nature medicine

SUPPLEMENTARY INFORMATION

https://doi.org/10.1038/s41591-018-0213-5

In the format provided by the authors and unedited.

The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care

Matthieu Komorowski[1,2,3], Leo A. Celi[3,4], Omar Badawi[3,5,6], Anthony C. Gordon[1*] and A. Aldo Faisal[2,7,8,9*]

[1]Department of Surgery and Cancer, Imperial College London, London, UK. [2]Department of Bioengineering, Imperial College London, London, UK. [3]Laboratory of Computational Physiology, Harvard-MIT Division of Health Sciences & Technology, Cambridge, MA, USA. [4]Beth Israel Deaconess Medical Center, Boston, MA, USA. [5]Department of eICU Research and Development, Philips Healthcare, Baltimore, MD, USA. [6]Department of Pharmacy Practice and Science, University of Maryland, School of Pharmacy, Baltimore, MD, USA. [7]Department of Computer Science, Imperial College London, London, UK. [8]Medical Research Council London Institute of Medical Sciences, London, UK. [9]Behaviour Analytics Lab, Data Science Institute, London, UK.
*e-mail: anthony.gordon@imperial.ac.uk; a.faisal@imperial.ac.uk

Supplementary Appendix

Supplement to: “The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care” by Komorowski et al.

Supplementary Tables

Table 1: Full description of the datasets

	MIMIC-III	eRI
Unique ICUs (N)	5 (NICU excluded)	128
Characteristics of hospitals, per number of ICU admissions (N, %)	Teaching tertiary hospital.	Non-teaching: 37,146 (47.0%) Teaching: 29,388 (37.2%) Unknown: 12,539 (15.9%)
Hospital location in the USA, per number of ICU admissions (N, %)	Boston, Massachusetts.	South: 32,878 (41.6%) Northeast: 15,280 (19.3%) Undocumented: 12,858 (16.3%) Midwest: 12,298 (15.6%) West: 5,759 (7.3%)
Type of ICUs (N, %) Medical-surgical ICU MICU CCU/CTICU/CSRU SICU/TICU Other	 - 8,131 (47.6%) 4,534 (26.5%) 4,418 (25.8%) -	 44,567 (56.4%) 11,191 (14.2%) 15,404 (19.5%) 5,544 (7.0%) 2,367 (3.0%)
Missing data after sample-and-hold	16.5%	9.3%
Unique ICU admissions (N)	17,083	79.073
Unique hospital admissions (N)	17,045	79,073
Unique patients (N)	14,493	79,073

Source of hospital admission (N, %)	ED: 7,620 (44.6%) Clinic referral: 3,990 (23.3%) Transfer from external hospital: 2,760 (16.2%) Physician referral: 2,572 (15.1%) Other: 141 (0.8%)	ED: 41,241 (52.2%) Undocumented: 17,544 (22.2%) Floor: 10,753 (13.6%) Transfer from external hospital: 3,533 (4.5%) Direct admission: 2,853 (3.6%) OR: 2,988 (3.8%) Other: 161 (0.2%)
Age, years (Mean, SD)	64.4 (16.9)	65.0 (16.7)
Male gender (N, %)	9,604 (56.2%)	40,949 (51.8%)
Premorbid status (N, %) Hypertension Diabetes CHF Cancer COPD/RLD CKD	 9,384 (54.9%) 4,902 (28.7%) 5,206 (30.5%) 1,803 (10.5%) 4,248 (28.7%) 3,087(18.1%)	 43,365 (54.8%) 25,290 (32.0%) 15,023 (19.0%) 11,807 (14.9%) 18,406 (23.3%) 14,553 (18.4%)
Primary ICD-9 diagnosis (N, %) Sepsis, including pneumonia Cardiovascular Other respiratory conditions Neurological Renal Others	 5,824 (34.1%) 5,270 (30.8%) 1,798 (10.5%) 1,590 (9.3%) 429 (2.5%) 2,172 (12.7%)	 41,396 (52.3%) 11,221 (14.2%) 9,127 (11.5%) 7,127 (9.0%) 1,454 (1.8%) 8,747 (11.1%)
Estimated time of onset of sepsis, after ICU admission, in hours (Median, IQR)	3.9 (-1.1 – 35.5)	1 (-0.7 – 2.8)
Initial OASIS (Mean, SD)	33.5 (8.8)	34.8 (12.4)
Initial SOFA (Mean, SD)	7.2 (3.2)	6.4 (3.5)

