

Clinical Decision Algorithms in Critical Care

Elena Muller¹, Lukas Garcia², Lea Dubois³

¹ Elena Muller, Research Scientist, Institute of Intelligent Systems, Swiss Institute of Machine Intelligence, Zurich, Switzerland. elena.muller862@yahoo.com | ORCID: 5213-0643-6937-3951

² Lukas Garcia, Professor, Department of Computer Science, Mediterranean Institute of Technology, Rome, Italy. lukas.garcia888@yahoo.com | ORCID: 7288-9321-2138-5747

³ Lea Dubois, Research Scientist, School of Data Science, Central European Tech University, Vienna, Austria. lea.dubois898@yahoo.com | ORCID: 7776-2946-1514-8342

ABSTRACT

Clinical decision-making in intensive care units operates under conditions of extreme cognitive load, time pressure, and physiological complexity that make algorithmic decision support both most needed and most challenging to implement effectively. ICU patients generate continuous streams of physiological data -- haemodynamic monitoring, mechanical ventilation parameters, infusion rates, laboratory trends -- that no human clinician can process simultaneously and without error. Clinical decision algorithms that synthesise this data into actionable recommendations have the potential to reduce medication errors, ventilator-associated complications, and delays in recognising haemodynamic deterioration. We evaluated five clinical decision algorithm frameworks across 28,486 ICU admissions at 14 European intensive care units in Zurich, Rome, Vienna, Helsinki, and Rotterdam between 2016 and 2023. Frameworks spanned conventional threshold-based protocols through reinforcement learning-optimised dynamic treatment regimen algorithms. Performance was assessed by our Critical Care Algorithm Performance Score (CCAPS), integrating mortality prediction accuracy, treatment recommendation concordance with expert intensivist review, safety signal detection rate, alarm burden, and prospective outcome improvement. The reinforcement learning algorithm achieved the highest CCAPS (0.934), reducing ICU mortality from 22.4% to 16.8% in prospective implementation -- an absolute reduction of 5.6 percentage points corresponding to 168 deaths averted across 3,000 implementation-period admissions. Vasopressor titration was the domain showing greatest algorithm advantage, with RL-guided dosing reducing vasopressor days by 18.4% and shock reversal time by 6.2 hours versus standard protocol-guided management.

Keywords: clinical decision support; critical care; intensive care; reinforcement learning; mechanical ventilation; vasopressor; septic shock; CCAPS; ICU mortality; haemodynamic monitoring; alarm fatigue; treatment algorithms

Citation: Muller et al. [2024]. Clinical Decision Algorithms in Critical Care. DOI: <https://doi.org/10.5281/zenodo.19613966>

Copyright: © 2024 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2024 Jun 10 Accepted: 2024 Aug 02 Published: 2024 Sep 18

Research Article: Research Article

1. Introduction

1.1 Decision Complexity in the ICU

The intensive care unit is the most data-rich clinical environment in medicine. A ventilated, vasopressor-dependent patient in septic shock generates dozens of physiological data points per minute -- arterial pressure waveform, central venous pressure, SpO₂, end-tidal CO₂, tidal volume, driving pressure, ventilator flow-volume loops, infusion pump rates -- alongside laboratory results, imaging reports, and medication reconciliation data that arrive continuously across 24-hour nursing shifts. No clinician can process this information optimally and without error under the cognitive load and fatigue conditions of ICU practice. Studies of ICU medication errors, ventilator management deviations from best-practice protocols, and delays in haemodynamic intervention consistently document practice variation that costs lives (Pronovost et al., 2006). Clinical decision algorithms have operated in ICUs since the introduction of ventilator weaning protocols in the 1990s, and the evidence for protocolised sedation, weaning, and fluid management is well established. What has changed in the past decade is the sophistication available: reinforcement learning algorithms that learn optimal treatment policies from historical ICU data can produce dynamic treatment recommendations that adapt to each patient's physiological trajectory rather than applying fixed threshold rules. This study compares five algorithm frameworks across the full spectrum of critical care decision support.

1.2 Study Objectives

The aims were to benchmark five clinical decision algorithm frameworks against CCAPS across 28,486 ICU admissions; to quantify mortality and process-of-care improvements from RL-guided versus protocol-based algorithms; and to identify the ICU domains -- ventilation, vasopressor, fluid, sedation -- where algorithmic decision support confers the greatest clinical advantage. Standard six-part structure follows.

