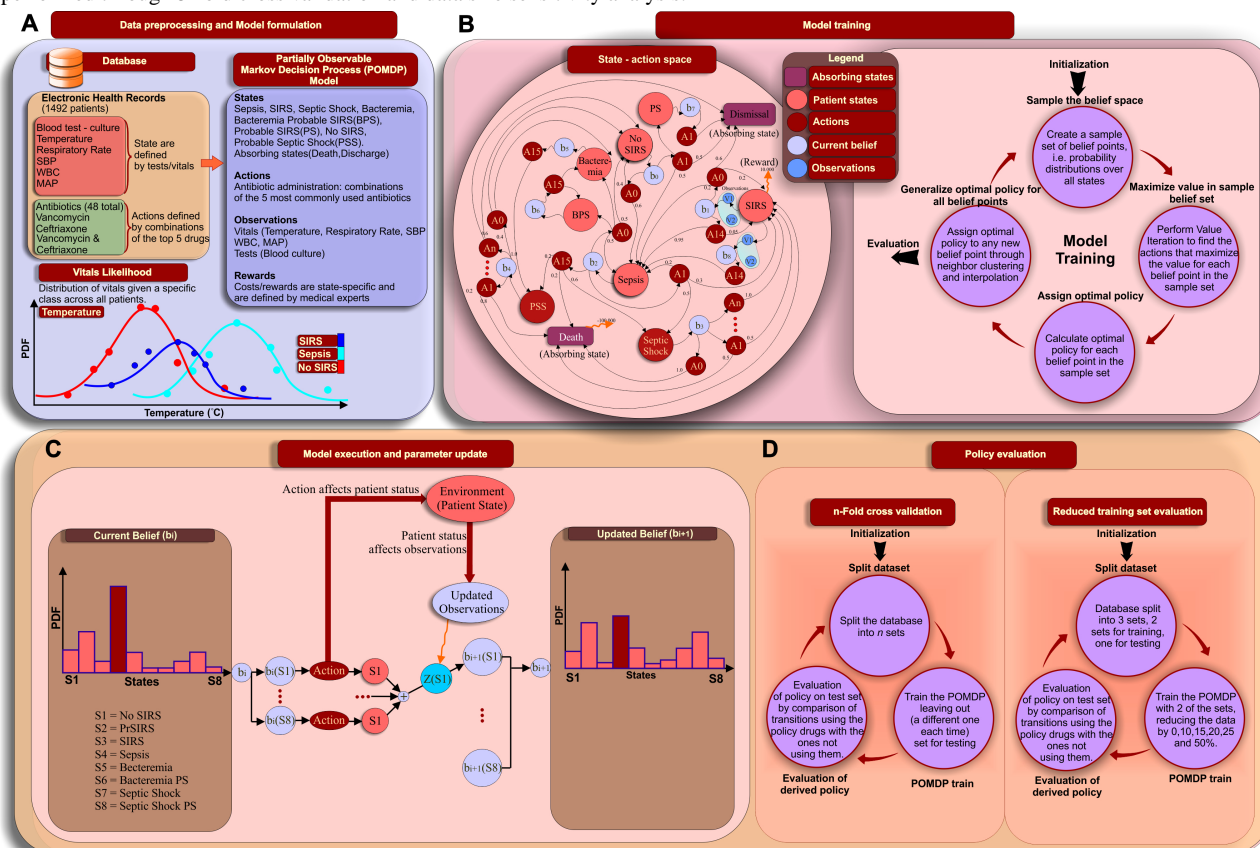
An EHR database containing 1492 adult patients ( $\geq 18$  years of age; ICU cases) with personal health information removed and meeting at least two Systemic Inflammatory Response Syndrome (SIRS) criteria [31] admitted to the University of California Davis Health System (UCDHS), was used for all the analyses, Figure 1 (a). Informed consent was obtained for all human subjects and the analysis was approved by the institutional review board of the University of California, Davis (IRB # 254575). Of the 1492 patients, 45.0% were female, the mean length of stay was 17.0 (SD 36.7) days, and 38.0% were admitted from the emergency department. Table 1 summarizes the dataset used; Figure 2 provides a histogram of the total/ICU

length of stay). UCDHS is a tertiary care, academic medical center that did not have an active EHR alert system for the diagnosis or treatment of sepsis during the study period. All data were abstracted retrospectively from the EHR via structured query language interrogation of a de-identified relational database. Patients were included in the database if they were hospitalized and discharged between 1 January 2010 and 31 December 2010. The following six variables were used as *observation variables* in our model Figure 1 (a): temperature, respiratory rate (RR), white blood count (WBC), mean arterial pressure (MAP), systolic blood pressure (SBP), and blood culture results, with explicit mention of the bacterial species that were present in the culture. The first five variables are measurements of a patient's condition recorded over time; temperature, RR, WBC, and MAP are part of the SIRS criteria.

**Table 1.** Database characteristics.

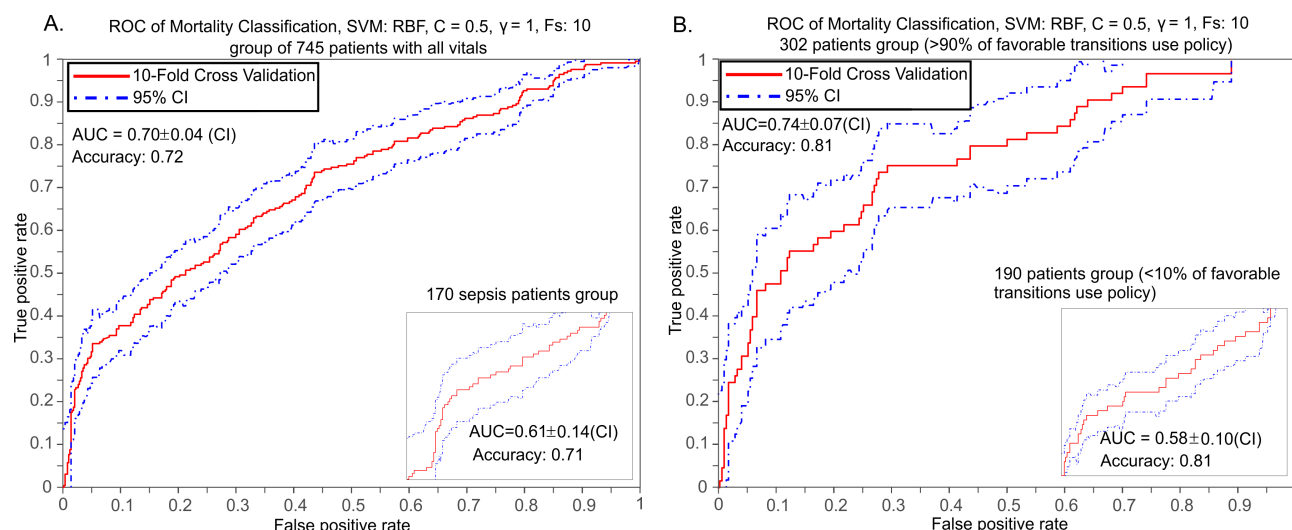
States and state transition distribution in the dataset		
State	Transitions (4200 total)	Patients (1492 total)
No SIRS	1300 (30.85%)	915 (61.33%)
SIRS	294 (6.98%)	264 (17.70%)
Sepsis	41 (0.97%)	38 (2.54%)
Septic Shock	17 (0.40%)	17 (1.14%)
PS (Probable SIRS)	1929 (45.79%)	939 (62.94%)
Bacteremia	157 (3.73%)	121 (8.11%)
BPS (Bacteremia Probable Sepsis)	323 (7.67%)	179 (11.99%)
PSS (Probable Septic Shock)	139 (3.30%)	135 (9.05%)
Demographics and Final Diagnosis		
Characteristic	Value	
Female Gender	677 (45.37%)	
Male Gender	815 (54.63%)	
Mortality Rate	376 (25.2%)	
Length of Stay (days)	Mean: 17 Median: 8	
Intensive Care Unit Stay (days)	Mean: 6.5 Median: 1	
Sepsis Diagnosis	188 (12.60%)	
Septic shock Diagnosis	21 (1.41%)	
Severe Sepsis Diagnosis	8 (0.53%)	