Procedures during the 72h of data collection: Mechanical ventilation (N, %) Vasopressors (N, %) Renal replacement therapy (N, %)	 9,362 (54.8%) 6,023 (35.3%) 1,488 (8.7%)	 39,115 (49.5%) 23,877 (30.2%) 6,071 (7.7%)
Fluid balance on admission documented (N, %)	9,317 (54.5%)	24,672 (31.2%)
Length of stay, days (Median, IQR)	3.1 (1.8 – 7)	2.9 (1.7 – 5.6)
ICU mortality	7.4%	9.8%
Hospital mortality	8.9%	16.4%
28-day mortality	11.3%	Not available
90-day mortality	18.9%	Not available

Supplementary Table 1. Description of the datasets. CCU: Coronary Care Unit; CHF: Congestive Heart Failure; CKD: Chronic Kidney Disease; COPD: Chronic Obstructive Pulmonary Disease; CSRU: Cardiac Surgery Recovery Unit; CTICU: Cardio-thoracic ICU; ED: Emergency Department; ICD-9: International Classification of Diseases version 9; IQR: Interquartile Range; MICU: Medical ICU; OASIS: Oxford Acute Severity of Illness Score; OR; Operating Room, RLD: Restrictive Lung Disease; SD: Standard Deviation; SICU: Surgical ICU; SOFA: Sequential Organ Failure Assessment; TICU: Trauma ICU.

Table 2: List of model features

Category	Items	Type	Available in MIMIC-III	Available in eRI
Demographics	Age	Cont.	+	+
	Gender	Binary	+	+
	Weight	Cont.	+	+
	Readmission to intensive care	Binary	+	+
	Elixhauser score (premorbid status)	Cont.	+	-
Vital signs	Modified SOFA*	Cont.	+	+
	SIRS	Cont.	+	+
	Glasgow coma scale	Cont.	+	+
	Heart rate, systolic, mean and diastolic	Cont.	+	+
	blood pressure, shock index	Cont.	+	+
	Respiratory rate, SpO_2	Cont.	+	+
	Temperature			
Lab values	Potassium, sodium, chloride	Cont.	+	+
	Glucose, BUN, creatinine	Cont.	+	+
	Magnesium, calcium, ionized calcium,	Cont.	+	+
	carbon dioxide	Cont.	+	+
	SGOT, SGPT, total bilirubin, albumin	Cont.	+	+
	Hemoglobin	Cont.	+	+
	White blood cells count, platelets	Cont.	+	+
	count, PTT, PT, INR			
	pH, PaO_2, $PaCO_2$, base excess,			
	bicarbonate, lactate, PaO_2/FiO_2 ratio			
Ventilation parameters	Mechanical ventilation	Binary	+	+
	FiO_2	Cont.	+	+
Medications and fluid balance	Current IV fluid intake over 4h	Cont.	+	+
	Maximum dose of vasopressor over 4h	Cont.	+	+
	Urine output over 4h	Cont.	+	+
	Cumulated fluid balance since	Cont.	+	+
	admission (includes preadmission data			
	when available)			
Outcome	Hospital mortality	Binary	+	+
	90-day mortality	Binary	+	-

Supplementary Table 2. Description of the variables included in the datasets. Cont.: continuous; INR: International Normalized Ratio; * Modified SOFA: SOFA based on values in the current 4h time step; PEEP: Positive End Expiratory Pressure; PT: Prothrombin Time; PTT: Partial Thromboplastin Time; SIRS: Systemic Inflammatory Response Syndrome; Shock index: systolic blood pressure/heart rate.

Table 3: Construction of the action space

Discretized action	IV fluids (mL/4 hours)		Vasopressors (mcg/kg/min)	
	Range	Median dose	Range	Median dose
1	0	0	0	0
2	]0-50]	30	]0-0.08]	0.04
3	]50-180]	85	]0.08-0.22]	0.13
4	]180-530]	320	]0.22-0.45]	0.27
5	>530	946	>0.45	0.68

Supplementary Table 1. Range and median doses of drugs in all discretized actions. Option 1 corresponds to "no drug given", and the remaining non-null doses were divided into 4 quartiles. The combination of the two drugs made up 5x5=25 possible actions. The median value in each dose bin represents the suggested dose mapping an optimal action.