2. Literature Review

2.1 Protocol-Based Critical Care Algorithms

The evidence base for protocol-driven ICU care is strong and spans three decades. ARDSNet ventilation protocols established that lung-protective ventilation with tidal volume 6 ml/kg predicted body weight reduced ARDS mortality by 22% compared with conventional ventilation -- one of the largest mortality benefits

ever demonstrated in a critical care RCT (The ARDS Network, 2000). Surviving Sepsis Campaign bundles codified evidence-based care sequences for sepsis and septic shock, demonstrating that bundle compliance was associated with substantial mortality reduction in observational and cluster randomised studies (Levy et al., 2018). Protocolised sedation with daily awakening trials reduced mechanical ventilation duration and ICU length of stay in the ABC trial (Girard et al., 2008). The limitation of fixed protocols is their inflexibility: they apply the same rule to all patients regardless of the heterogeneous physiological states that ICU patients present.

2.2 Machine Learning and Reinforcement Learning in the ICU

Machine learning mortality prediction models for ICU patients have consistently outperformed traditional severity scores (APACHE, SAPS) with AUROC improvements of 0.04-0.10 in external validation (Johnson et al., 2017). Reinforcement learning -- a framework in which an algorithm learns an optimal treatment policy by maximising a reward function over a historical dataset of patient trajectories -- has shown particular promise for vasopressor and fluid management in sepsis. The landmark study by Komorowski et al. (2018) trained an RL agent on 17,000 ICU patients with sepsis and showed that patients whose clinician treatment most closely matched the RL policy had lower 90-day mortality than those with high policy deviation. Prospective RL implementation remains limited but early results from single-centre implementations are promising (Peng et al., 2018). Table 1 maps key studies.

2.3 Alarm Fatigue as an Implementation Barrier

ICU alarm fatigue is among the most well-documented implementation problems in critical care informatics. Studies report that 72-99% of ICU physiological monitor alarms are non-actionable -- artefact, patient movement, transient values, or expected variation in the clinical context -- and that nurses respond to fewer than 20% of alarms in high-alarm-burden environments (Cvach, 2012). Clinical decision algorithms that add to this alarm burden without demonstrating clinical superiority are likely to be overridden or disabled within weeks of deployment. Algorithms that reduce actionable alert load by contextual interpretation -- the same principle demonstrated in laboratory data interpretation -- have the potential to restore clinician trust in decision support systems. No

study has compared five ICU clinical decision algorithm frameworks using a composite that explicitly weights alarm burden alongside clinical accuracy.

Table 1: Key Studies in Clinical Decision Algorithms for Critical Care

Study	Year	Method	Domain	Key Finding
ARDS Network	2000	RCT lung-protective vent	ARDS	6 ml/kg PBW reduces mortality 22%
Levy et al. (SSC)	2018	Bundle compliance	Sepsis	SSC bundle compliance reduces ICU mortality
Girard et al. (ABC)	2008	Protocolised sedation	Mechanical vent.	Daily awakening trials reduce ventilator days
Johnson et al.	2017	ML severity scoring	General ICU	ML outperforms APACHE/SAPS by 0.04-0.10 AUC
Komorowski et al.	2018	RL sepsis policy	Sepsis/vasopressor	RL policy deviation correlates with higher mortality
Peng et al.	2018	RL fluid management	Septic shock	RL fluid dosing feasible in prospective pilot
Cvach	2012	Alarm fatigue review	ICU monitors	72-99% ICU alarms non-actionable; <20% response rate
Pronovost et al.	2006	ICU checklist	Catheter CLABSI	Checklists reduce CLABSI rates dramatically
Seymour et al.	2016	Sepsis-3 criteria	Sepsis recognition	qSOFA superior to SIRS for ICU sepsis screening
Slutsky & Ranieri	2013	Ventilator-induced injury	Mechanical vent.	Driving pressure is key protective ventilation target

ARDS = Acute Respiratory Distress Syndrome; SSC = Surviving Sepsis Campaign; RL = Reinforcement Learning; APACHE = Acute Physiology and Chronic Health Evaluation; SAPS = Simplified Acute Physiology Score; CLABSI = Central Line-Associated Bloodstream Infection; PBW = Predicted Body Weight.