**Figure 1.** Development and evaluation of a clinical decision support system (CDSS) for Sepsis. (A) Synopsis of the EHR database, distribution of vitals, states, actions. Likelihood functions were used to calculate state-specific transition and observation probabilities. (B) The state-action diagram describes underlying patient states, possible transitions and beliefs based on the values of the observed variables. There are two absorbing states, “Death” and “Dismissal”. The training of the POMDP model that is applied on the state-action space performs value iteration updates on a sample set of beliefs, effectively using a Monte Carlo approach for sampling together with dynamic programming for the calculation of the value iteration. (C) Belief (ie, probability distribution of the patient states) is updated based on the action taken and the new observations. At each time step patient vitals are observed and the action that corresponds to the optimal policy is taken. A new set of observations (vitals, tests) will lead to an updated belief that may lead to a new action to be undertaken. The update is asynchronous, as it is calculated on-the-fly as new information arrives. (D) Evaluation of the CDSS framework was performed through 5-fold cross validation and data size sensitivity analysis.





**Figure 2.** ROC curves of mortality classification/prediction. Support Vector Machine training results, given the provided vitals and mortality model, in a 10-fold cross validation scheme. The features used for the classification are temperature, respiratory rate, WBC, MAP and lactate levels and 745 patients for whom all seven variables were available were considered out of the 1492 of the DB. Radial Basis Function (RBF) kernel was used for the SVM training. The five measurement variables were summarized by their mean and standard deviation (STD) across the trajectory of each patient. Principal Component Analysis (PCA) was also used to assess whether linear transformation of the feature space and dimensionality reduction can be achieved in this case. A filter method was applied using the Area Under the Curve (AUC) of the Receiver Operating Characteristic (ROC) curves as a ranking criterion. The positive/negative classes for mortality prediction are defined as alive/deceased respectively. (A) Classification using the patients that have all features available. The maximum prediction accuracy when all patients with available all vitals are used, is 0.72 (72%) and the AUC is 0.70. (Inset) Classification using the 170 patients that have a diagnosis related to Sepsis. The maximum prediction accuracy is 0.71 (71%) and the AUC is 0.61. (B) SVM classification performance when the patient group that have ratio of transitions to better states with policy drugs vs all transitions to better states  $> 0.9$  (302 patients out of 745) is used. (Inset) SVM classification performance for the patient group that have ratio of transition to better states with policy drugs vs all transitions to better states less or equal to 0.1 (190 patients out of 745).



## Sepsis Model

### State Definition

The states that we include in our formulation were selected based on the existence of well-defined criteria and expert opinion, Figure 1 (a). Each of these states is defined by a number of features as illustrated in Table 2. The SIRS criteria that define the respective state are HR  $> 90$  beats per minute, RR  $> 20$  breaths per minute (or partial pressure of arterial  $\text{CO}_2 < 32$  mm Hg), temperature either  $> 38^\circ\text{C}$  or  $< 36^\circ\text{C}$ , and WBC either  $> 12,000$  or  $< 4000$  cells/mm $^3$  (or  $> 10\%$  bands). For a patient to be diagnosed with sepsis, at least two of the SIRS criteria need to be present and a suspected infection should be present (eg, evident through blood test results). Septic shock occurs

when there is sepsis-induced hypotension (where either SBP is below 90 mm Hg,  $< 40$  mm Hg below baseline, or MAP is below 70 mm Hg) that persists despite adequate fluid resuscitation. Additionally, we included states that cannot be fully determined during the training phase because of missing database information, such as missing vitals measurements and time stamps. These states are “Probable Septic Shock” (PSS: Hypotension, positive blood test and no adequate vitals to determine SIRS and/or Sepsis), “Probable SIRS” (PS: no infection and no vitals to determine SIRS), “Bacteremia Probable SIRS” (BPS: infection but no vitals to determine SIRS and/or Sepsis). We have not included the “Severe Sepsis” state in this work, as the current dataset does not offer enough information required to incorporate the organ failure into the patient state definition.

**Table 2.** States and their definition based on vitals and blood tests.

State	Features
No SIRS	
SIRS	heart rate (HR) >90 beats per minute respiratory rate (RR) >20 breaths per minute (or partial pressure of arterial CO <sub>2</sub> < 32) temperature either >38°C or <36°C white blood cell count (WBC) either >12,000 or <4000 cells/mm <sup>3</sup> (or > 10% bands)
Sepsis	SIRS and Infection (blood test result)
Severe Sepsis	Sepsis and Organ failure (shown in ICL code)
Septic Shock	Sepsis and Hypotension (systolic blood pressure (SBP) is below 90 mm Hg, < 40 mm Hg below baseline, or the mean arterial pressure (MAP) is below 70 mm Hg)
PS	No infection– no vitals to determine SIRS
Bacteremia	Infection (blood test result) & No SIRS
BPS	Infection – no vitals to determine SIRS (thus Sepsis)
PSS	Hypotension, positive blood test and no vitals to determine SIRS (thus Sepsis)