Supplementary Figures

Figure 1: Patient inclusion diagrams

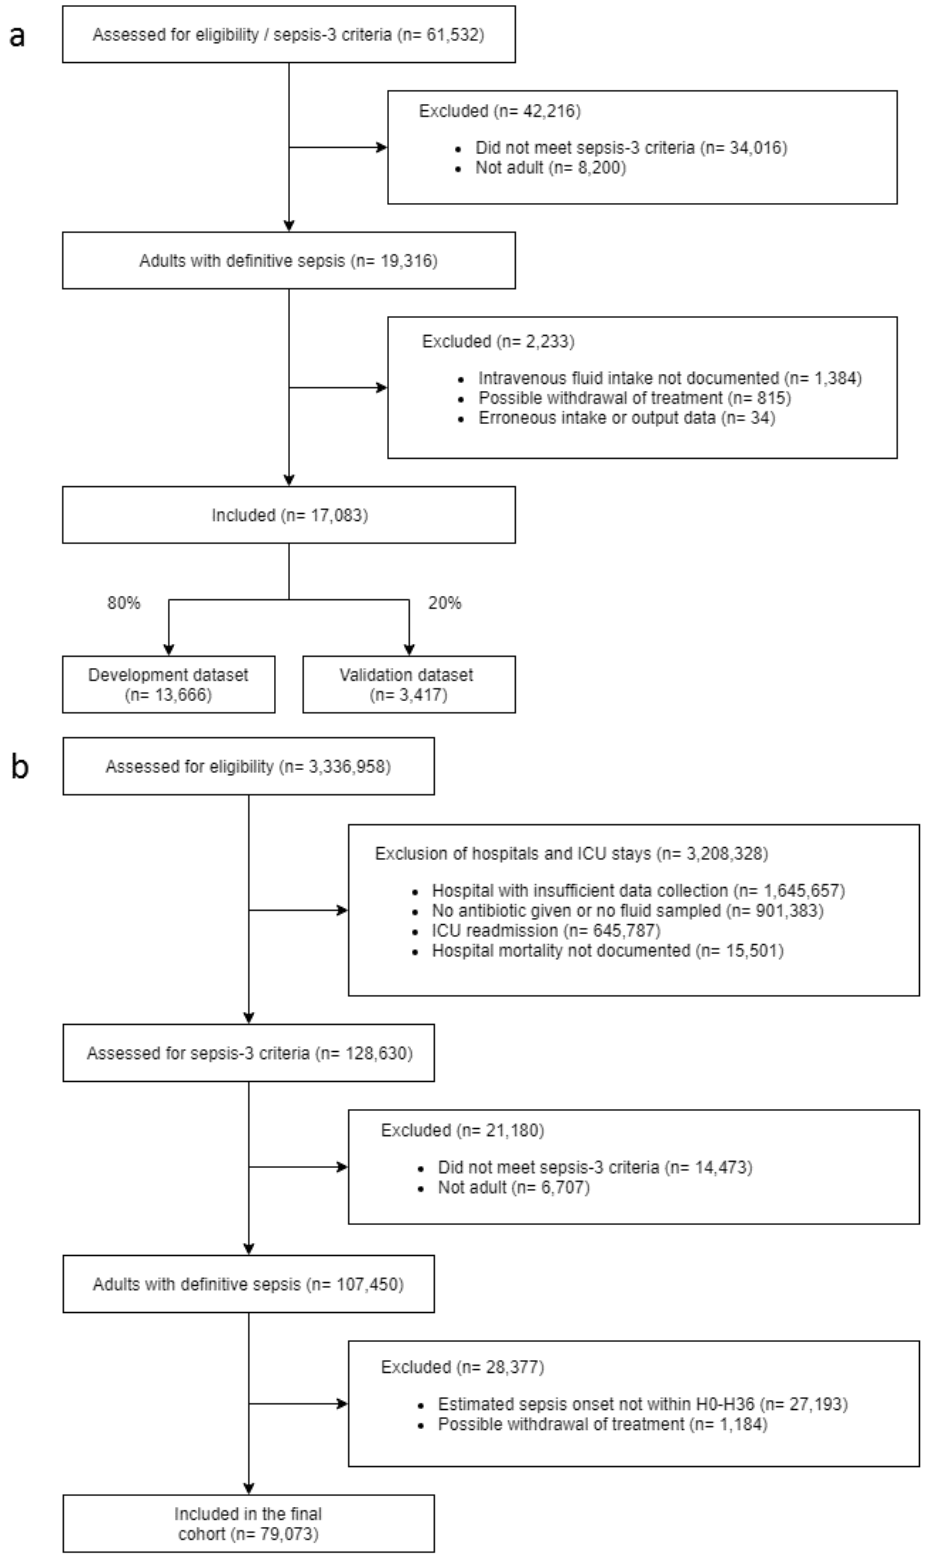

Supplementary Figure 1. Patient inclusion diagrams in MIMIC-III (a) and eRI (b).