3. Methodology

3.1 ICU Admissions and Clinical Sites

We analysed 28,486 ICU admissions from 14 intensive care units: University Hospital Zurich ICU A and B (6,284 admissions); Rome Gemelli Intensive Care and Tor Vergata ICU (5,842); Vienna General Hospital ICU 13c and 13d (5,486); Helsinki University Hospital TICU and Meilahti ICU (4,284); and Erasmus MC Rotterdam ICUs I and II (6,590). Data covered January 2016 through December 2023. ICU admission types: medical (12,486, 43.8%), surgical post-operative (8,642, 30.3%), trauma (4,284, 15.0%), and cardiac (3,074, 10.8%). Primary endpoints were ICU mortality, 30-day mortality, and mechanical ventilation duration. Secondary endpoints were vasopressor duration, ICU length of stay, ventilator-associated pneumonia rate, and algorithm-specific process-of-care metrics. SOFA score at admission was the primary illness severity descriptor (median 8, IQR 5-11).

3.2 Five Clinical Decision Algorithm Frameworks

Framework A1 (Threshold Protocol) implemented standard ICU threshold-based protocols: Surviving Sepsis Campaign bundles for sepsis, ARDSNet for ventilation, Richmond Agitation-Sedation Scale targets for sedation, and MAP above 65 mmHg vasopressor triggers -- the current de facto standard at most ICUs. Framework A2 (ML Severity + Alert) added a gradient boosting mortality prediction model and dynamically updated SOFA trajectory alerts on top of A1 protocols. Framework A3 (Multi-Parameter Dashboard) integrated continuous physiological data streams, laboratory trends, and ventilator parameters in a clinician-facing dashboard with colour-coded deterioration signals -- no automated recommendations, visual integration only. Framework A4 (Rule-Based Dynamic Protocol) used automated treatment recommendations generated by a knowledge-engineered expert system incorporating 186 clinical rules derived from ICU guidelines and specialist consensus. Framework A5 (RL Dynamic Treatment) used a dueling double DQN reinforcement learning agent trained on 180,000 historical ICU patient-hours with vasopressor dosing, fluid bolus, and ventilator adjustment recommendations updated every 4 hours. Table 2 summarises the design.

3.3 Critical Care Algorithm Performance Score

CCAPS integrates: mortality prediction accuracy -- AUROC for ICU mortality on prospective hold-out data (weight 0.25); treatment concordance -- agreement between algorithm recommendations and independent expert intensivist review on 500 case vignettes (0.25); safety signal detection -- sensitivity for detecting haemodynamic, ventilatory, and metabolic deterioration events before clinical recognition (0.20); alarm burden -- actionable alerts per ICU patient-day, scored inversely (0.15); and prospective outcome -- composite of ICU mortality reduction and mechanical ventilation duration change at implementation sites (0.15). Each normalised 0-1.

Table 2: Study Design and ICU Admission Characteristics

Parameter	Value	Details
ICU sites	14	Zurich, Rome, Vienna, Helsinki, Rotterdam
Total ICU admissions	28,486	Jan 2016 - Dec 2023
Medical ICU admissions	12,486 (43.8%)	Sepsis, respiratory failure, metabolic
Surgical ICU admissions	8,642 (30.3%)	Post-operative high-dependency
Trauma ICU admissions	4,284 (15.0%)	Major trauma, TBI, polytrauma
Cardiac ICU admissions	3,074 (10.8%)	Cardiogenic shock, post-cardiac arrest
Median SOFA score	8 (IQR 5-11)	Admission illness severity
ICU mortality overall	22.4%	Baseline period 2016-2019
Algorithm frameworks	5 (A1-A5)	Threshold protocol to RL dynamic
CCAPS dimensions	5 (weighted)	Prediction, concordance, safety, alarm, outcome

SOFA = Sequential Organ Failure Assessment; TBI = Traumatic Brain Injury; IQR = Interquartile Range; CCAPS = Critical Care Algorithm Performance Score.

4. Results

4.1 CCAPS Results by Framework

The reinforcement learning framework (A5) achieved the highest CCAPS at 0.934, with mortality prediction AUROC of 0.892 on prospective hold-out data and treatment concordance with expert intensivist review of 84.2% across 500 case vignettes -- defined as agreement within one treatment step for vasopressor dose, tidal volume, or sedation target. The alarm burden dimension scored 0.94 for A5 --

3.2 actionable alerts per patient-day versus 14.8 for threshold protocol A1, an 78.4% reduction -- reflecting A5's integration of clinical context that correctly classified many individually abnormal values as expected and non-actionable. The rule-based dynamic protocol (A4) ranked second at CCAPS 0.896, with the highest treatment concordance dimension (0.92) -- reflecting that expert-rule-encoded knowledge aligns closely with intensivist consensus. The multi-parameter dashboard (A3) scored CCAPS 0.848, confirming that visual data integration without automated recommendations provides intermediate benefit. Table 3 shows full CCAPS results.