## Actions

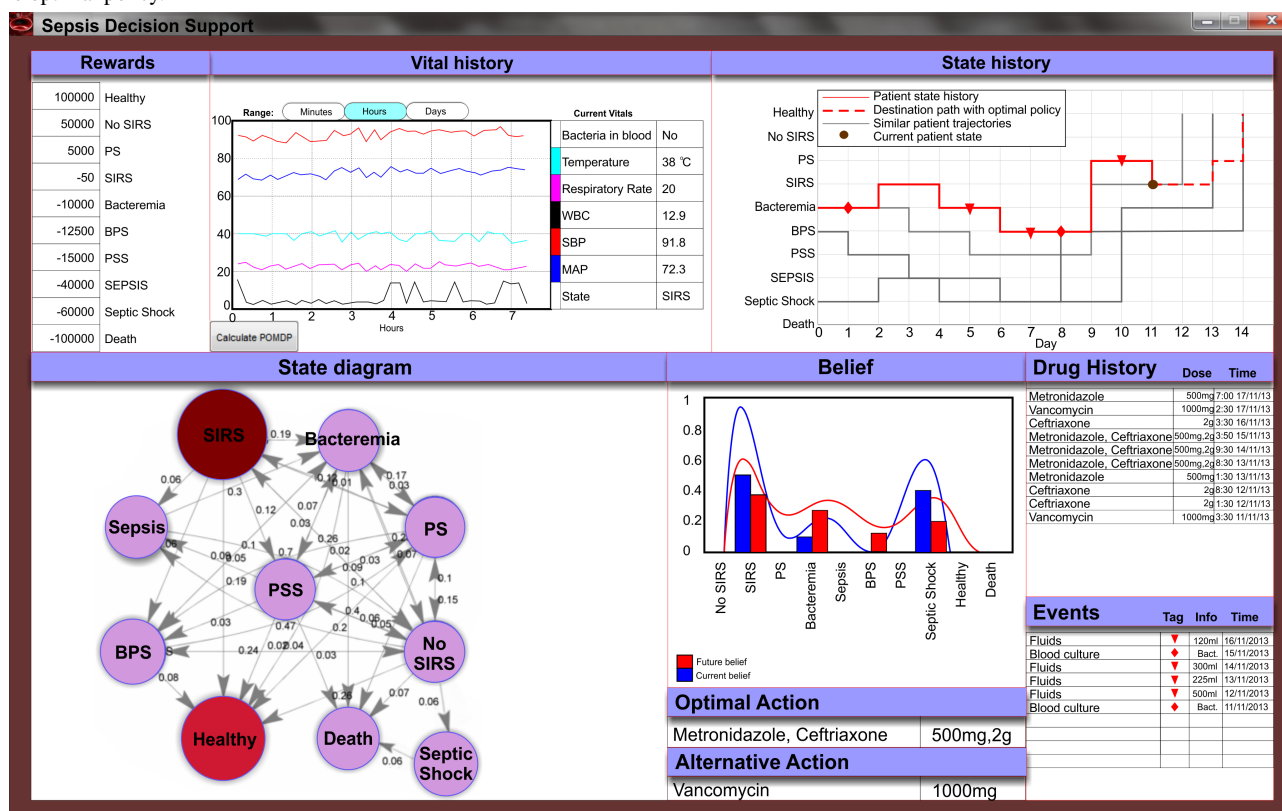
A policy is one or more actions that are followed. Each antibiotic combination is considered a possible action. A total of 48 antibiotics have been included in the patients' EHR that we analyzed. Here we consider the top five more frequently used antibiotics (Vancomycin, Cefepime, Metronidazole, Ceftriaxone, and Meropenem) and all their possible combinations, plus one more combination that encapsulates all other antibiotics that may have been used. This results in 32 possible combinations, each one defining a different action of the five most frequently used antibiotics. Vancomycin is a glycopeptide antibiotic that inhibits the cell wall synthesis of gram-positive bacteria, although it is avoided due to its nephrotoxicity and ototoxicity. Both Cefepime and Ceftriaxone are cephalosporin antibiotics that have activity against both Gram-negative and Gram-positive bacteria and they are especially used to treat moderate-severe pneumonia. Cefepime is also used to treat infections caused by

multi-drug resistant microbial strains. Metronidazole is a nitroimidazole antibiotic that is used particularly for anaerobic bacteria and some protozoa. Meropenem is an ultra-broad spectrum antibiotic and beta-lactam that inhibits bacterial wall synthesis. Combinations of all states with all possible actions comprise the state-action space for our sepsis model, [Figure 1](#) (b, left side).

## Rewards

Reward/cost values have been provided empirically by physicians, based on the severity of each state. These are, from best to worse: Healthy (100,000), No SIRS (50,000), Probable Sepsis (PS, 5000), SIRS (-50), Bacteremia (-10,000), Bacteremia Possible Sepsis (BPS, -12,500), Possible Septic Shock (PSS, -15,000), Sepsis (-40,000), Septic Shock (-60,000), Death (-100,000). This information is also depicted in [Figure 3](#), first panel.

**Figure 3.** A graphical user interface (GUI) for the clinical decision support system. Physicians have access to real-time and historical vital history (upper left), as well as state history (upper right) for a given patient. The state history displays events, most likely path if optimal policy is to be adopted and the past trajectories of the top three patients that had similar profiles (ie, past states and vitals) to the current patient. The state diagram (bottom left) depicts the state transition probabilities and its updates based on the current state. The GUI also displays the belief distribution, optimal action given the current belief and the second best alternative action. Drug history and significant events are displayed (bottom right). The tool updates values automatically with new data, although the user has the flexibility to revise the desired reward values (upper left) and then manually trigger a recalculation of the optimal policy.



## Calculation of Transition and Observation Probabilities

Transition probabilities are calculated as the frequency of State-Action-Next State patterns in each State-Action combination. Similarly, observation probabilities are calculated based on the frequency of observing vital value (in bins) combinations in any given state. In order to extract the probability of the observation combinations, we fit the distribution that best describes the data for each vital and state and then divide it in five non-overlapping, equally-sized bins. The number of bin was selected such that sufficient data (>10 samples) will be present in each bin. For each state, the distributions of the five vitals are divided in five equal parts across their min-max range (each assigned a probability from the distribution) with blood test being modeled with a binary variable (presence/absence of bacteria). This leads to 6250 combinations ( $2 \cdot 5^5$ ).

## POMDP Formulation

A POMDP is defined as a 8-tuple  $(S, A, Z, T, O, R, b_0, \gamma)$ , where  $S$  is a finite set of states,  $A$  is a finite set of actions,  $T: S \times A \times$

$S \rightarrow P(S)$  is the state transition function,  $T(s, a, s')$  denotes the probability  $P(s' | s, a)$  of reaching state  $s'$  from state  $s$  by taking action  $a$ ,  $R: S \times A \rightarrow \mathbb{R}$  is the reward function,  $P(s', a)$  denotes the immediate reward of executing action  $a$  in state  $s$ ,  $\gamma \in [0, 1]$  is the discount factor,  $Z$  is the finite set of observations,  $O: S \times A \rightarrow P(Z)$  is the observation function,  $O(s, a, z)$  denotes the probability  $P(z | s, a)$  of perceiving observation  $z$  when taking action  $a$  and arriving in state  $s$ ,  $b_0$  is the initial state probability distribution,  $b$  is the state probability distribution and  $b_0(s)$  denotes the probability of starting in state  $s$ . A policy  $\pi$  for a POMDP problem is defined as  $\pi(b) \rightarrow a$  and its value is the expected cumulative discounted reward that we will receive if we perform actions  $a$  when we have belief  $b$ . The policy that maximizes this cumulative value is called the *optimal policy*. We used a POMDP model training methodology that is based on a Monte Carlo approximation for solving the Value Iteration method over a sampling belief space, Figure 1 (b, right side). Value iteration calculates the value of each state by solving the Bellman equations (Figure 4).

If we have belief  $b(s)$  of being in state  $s$ , perform action  $a$  and we observe  $z$ , the updated belief for being in state  $b(s')$  is shown in Figure 1 (c) and Figure 5.