Figure 2: Feature importance in the clinicians' policy and the AI policy

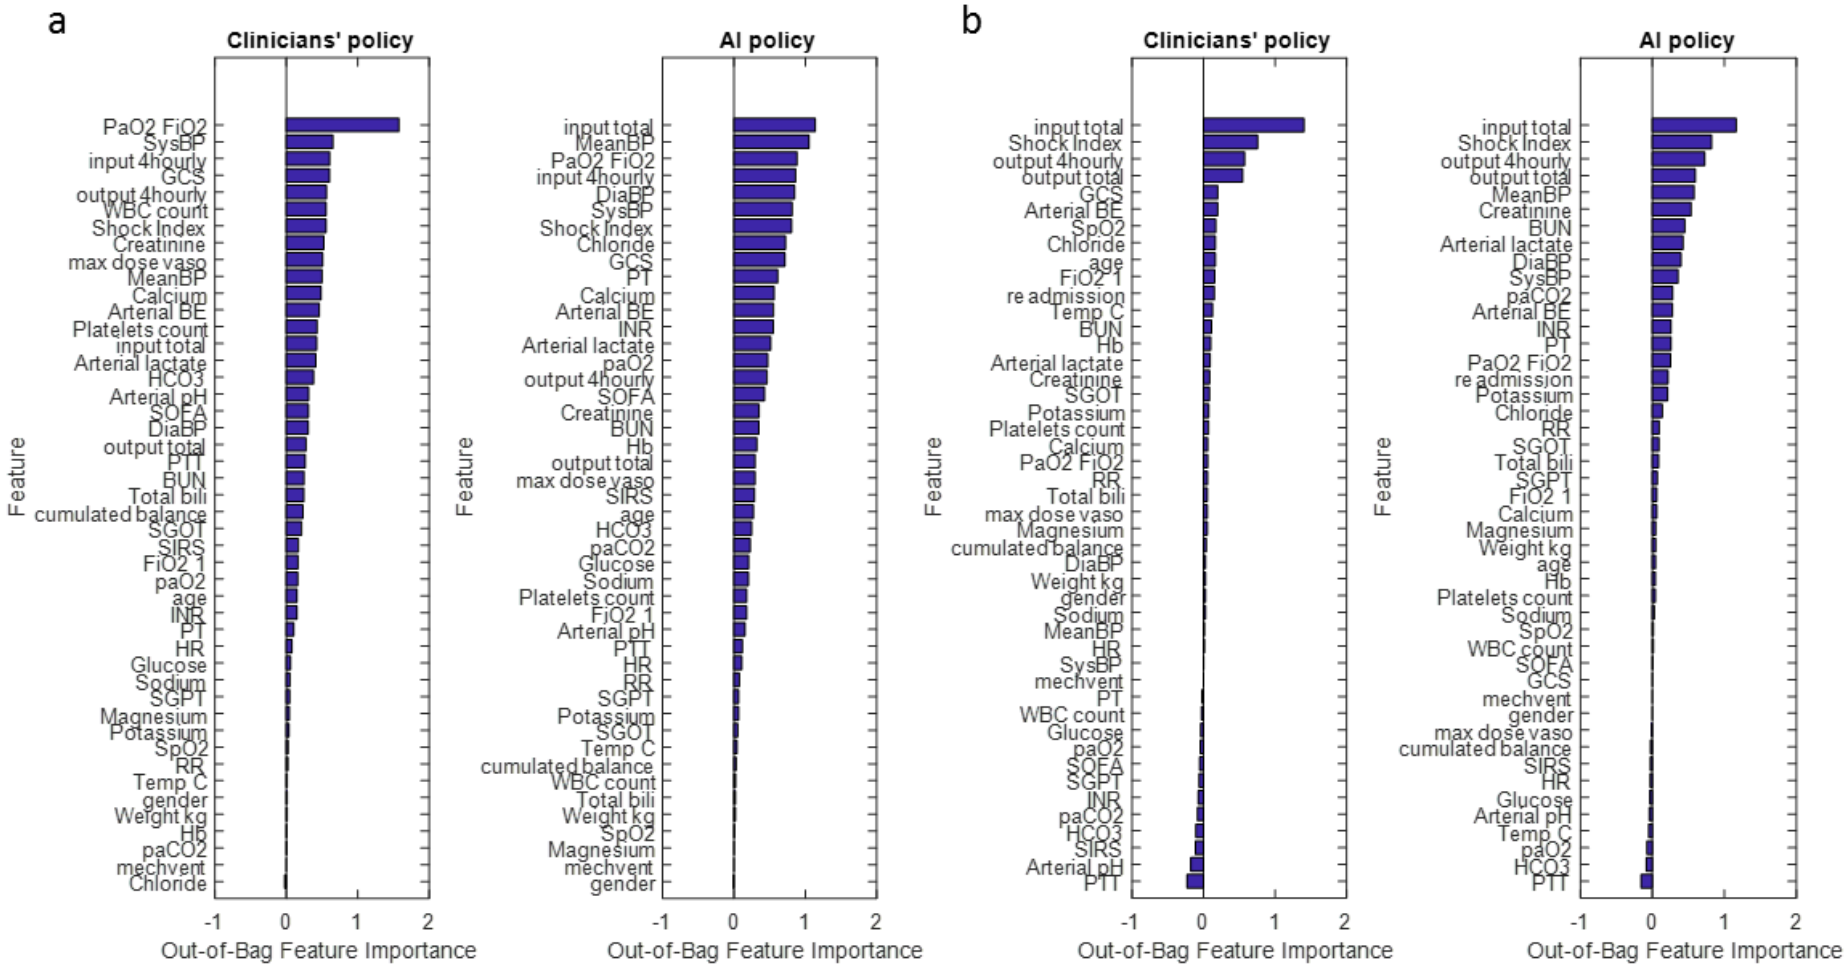

Supplementary Figure 2. Model interpretability: feature importance underlying the treatment strategies of clinicians and our AI algorithm.

We built classification random forest models to predict whether the medications were (clinicians' policy) or should have been (AI policy) administered (regardless of the dose), using patient variables as input data. The current dose of vasopressor or intravenous fluid was discarded from the input data in the respective models. Then, the relative importance of each variable was estimated using an out-of-bag technique, where we measured the loss of prediction ability (an increase in the mean squared error on prediction) while we permuted the values of each variable across every observation in the dataset. When permuted, important variables led to large increases in the error on prediction. Then, we plotted the estimated variable importance averaged over all trees in the random forest ensemble, for intravenous fluids (a) and vasopressors (b), for both the clinicians' and the AI policy. This confirmed that the decisions suggested by the reinforcement learning algorithm were clinically interpretable and relied primarily on sensible clinical and biological parameters, such as arterial lactate, mean blood pressure, or urine output.