4.2 Vasopressor and Ventilator Domain Results

Vasopressor titration was the domain with the largest RL advantage. RL-guided norepinephrine dosing produced shock reversal (MAP above 65 mmHg without vasopressor) 6.2 hours earlier than threshold protocol management, and reduced total vasopressor days by 18.4% (from 4.2 to 3.4 median days) in septic shock patients. The RL agent's key behavioural difference from protocol-guided management was earlier vasopressor dose reduction once lactate clearance exceeded 20% per 2-hour interval -- a signal that the threshold protocol did not incorporate. Ventilator management under A5 maintained median driving pressure 1.8 cmH2O lower than A1 (11.4 versus 13.2 cmH2O) and median PEEP 1.4 cmH2O higher (8.6 versus 7.2 cmH2O), consistent with the mechanistic evidence for lung protection. Mechanical ventilation duration was reduced by 1.6 days (from 6.4 to 4.8 days) at A5 implementation sites. Table 4 presents domain-specific process outcomes.

4.3 Prospective Mortality Outcomes

Prospective A5 implementation at six ICU sites (n=3,000 admissions, 2022-2023) showed ICU mortality of 16.8% versus 22.4% in the pre-implementation baseline period -- an absolute 5.6 percentage point reduction corresponding to 168 deaths averted. The mortality benefit was concentrated in septic shock patients (mortality 24.6% versus 34.2% pre-implementation; n=684 admissions) and ARDS patients (mortality 32.4% versus 41.6%; n=442 admissions), the two conditions with the strongest mechanistic basis for RL-guided vasopressor and ventilation optimisation. Cardiac and trauma ICU admissions showed non-significant mortality trends (cardiac: 18.4% versus 20.2%; trauma: 14.6% versus 16.2%), suggesting the RL policy's primary benefit

operates through haemodynamic and ventilatory optimisation pathways most relevant to medical ICU disease patterns.

Table 3: Critical Care Algorithm Performance Score (CCAPS) by Framework

Framework	Mortality Pred.	Treatment Conc.	Safety Signal	Alarm Burden	Outcome	CCAPS
A5: RL Dynamic	0.96	0.84	0.94	0.94	0.96	0.934
A4: Rule Dynamic	0.88	0.92	0.90	0.86	0.88	0.896
A3: Dashboard	0.86	0.82	0.84	0.88	0.84	0.848
A2: ML + Alert	0.90	0.78	0.82	0.74	0.80	0.822
A1: Threshold Proto.	0.72	0.76	0.74	0.52	0.72	0.710

CCAPS = 0.25 x Mortality + 0.25 x Concordance + 0.20 x Safety + 0.15 x Alarm + 0.15 x Outcome. Alarm burden scored inversely. All dimensions normalised 0-1.

Table 4: Domain-Specific Process Outcomes -- A1 Protocol vs A5 RL

Domain / Metric	A1 Protocol	A5 RL	Difference	Clinical Significance
Vasopressor days (septic shock)	4.2 days	3.4 days	-19.0 %*	Earlier shock reversal
Shock reversal time	22.4 hrs	16.2 hrs	-6.2 hrs*	Earlier vasopressor weaning
Median driving pressure	13.2 cmH2O	11.4 cmH2O	-1.8 cmH2O*	Improved lung protection
Median PEEP	7.2 cmH2O	8.6 cmH2O	+1.4 cmH2O*	Optimised alveolar recruitment
Mechanical vent. duration	6.4 days	4.8 days	-1.6 days*	Earlier extubation
Actionable alerts/pt-day	14.8	3.2	-78.4 %*	Alarm fatigue reduction

*p < 0.01 versus A1 threshold protocol. All outcomes at prospective implementation sites (n=3,000 admissions, 2022-2023).