**Figure 4.** Value iteration calculates the value of each state by solving the Bellman equations.

$$V_{k+1}(b) = \max_{a \in A} \left[ \sum_{s \in S} b(s) R(s, a) + \gamma \sum_{z \in Z} P(z|b, a) V_k(\tau(b, a, z)) \right]$$

$$\pi_{k+1}(b) = \operatorname{argmax}_{a \in A} \left[ \sum_{s \in S} b(s) R(s, a) + \gamma \sum_{z \in Z} P(z|b, a) V_k(\tau(b, a, z)) \right]$$

**Figure 5.** Belief  $b(s)$  of being in state.

$$b(s') = \tau(b(s), a, z) = P(s'|b, a, z) = \frac{O(s', a, z) \sum_s T(s, a, s') b(s)}{P(z|b, a)}$$

## POMDP Statistical Evaluation

We use *Perseus* [32], a randomized point-based value iteration algorithm to extract the optimal policy with default settings. We use 5-fold cross-validation to evaluate the generalization error of the policy, Figure 1 (d). For a specific transition, a CDSS-derived policy is deemed as followed if one or more of the antibiotics in the policy have been administered to the patient. Comparison is performed between treatments that are in agreement with the CDSS-derived policy for each transition and those that are not (non CDSS-derived policies), measuring the percentage of transitions that lead to a better/same/worse state in both cases. At a trajectory level, we compare what percentage of trajectories move to a better state on average with and without the CDSS-derived policy. The robustness of the algorithm to perform in the context of reduced data is tested by splitting the patient data into 3 parts, where only the 2 out of the 3 parts are used for training. The algorithm is trained on randomly selected subsets of the 2/3 training data that correspond to the desired percentage of the total training set. The performance is always evaluated in the same testing set (ie 1/3 of the total dataset) and the whole process is repeated 10 times to reduce bias, Figure 1 (d). All  $P$  values are calculated using hypergeometric distribution with multiple hypothesis correction (Benjamini-Hochberg).

## Length-of-Stay and Mortality Inference

For the length of stay (LOS) prediction, we first split the patients in equally-sized bins, based on their LOS distribution, thus maximizing entropy and to avoid bias on the training data. Then we used Support Vector Machines (SVM) with different kernel functions [33,34]. In SVM classification, the training feature vectors are mapped to a higher dimension space, in which the SVM determines a linearly separating hyperplane given by a maximal margin [35]. In order to predict survivability results given a patient's vital signs, an SVM classification method was used, taking into account five features (temperature, respiratory rate, WBC, MAP, and lactate levels) and the mortality state of each patient, for both binary and multi-class classification. For the latter, we consider both combinatorial pairwise and one-versus-all schemes [36], where we observed no significant

difference in our results. We evaluated the performance of the classifier by performing cross-validation (CV) and calculating the Receiver-Operator-characteristic (ROC) curves, confidence intervals (CI), and the area-under-the-curve (AUC).

## Results

### Optimal Policy Prediction by Data-driven Machine Learning Approaches

We performed 5-fold cross validation (CV) to evaluate the generalization error of our approach, with similar results across all folds, Figure 6 (a) and Tables 2 and 3. The CDSS-derived optimal policies in each transition led in significantly more occurrences to better states than when the treatment that was followed by the physicians was not in agreement with the CDSS-derived policy (49% of transitions to better states when the CDSS-derived policy was followed vs 37% otherwise;  $P=1.3e-13$ ). Interestingly, when non-CDSS policies were used, patients tend to stay in the same condition (35% in non-CDSS policies vs 25% in CDSS policies,  $P=5.1e-13$ ) while the difference between CDSS-derived policies and non-CDSS policies is not statistically significant for transitions to a worse condition (28% in non-CDSS policies vs 26% in CDSS policies,  $P=4.2e-1$ ). We then analyzed each patient trajectory independently, to estimate the number of transitions within a trajectory that lead to better states, with and without following the policy. Results show that when the optimal policy prediction (ie, the policy that maximizes the expected cumulated reward, as defined in the methods section) is followed, there exists a significant shift to trajectories that have more than 90% of their transitions leading to a better state, Figure 6 (b). When the policy was followed, 387 patients (25.9%) have 90% of their transitions to better states and 503 patients (33.7%) patients had 90% of their transitions to worse states ( $P=4.0e-06$ ), while in the non-policy cases, these numbers are 192 (12.9%) and 764 (51.2%) patients ( $P=4.6e-117$ ), respectively. Furthermore, the percentage of transitions within a trajectory that lead to a better or better/same state are significantly higher by following the policy than for non-policy cases. Indeed, 605 versus 344 patients ( $P=8.6e-25$ ) have 90% of their transitions to a better state with versus without following a CDSS-derived policy, Figure 1 (a).

This result was observed in all five runs of the 5-fold CV and on the full dataset, hence it holds for different data distributions.

Next, we evaluated the robustness of the POMDP framework to decreasing sets of training data. To perform this analysis, we iteratively reduced the training set and we evaluated in the same testing data set (see Methods). Results demonstrate the method is robust to decreasing amount of data in the set, [Figure 6](#) (c) as well as [Figures 3](#) and [7](#), mainly due to the significant overlap of the various antibiotics in each combination proposed by the optimal policy. To gain more insight on how the policy changes by decreasing the training set we constructed a comprehensive map of the optimal policy for each state, [Figure 6](#) (d). The resulting map provides the CDSS-derived drug combination that led to more favorable outcomes in each state, [Table 3](#). It is