Figure 3: State representation and Markov property

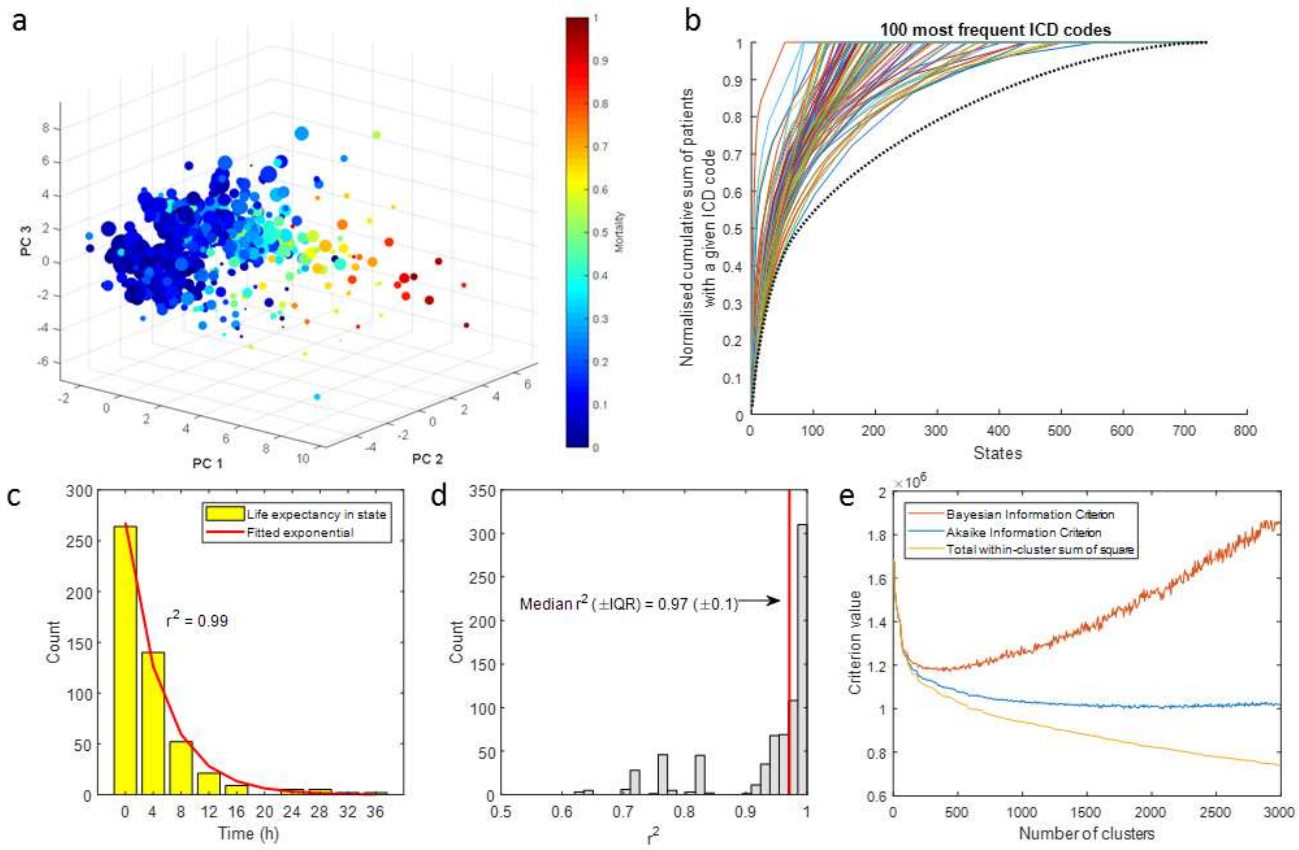

Supplementary Figure 3. State representation and Markov property.

a, Mapping of the 750 health states to low-dimensional projection. Each spot represents a cluster centroid, whose size and colour correspond to, respectively, the number of records and the average mortality in the state. PC: principal component. A spontaneous gradient can be seen in the state mortality. This demonstrates that state membership was indeed associated with outcome (mortality), which supports the validity of the state representation.

b. Capture of clinical concepts and past medical history in the model states, using the 100 most frequently used International Classification of Diseases (ICD) codes as a surrogate. We measured how many patients have a given ICD code in all the states, and show the cumulative sum across states of patients, starting with the states with the highest number of patients with a given code. The black dotted line shows the cumulative sum of patients in the states, with states ordered by descending size. This is the theoretical distribution we should obtain if the ICD codes were randomly assigned to the clusters (the proportion of a given ICD code should be globally equal in all the states). The majority of patients are found in less states than if the codes were randomly distributed. For example, 50% of the patients with the ICD code "Coronary atherosclerosis of native coronary artery" code are found in only 23 states. This number would be 79 if the code was randomly distributed in the states.

c-d, Verification of the Markov property. We measured the life expectancy in each state, by observing how long an agent would remain in a given state when following the transition matrix, in 500 trials. c, Example of the life expectancy in one state, with fitted exponential decay function. The correlation coefficient r^2 between the data and the fitted function is 0.99 in this example. d, Distribution of the correlation coefficients between life expectancy and exponential decay functions in the 750 states of the model. The high median correlation coefficient of 0.97 (interquartile range 0.1) confirmed that the life expectancy in most states was indeed memoryless.

e. Selection of the number of clusters in the model. Bayesian information criterion (BIC), Akaike information criterion (AIC), and total within-cluster sum of square was computed as a function of the number of clusters. The minimum is found for k around 2,000 for AIC and k=400 for BIC.