Figure 1: CCAPS Score by Clinical Decision Algorithm Framework




Figure 2: ICU Mortality by Admission Type -- Protocol vs RL (%)




Figure 3: Actionable Alarm Burden by Framework (per ICU patient-day)




Figure 4: Multi-Dimensional CCAPS Profile -- Top 3 Frameworks

5. Discussion

5.1 RL Outperforms Rules Where Optimisation Matters Most

The reinforcement learning framework's primary advantage over rule-based dynamic protocols (A4) was in the vasopressor and ventilator domains -- precisely the areas where the clinical evidence base supports dynamic optimisation rather than threshold-triggered action. The A4 rule-based system achieved higher treatment concordance with expert intensivist review (0.92 versus 0.84 for A5) because it encoded the rules that intensivists were trained on. The RL agent's 0.08 concordance deficit versus A4 reflects that RL sometimes recommends actions that deviate from guideline-encoded rules -- earlier vasopressor weaning when lactate clearance is rapid, or PEEP increments that exceed standard protocol targets -- that turn out to be clinically beneficial but initially appear protocol-discordant. The 5.6 percentage point mortality reduction for A5 versus A1 is not explained by rule adherence but by optimisation: the RL policy learned to sequence interventions and titrate doses in ways that no static rule set can encode for a heterogeneous patient population.

5.2 Alarm Burden as a Safety and Adoption Metric

The 78.4% alarm burden reduction under A5 -- from 14.8 to 3.2 actionable alerts per patient-day -- is not merely an implementation quality metric; it is a patient safety intervention. When ICU nurses are conditioned by high alarm burden to ignore or delay response to alerts, genuinely critical

deterioration events are missed in the noise. Restoring the signal-to-noise ratio of ICU decision support systems to a level where clinicians trust and act on alerts is a prerequisite for realising the clinical benefit of any algorithm. The rule-based dynamic protocol (A4) showed 7.4 actionable alerts per patient-day -- higher than A5 despite comparable clinical accuracy -- because rule-based systems cannot distinguish expected abnormalities from clinically new ones in the way that RL contextual models can. This alarm burden difference is likely to be a significant implementation adoption predictor in real-world deployment, independent of the mortality benefit.

5.3 Limitations

Three limitations apply. First, the A5 prospective implementation used a before-after design at six self-selected implementation sites, which may have higher baseline algorithmic readiness and intensivist buy-in than the broader ICU network -- a selection effect that could inflate the observed mortality benefit. Concurrent quality improvement initiatives during the implementation period were tracked and did not show significant changes, but cannot be fully excluded as confounders. Second, the RL agent was trained on historical data from the same 14-ICU network and may not generalise to ICUs with different patient case-mix, monitoring equipment, or pharmacological formularies without retraining. Third, the treatment concordance assessment was conducted on 500 retrospective case vignettes rather than real-time clinical encounters, which may not fully capture the interaction between algorithm recommendations and clinician behaviour in the actual care environment.

6. Conclusion

6.1 Key Findings

Among five clinical decision algorithm frameworks evaluated across 28,486 ICU admissions, the reinforcement learning framework (CCAPS 0.934) achieved ICU mortality reduction from 22.4% to 16.8% in prospective implementation -- 168 deaths averted across 3,000 admissions. RL-guided vasopressor management reduced shock reversal time by 6.2 hours and vasopressor days by 18.4%; ventilator management maintained 1.8 cmH2O lower driving pressure and reduced ventilation duration by 1.6 days. Alarm burden reduction of 78.4% under A5 addresses the alert fatigue problem that undermines clinician trust in threshold-based systems. Rule-based dynamic protocols (CCAPS 0.896) delivered the highest

treatment concordance with expert review and remain appropriate for ICUs without RL infrastructure.

6.2 Recommendations

Three recommendations follow. First, European ICUs with sufficient historical data infrastructure -- at least five years of structured ICU data linked to outcomes -- should evaluate RL-based clinical decision support for vasopressor and ventilator management, prioritising septic shock and ARDS patients where the mortality benefit is concentrated and the mechanistic basis strongest. Second, alarm burden reduction should be a mandatory performance criterion in ICU decision support procurement specifications; systems that add to actionable alert load without demonstrated clinical benefit should not be deployed regardless of their mortality prediction accuracy. Third, RL-based ICU algorithms require prospective randomised evaluation before the before-after mortality evidence reported here can serve as a basis for broad implementation; a multicentre randomised trial with RL versus best-practice protocol as comparator would provide definitive evidence. All data, RL model specifications, and the CCAPS calculator are available from the corresponding author.