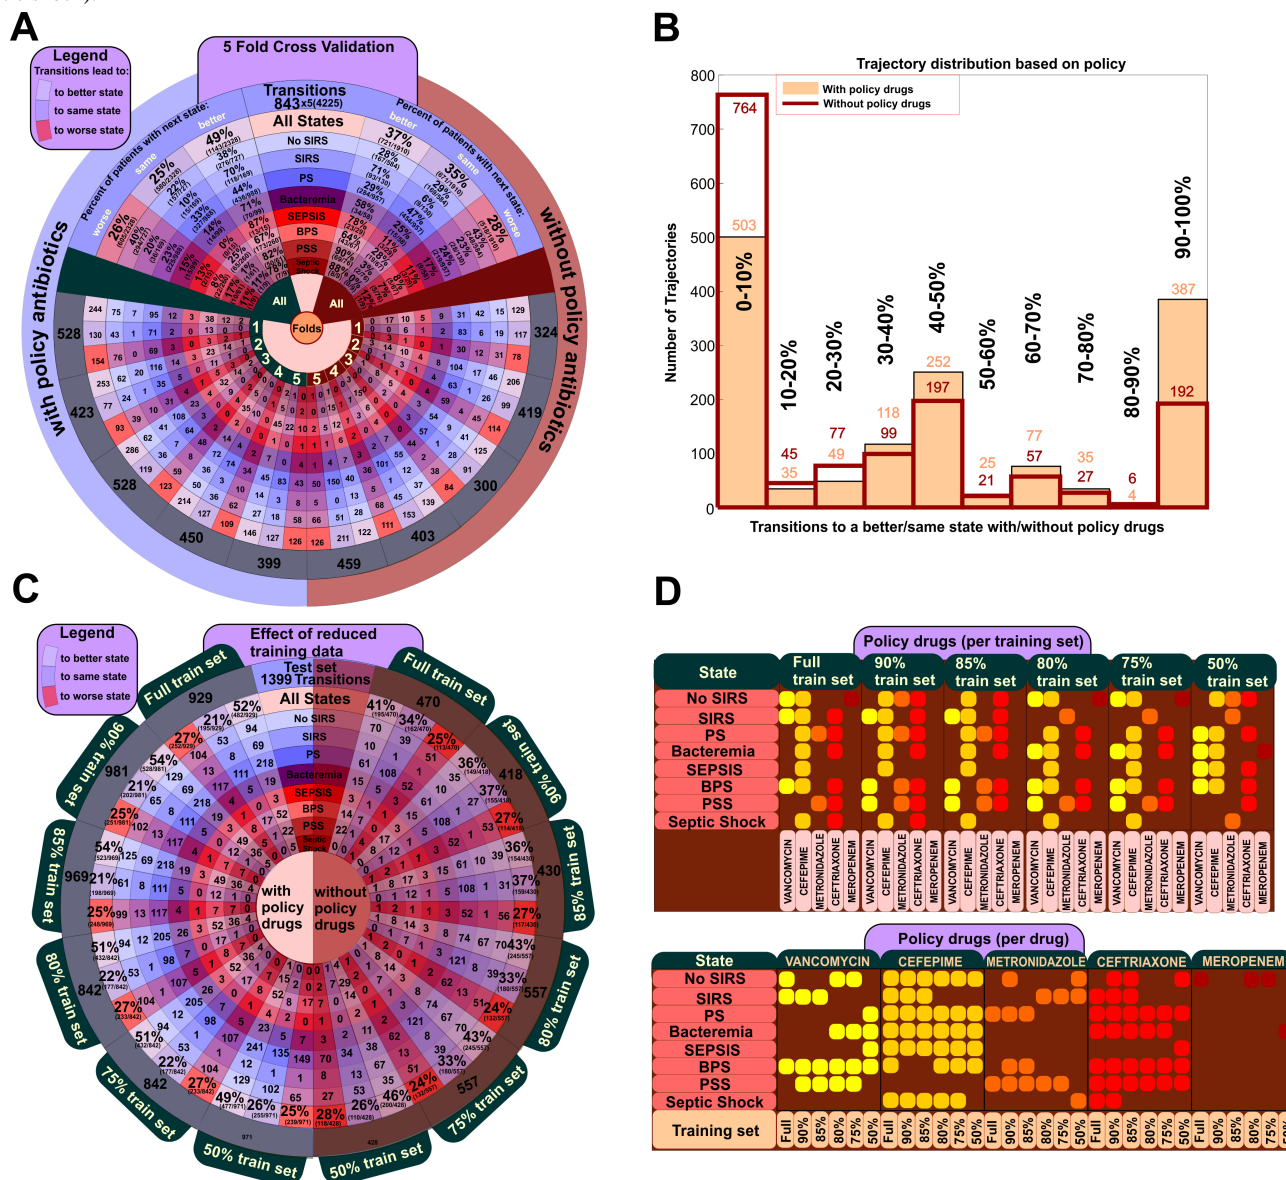
important to note, that these policies correspond to a definite knowledge that the patient is that specific state (belief/probability of 1), which is almost never the case as his/her prior history (previous states, clinical information, etc) shapes the belief distribution across all states at any given time. Additionally, the depicted drug combinations are associated with a better outcome overall, and not that are necessarily the optimal combination under any condition when a patient is in that specific state, since potent drug combinations are used to more severe cases, which have a higher probability to transition to a worse state. The optimal decision for any state will ultimately be a function of all observations (vitals, blood results, etc). Their associations will depend on the structure of the data that were used for the CDSS training.

**Table 3.** Optimal Policy based on from the POMDP CDSS tool. Note that this result assumes that current state is known and given in the state column (Belief/Probability of 1).

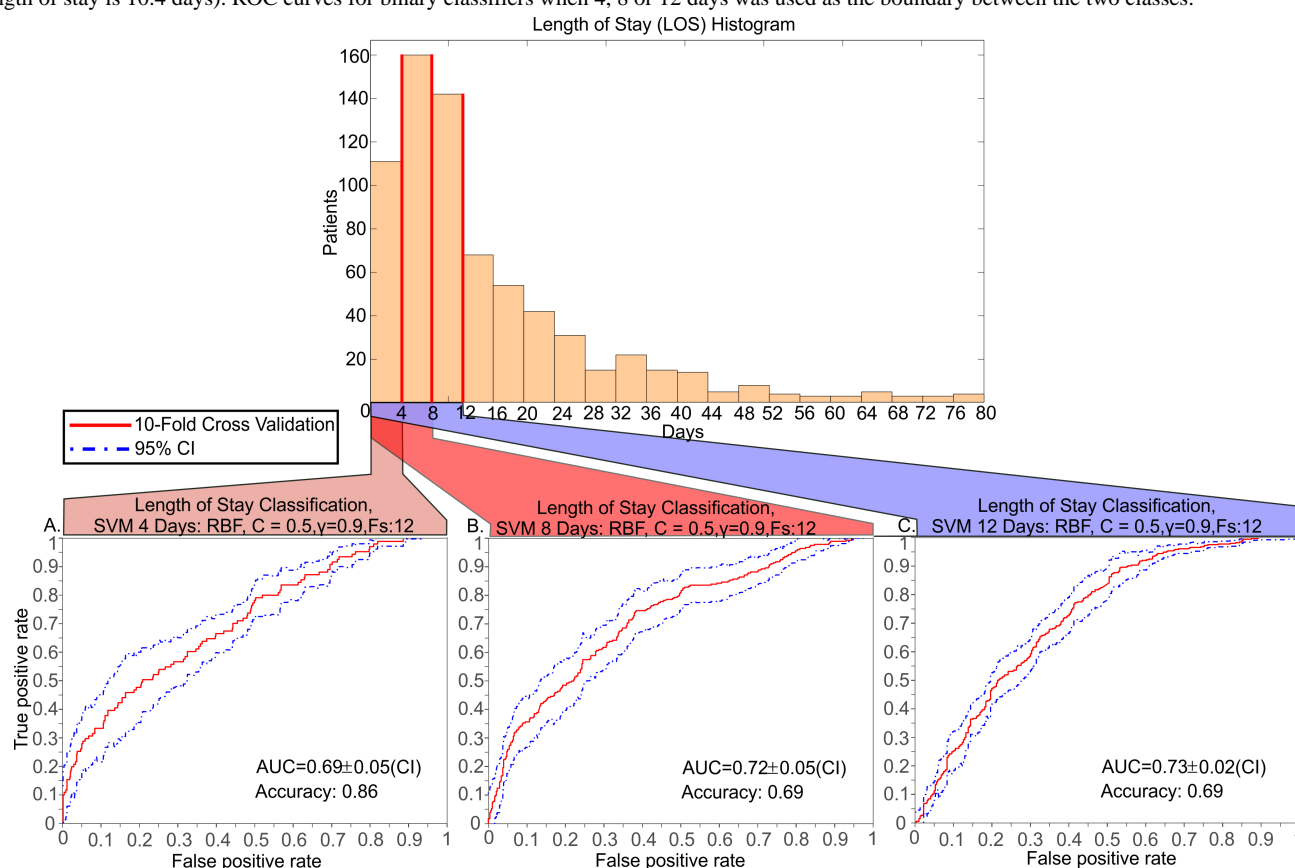
State	Drugs
SIRS	CEFEPIME,METRONIDAZOLE,CEFTRIAXONE
PS	METRONIDAZOLE,CEFTRIAXONE,MEROPENEM
Bacteremia	CEFEPIME,CEFTRIAXONE
Sepsis	CEFTRIAXONE
BPS	VANCOMYCIN,CEFEPIME,CEFTRIAXONE
PSS	METRONIDAZOLE,CEFTRIAXONE
Septic Shock	CEFTRIAXONE