Figure 4: Internal validation in the MIMIC-III test set: model optimizing hospital mortality

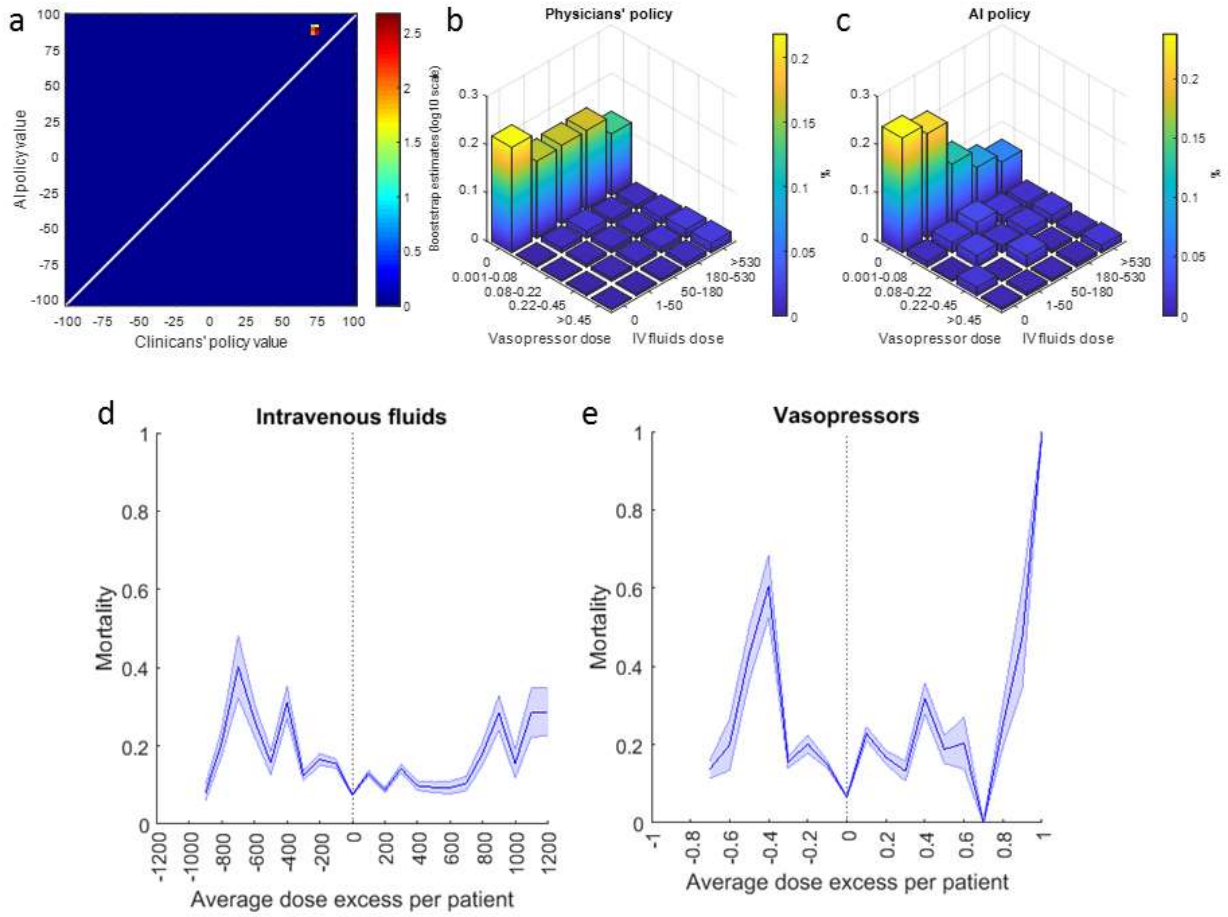

Supplementary Figure 4. Internal validation in the MIMIC-III test set: model optimizing hospital mortality.

a: Distribution of the estimated value of the clinicians' policy and the AI policy in the selected model, built by bootstrapping with 2,000 resamplings.

b-c: Visualization of the clinicians' and AI policies. All actions were aggregated over all time steps for the 5 dose bins of both medications. On average, patients were administered more intravenous fluid and less vasopressor medications than recommended by the AI policy. Vasopressor dose is in mcg/kg/min of norepinephrine equivalent and intravenous fluids dose is in mL/4 hours.

d-e, Average dose excess received per patient of intravenous fluids (d) and vasopressors (e), and corresponding mortality. The figure is built by bootstrapping with 2,000 resamplings. The lowest expected mortality was found when the dose actually administered to the patients matched the dose suggested by the AI policy