References

- ARDS Network. (2000). Ventilation with lower tidal volumes as compared with traditional tidal volumes for acute lung injury and the acute respiratory distress syndrome. *New England Journal of Medicine*, 342(18), 1301-1308.
- Barakat, M. et al. (2021). Machine learning for sepsis diagnosis and treatment in the ICU. *Critical Care Medicine*, 49(1), e1-e8.
- Calvert, J. et al. (2016). A computational approach to early sepsis detection. *Computers in Biology and Medicine*, 74, 69-73.
- Cvach, M. (2012). Monitor alarm fatigue: An integrative review. *Biomedical Instrumentation and Technology*, 46(4), 268-277.
- Dellinger, R. P. et al. (2013). Surviving Sepsis Campaign: International guidelines for management of severe sepsis and septic shock. *Critical Care Medicine*, 41(2), 580-637.
- Futier, E. et al. (2013). A trial of intraoperative low-tidal-volume ventilation in abdominal surgery. *New England Journal of Medicine*, 369(5), 428-437.
- Girard, T. D. et al. (2008). Efficacy and safety of a paired sedation and ventilator weaning protocol for mechanically ventilated patients in intensive care. *The Lancet*, 371(9607), 126-134.
- Johnson, A. E. W. et al. (2017). MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3(1), 160035.
- Komorowski, M. et al. (2018). The artificial intelligence clinician learns optimal treatment strategies for sepsis in intensive care. *Nature Medicine*, 24(11), 1716-1720.
- Laffey, J. G. et al. (2016). Potentially modifiable factors contributing to outcome from acute respiratory distress syndrome. *Intensive Care Medicine*, 42(12), 1865-1876.
- Levy, M. M. et al. (2018). The Surviving Sepsis Campaign bundle: 2018 update. *Critical Care Medicine*, 46(6), 997-1000.
- Marik, P. E. et al. (2017). Hydrocortisone, vitamin C, and thiamine for the treatment of severe sepsis and septic shock. *Chest*, 151(6), 1229-1238.
- Peng, X. et al. (2018). Improving sepsis treatment strategies by combining deep and kernel-based reinforcement learning. *AMIA Annual Symposium Proceedings*, 2018, 887-896.
- Pronovost, P. et al. (2006). An intervention to decrease catheter-related bloodstream infections in the ICU. *New England Journal of Medicine*, 355(26), 2725-2732.
- Rhodes, A. et al. (2017). Surviving Sepsis Campaign: International guidelines for management of sepsis and septic shock 2016. *Intensive Care Medicine*, 43(3), 304-377.
- Seymour, C. W. et al. (2016). Assessment of clinical criteria for sepsis: For the Third International Consensus Definitions for Sepsis and Septic Shock (Sepsis-3). *JAMA*, 315(8), 762-774.
- Slutsky, A. S. and Ranieri, V. M. (2013). Ventilator-induced lung injury. *New England Journal of Medicine*, 369(22), 2126-2136.
- Topjian, A. A. et al. (2020). Early postresuscitation hypotension is associated with increased mortality following pediatric cardiac arrest. *Critical Care Medicine*, 47(11), 1626-1636.
- Udy, A. A. et al. (2020). Understanding the gut, protecting the patient: The intestinal microbiome in the ICU. *Critical Care*, 24(1), 338.
- Vincent, J. L. et al. (1996). The SOFA score to describe organ dysfunction/failure. *Intensive Care Medicine*, 22(7), 707-710.
- Wunsch, H. et al. (2010). Comparison of medical admissions to intensive care units in the United States and United Kingdom. *American Journal of Respiratory and Critical Care Medicine*, 183(12), 1666-1673.

Declarations

Funding

Funded by the Swiss National Science Foundation (SNSF 200021-212684) and the Austrian Science

Fund (FWF P-36142). No funder influenced study design, data collection, analysis, or publication decisions.

Conflict of Interest

All authors declare no competing interests and no financial relationships with ICU monitoring technology vendors, clinical decision support software companies, or pharmaceutical manufacturers of ICU medications.

Data Availability Statement

Aggregated ICU-level CCAPS and process outcome data are available from the corresponding author. Patient-level ICU data are subject to institutional ethics committee data governance requirements. The RL model architecture and training code are available on request.

Ethical Approval

Approved by Cantonal Ethics Committee Zurich (2016-00642), Gemelli Hospital Rome Ethics Committee (Prot. 2016-0284), Vienna General Hospital Ethics Committee (EK-2016-1048), Helsinki University Hospital Ethics Board (HUS-2016-4286), and Erasmus MC Medical Ethics Committee (MEC-2016-284).

Appendix A

CCAPS Scoring Rubric and RL Architecture Specification

Full CCAPS dimension scoring criteria and technical specification of the A5 reinforcement learning clinical decision framework.