**Figure 6.** Performance and Robustness of the POMDP Clinical Decision Support System for Sepsis. (A) 5-fold cross-validation results depict the performance of each fold with (left, blue) and without (right, brown) using the policy-proposed antibiotic combination. Each cell contains the number of transitions (4225 total transitions; 843 transitions per test fold) that lead to worse, equal or better states in each case. A state-specific percentage across all folds allows for comparison between the different policy strategies. (B) Number of patient trajectories vs. the percentage of their transitions that lead to a better state. (C) Dependency of CDSS performance on data size based on stratified reduction of the dataset. Outcome is shown for policy-proposed antibiotic combinations (left) and all other combinations (right) for different states (D) Changes in antibiotic combinations proposed in calculated optimal policy as a function of data-size reduction. Each row is a state and each column represents a drug-training set combination. The two tables depict which drug combinations were found to lead to better outcomes when in the perspective state, in the general case. States are as defined in the Methods sections, with three states denoting uncertainty due to missing data (PS: probable sepsis; BPS: bacteremia, probable sepsis; PSS: probable septic shock).



**Figure 7.** Predicting a patient's Length of Stay (LOS). Histogram of the LOS for the 745 patients in the database that have a complete record (median length of stay is 10.4 days). ROC curves for binary classifiers when 4, 8 or 12 days was used as the boundary between the two classes.



## Mortality Prediction

Clinical outcome was found to be accurately classified by using Support Vector Machines (see Methods). In order to predict patient mortality, we used the five features (temperature, respiratory rate, WBC, MAP, and lactate levels) as well as the final outcome for each patient. This led to a dataset of 745 patients out of the 1492 total for whom all six variables were available. The five measurement variables were summarized by their mean and standard deviation (STD) across the trajectory of each patient [37]. Principal Component Analysis (PCA) was also used to assess whether linear transformation of the feature space and dimensionality reduction can be achieved in this case [38]. A filter method was applied using the Area-Under-the-Curve (AUC) of the Receiver Operating Characteristic (ROC) curves as a ranking criterion [39]. Radial Basis Function (RBF) kernel was used for the SVM training in a 10-fold cross validation scheme. The positive/negative classes for mortality prediction are defined as alive/deceased, respectively. Mortality classification of 745 patients has an AUC of 0.70 (SD 0.04; 95% CI) and accuracy of 0.72, Figure 2 (a). When the test set focuses on the 170 patients with sepsis group, it achieves an AUC of 0.61 (SD 0.14) and accuracy of 0.71, Figure 2 (a). For the group of patients with a better-to-all transition ratio larger than 0.9 (302 patients out of 745) the AUC is higher at 0.74 (SD 0.07) and the accuracy is 0.81, Figure 2 (b), while AUC drops to 0.58 (SD 0.10) for patients with a ratio of better-to-all smaller than 0.1. Comparison of these results argue that the trained classifier performs better and is more accurate in cases where the proposed policy has been used and

its discrimination power is higher when these policies lead to a favorable outcome.

## Length of Stay Prediction

To predict the length-of-stay (LOS), we trained SVM classifiers with two extra features that were found to be informative: the occurrences of a positive blood culture and the number of policy drug administration during a patient's stay. We then defined and classified the patients in two classes, based on the length of stay. The threshold for that discrimination was driven by the median length of stay in the hospital (10.4 days) and hence we selected thresholds of 4, 8 and 12 days. A 10-fold cross validation scheme was used for the evaluation of the classifier. The AUC of the classifiers were 0.69 to 0.73 with small deviations in the CI (0.02-0.05) and accuracies from 0.69 to 0.82, Figure 7. Multi-class classification for multiple length-of-stay bins (0-3, 3-6, 6-12, 12+ days) had similar results, although the AUC drops to 0.53 when predicting the two intermediate (3-6 and 6-12) classes, Figure 6.

## Discussion

In this work, we used the EHR of 1492 patients to build a decision support tool and predictive classifiers for patients with sepsis. Despite the fact that the dataset was limited in both number of patients and features available, the CDSS methodology resulted in data-driven policies that led to significantly improved patient outcomes. Similarly, we demonstrated that time-stamped EHR observational data, such as patient vitals and blood results, can be used to predict

mortality and length of stay intervals, with increased accuracy and discriminative performance.

Given the vast combinatorial space of treatments and outcomes, one of the main challenges in the development of a statistical decision support tool is the definition of states and actions in a way that is both clinically relevant and computationally feasible. To create a framework that balances these trade-offs, we used expert knowledge and statistical methods to efficiently represent clinical cases within the POMDP framework, while at the same time making sure that each state-action combination has sufficient data for model training and testing. As the size of clinical databases scales up, an automated state and action definition technique can be applied, which might lead to interesting insights on what is medically relevant in each case. Our robustness analysis argues that a POMDP-based tool is quite robust even for small sample sizes and it remains to be seen the generalization boundaries of such approach for larger, integrative datasets, more sophisticated state-action spaces with additional features, and complex clinical histories. As with any data-driven predictive approaches, the generalization error and applicability of the results is dependent on the extend a model can capture the real state and action space, as well as the various biases that arise due to limited sample sizes, data quality and precision. For instance, differences in patient populations, microbiological resistance patterns of the wards, anti-infective pre-treatment of patients, administration of drugs (eg, vasopressors) or treatments (e.g, ventilator support) that are currently not captured limit the applicability of this study as they can substantially change the proposed policies and actions. To address these issues, the methods that are proposed here can be applied in larger datasets that can support a more extensive modeling for states, actions and observables, while correcting for possible biases across different attributes.

This initial study paves the way for several interesting directions towards a predictive CDSS for sepsis treatment. In addition to the SIRS criteria and the indication of an infection, it would be useful to take into account possible dysfunctional organs, a set of information that was not present in the database that we used here. As such, we can define nine states for organ dysfunction: absent, Respiratory, Coagulation, Liver, CNS, Renal, Metabolic, Cardiovascular, Multi-organ dysfunction (the latter defined as two or more organ failures). The Sequential Organ Failure Assessment (SOFA) criteria and score [40] can be used for this

purpose. The action space can similarly be expanded to include several other actions that are important for sepsis treatment, such as the administration and dosage of IV fluids, vasoactive medications, initiation of mechanical ventilation, oxygen therapy, hemodialysis, the use of sepsis order set, and other admission and/or transfer decisions. To this end, a more extensive dataset, both in terms of features and patients is crucial so that the state/action combinatorial space will have adequate training samples. Furthermore, it is important to extend the number of composite features that one will investigate for complex traits, as in this study we only considered up to seven features. Such extension would likely lead to more accurate predictions regarding mortality and patient's length-of-stay. From a technical perspective, it would be important to work towards an algorithmic framework that can distinguish patients that have reached a state from different trajectories, as the optimal treatment in each of these scenarios can be substantially different. Although this would violate the Markov property in the general case, one can investigate finite-memory models that can accommodate such setting.

Finally, an important aspect of any CDSS tool is an intuitive and interactive visualization of the patient status, past history and decision space. For this work, we have developed an interactive graphical user interface (GUI) that is connected with the POMDP solver and the database and can display the vitals, drug and state history, state belief, state diagram with all possible transitions and their probabilities, as well as the optimal/near-optimal actions given the current belief for the patient's state, [Figure 3](#). The user can also define the rewards for each state and re-calculate the POMDP-derived optimal policy. In addition, the patient's trajectory is compared on-the-fly with other patient's trajectories in the database for comparison and visualization of potential outcomes. Novel visualization methods and interactive tools, such as head-mounted displays that are non-obtrusive are promising candidates to pair with the proposed CDSS, both as display and acquisition devices. The ultimate goal should be to use real time learning and analysis obtained from readily available EMR data, to warn the clinician of important changes in patient "states" and the need for new "actions" to improve the outcome of severe sepsis patients. As such, the integration of "big data" analytics with ubiquitous computing has the potential to revolutionize emergency and intensive care medicine as we know it.

## Acknowledgments

We would like to thank Dr Jeffrey Green, Dr Hien Nguyen and Dr Jason Adams for the helpful discussions related to the clinical relevance of the CDSS tool, Dr Aaron Bair for his advice on appropriate clinical environments for testing the CDSS tool, Dr Sergey Levine and Dr Kee-Eung Kim for their advice related to POMDP methods. This work was supported by the Center for Information Technology Research in the Interest of Society (CITRIS) through seed grant #2469085 and the National Center for Advancing Translational Sciences of the National Institutes of Health (grant #UL1 TR000002) to IT.

## Authors' Contributions

AT performed all experiments and analyzed the data. TA advised on the clinical approach and relevance of the study. IT advised on all computational aspects and supervised the data analysis of the study. AT, TA, and IT wrote the manuscript.

## Conflicts of Interest

None declared.

## Multimedia Appendix 1

Supplementary Tables and Figures.

[[PDF File \(Adobe PDF File\), 453KB - medinform\\_v3i1e11\\_app1.pdf](#)]

## References

- Wood KA, Angus DC. Pharmacoeconomic implications of new therapies in sepsis. *Pharmacoeconomics* 2004;22(14):895-906. [Medline: [15362927](#)]
- Hall MJ, Williams SN, DeFrances CJ, Golosinskiy A. Inpatient Care for Septicemia or Sepsis: A Challenge for Patients and Hospitals. In: NCHS Data Brief. Hyattsville, MD: National Center for Health Statistics; 2011.
- Sepsis facts. Sepsis facts. URL: <http://world-sepsis-day.org/.../?MET=HOME&vLANGUAGE=EN> [accessed 2015-02-09] [WebCite Cache ID 6WE7B0X2g]
- Gunter TD, Terry NP. The emergence of national electronic health record architectures in the united states and australia: Models, costs, and questions. *Journal of Medical Internet Research* 2005;7(1).
- Fonkych K, Taylor R. The state and pattern of health information technology adoption. Rand Corporation 2005;409.
- Kellermann AL, Jones SS. What it will take to achieve the as-yet-unfulfilled promises of health information technology. *Health Aff (Millwood)* 2013 Jan;32(1):63-68. [doi: [10.1377/hlthaff.2012.0693](#)] [Medline: [23297272](#)]
- Ouali A, Ramdane-Cherif Z, Ramdane-Cherif A, Levy N, Krebs M. Agent paradigm in clinical large-scale data mining environment. 2003 Presented at: 2nd IEEE International Conference on Cognitive Informatics (ICCI); August 18-20, 2003; London, UK p. 143-150.
- Catley C, Frize M. A prototype XML-based implementation of an integrated intelligent neonatal intensive care unit. 2003 Aug 18 Presented at: 4th International IEEE Conference on Information Technology Applications in Biomedicine; 2003; London, UK p. 322-325.
- Balter J, Labarre-Vila A, Ziébelin D, Garbay C. A knowledge-driven agent-centred framework for data mining in EMG. *C R Biol* 2002 Apr;325(4):375-382. [Medline: [12161917](#)]
- Clermont GC, Angus DC, DiRusso SM, Griffin M, Linde-Zwirble WT. Predicting hospital mortality for patients in the intensive care unit: a comparison of artificial neural networks with logistic regression models. *Crit Care Med* 2001 Feb;29(2):291-296. [Medline: [11246308](#)]
- Foster D, McGregor C, El-Masri S. A survey of agent-based intelligent decision support systems to support clinical management and research. 2005 Jul 25 Presented at: Fourth International Joint Conference on Autonomous Agents and Multiagent Systems; 2005; Utrecht, Netherlands p. 16-34.
- Wears RL, Berg M. Computer technology and clinical work: still waiting for Godot. *JAMA* 2005 Mar 9;293(10):1261-1263. [doi: [10.1001/jama.293.10.1261](#)] [Medline: [15755949](#)]
- Bellman RE. A markovian decision process. In: No. P-1066. SANTA MONICA CA: RAND CORP; 1957.
- Kaelbling LP, Littman ML, Cassandra AR. Planning and acting in partially observable stochastic domains. *Artificial intelligence* 1998;101(1):99-134.
- Argall BD, Chernova S, Veloso M, Browning B. A survey of robot learning from demonstration. *Robotics and Autonomous Systems* 2009;57(5):469-483.
- Howard RA. Dynamic programming and markov processes. In: Technology Press. Boston, MA: Massachusetts Institute of Technology; 1960.
- Watkins CJ, Dayan P. Q-learning. *Mach Learning* 1992;8(3-4):279-292.
- Alagoz O, Maillart LM, Schaefer AJ, Roberts MS. The optimal timing of living-donor liver transplantation. *Management Science* 2004;50(10):1420-1430.
- Alagoz O, Maillart LM, Schaefer AJ, Roberts MS. Determining the acceptance of cadaveric livers using an implicit model of the waiting list. *Oper Res* 2007;55(1):24-36.
- Alagoz O, Maillart LM, Schaefer AJ, Roberts MS. Choosing among living-donor and cadaveric livers. *Management Science* 2007;53(11):1702-1715.
- Shechter SM, Bailey MD, Schaefer AJ, Roberts MS. The optimal time to initiate HIV therapy under ordered health states. *Oper Res* 2008;56(1):20-33.
- Maillart LM, Ivy JS, Ransom S, Diehl K. Assessing dynamic breast cancer screening policies. *Oper Res* 2008;56(6):1411-1427.
- Faissol DM, Griffin PM. Timing of testing and treatment of hepatitis C and other diseases. In: Proceedings. 2007 Presented at: INFORMS international meeting; July 2007; Puerto Rico p. 11.
- Denton BT, Kurt M, Shah ND, Bryant SC, Smith SA. Optimizing the start time of statin therapy for patients with diabetes. *Med Decis Making* 2009;29(3):351-367. [doi: [10.1177/0272989X08329462](#)] [Medline: [19429836](#)]



25. Alagoz O, Hsu H, Schaefer AJ, Roberts MS. Markov decision processes: a tool for sequential decision making under uncertainty. *Med Decis Making* 2010;30(4):474-483 [FREE Full text] [doi: [10.1177/0272989X09353194](https://doi.org/10.1177/0272989X09353194)] [Medline: [20044582](https://pubmed.ncbi.nlm.nih.gov/20044582/)]
26. Hauskrecht M, Fraser H. Modeling treatment of ischemic heart disease with partially observable Markov decision processes. In: *Proc AMIA Symp.* 1998 Presented at: *Proc AMIA Symp*; 1998 p. 538-542 URL: <http://europepmc.org/abstract/MED/9929277> [WebCite Cache]
27. Hauskrecht M, Fraser H. Planning treatment of ischemic heart disease with partially observable Markov decision processes. *Artif Intell Med* 2000 Mar;18(3):221-244. [Medline: [10675716](https://pubmed.ncbi.nlm.nih.gov/10675716/)]
28. Ayer T, Alagoz O, Stout NK. A POMDP approach to personalize mammography screening decisions. *Operations Research* 2012 Jun 12;60(5):1019-1034.
29. Kreke JE. Modeling disease management decisions for patients with pneumonia-related sepsis. Pittsburgh, PA: University of Pittsburgh; 2007.
30. Kreke JE, Bailey MD, Schaefer AJ, Angus DC, Roberts MS. Modeling hospital discharge policies for patients with pneumonia-related sepsis. *IIE Transactions* 2008;40(9):853-860.
31. Bone RC, Balk RA, Cerra FB, Dellinger RP, Fein AM, Knaus WA, et al. Definitions for sepsis and organ failure and guidelines for the use of innovative therapies in sepsis. *Crit Care Med* 1992;20(6):864-874.
32. Spaan MTJ., Vlassis N.. Perseus: Randomized Point-based Value Iteration for POMDPs. *Journal of Artificial Intelligence Research* 2005 Jul 01;24(1):195-220 [FREE Full text] [WebCite Cache]
33. Wang SL, Wu F, Wang BH. Prediction of severe sepsis using SVM model. *Adv Exp Med Biol* 2010;680:75-81. [doi: [10.1007/978-1-4419-5913-3\\_9](https://doi.org/10.1007/978-1-4419-5913-3_9)] [Medline: [20865488](https://pubmed.ncbi.nlm.nih.gov/20865488/)]
34. Tang CH, Middleton PM, Savkin AV, Chan GS, Bishop S, Lovell NH. Non-invasive classification of severe sepsis and systemic inflammatory response syndrome using a nonlinear support vector machine: a preliminary study. *Physiol Meas* 2010 Jun;31(6):775-793. [doi: [10.1088/0967-3334/31/6/004](https://doi.org/10.1088/0967-3334/31/6/004)] [Medline: [20453293](https://pubmed.ncbi.nlm.nih.gov/20453293/)]
35. Vapnik V. *Statistical Learning Theory*. New York: John Wiley & Sons; 1998.
36. Hsu CW, Lin CJ. A comparison of methods for multiclass support vector machines. *IEEE Trans Neural Netw* 2002;13(2):415-425. [doi: [10.1109/72.991427](https://doi.org/10.1109/72.991427)] [Medline: [18244442](https://pubmed.ncbi.nlm.nih.gov/18244442/)]
37. Saeyns Y, Inza I, Larranaga P. A review of feature selection techniques in bioinformatics. *Bioinformatics* 2007;23:A.
38. Jolliffe IT. *Principle Component Analysis*. New York: Springer; 2005.
39. Wang R, Tang K. Feature selection for maximizing the area under the ROC curve. 2009 Dec 06 Presented at: *IEEE International Conference on Data Mining Workshops*; 2009; Florida, USA p. 400-405.
40. Vincent JL, Moreno R, Takala J, Willatts S, De Mendonça A, Bruining H, et al. The SOFA (Sepsis-related Organ Failure Assessment) score to describe organ dysfunction/failure. On behalf of the Working Group on Sepsis-Related Problems of the European Society of Intensive Care Medicine. *Intensive Care Med* 1996 Jul;22(7):707-710. [Medline: [8844239](https://pubmed.ncbi.nlm.nih.gov/8844239/)]

## Abbreviations

**AUC:** area under the curve  
**BPS:** bacteremia probable SIRS  
**CDSS:** clinical decision support system  
**CI:** confidence interval  
**CV:** cross-validation  
**EHR:** electronic health records  
**ICU:** intensive care unit  
**LOS:** length of stay  
**MAP:** mean arterial pressure  
**MDP:** Markov Decision Process  
**PCA:** principal component analysis  
**POMDP:** Partially Observable Markov Decision Process  
**PS:** probable SIRS  
**PSS:** probable septic shock  
**RBF:** radial basis function  
**ROC:** receiver-operator-characteristic  
**RR:** respiratory rate  
**SBP:** systolic blood pressure  
**SIRS:** Systemic Inflammatory Response Syndrome  
**SOFA:** sequential organ failure assessment  
**SVM:** support vector machines  
**UCDHS:** University of California Davis Health System  
**WBC:** white blood count



*Edited by G Eysenbach; submitted 03.04.14; peer-reviewed by G Rättsch, S Tafelski; comments to author 02.07.14; revised version received 26.08.14; accepted 11.10.14; published 24.02.15*

*Please cite as:*

*Tsoukalas A, Albertson T, Tagkopoulos I*

*From Data to Optimal Decision Making: A Data-Driven, Probabilistic Machine Learning Approach to Decision Support for Patients With Sepsis*

*JMIR Med Inform 2015;3(1):e11*

URL: <http://medinform.jmir.org/2015/1/e11/>

doi: [10.2196/medinform.3445](https://doi.org/10.2196/medinform.3445)

PMID: [25710907](https://pubmed.ncbi.nlm.nih.gov/25710907/)

©Athanasios Tsoukalas, Timothy Albertson, Ilias Tagkopoulos. Originally published in JMIR Medical Informatics (<http://medinform.jmir.org>), 24.02.2015. This is an open-access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/2.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work, first published in JMIR Medical Informatics, is properly cited. The complete bibliographic information, a link to the original publication on <http://medinform.jmir.org/>, as well as this copyright and license information must be included.