Human-in-the-Loop Multi-Agent Ventilator Decision Support with Contextual Bandit Preference Learning

Sijia Li[1,*], Xiaoyu Tan[2,*], Qixing Wang[3], Weiyi Zhao[1], Chen Zhan[1], Teqi Hao[1], Xuemin Wang[4], Lei Gu[5], Roland Eils[6,7], and Xihe Qiu[1,✉]

[1] Shanghai University of Engineering Science, Shanghai, China
qiuxihe1993@gmail.com
[2] Tencent Youtu Lab, Tencent, China
[3] Department of Critical Care Medicine, Shanghai Tenth People's Hospital, Tongji University School of Medicine, Shanghai, China
[4] Department of Emergency and Critical Disease, Songjiang Hospital Affiliated to Shanghai Jiao Tong University School of Medicine, Shanghai, China
[5] Max Planck Institute for Heart and Lung Research, Bad Nauheim, Germany
[6] Fudan University, Shanghai, China
[7] BIH at Charité – Universitätsmedizin Berlin, Berlin, Germany

Abstract. Ventilator decision support requires sequential decisions that track evolving physiology and disease trajectories while respecting safety boundaries and clinician specific tuning styles. Rule based approaches rarely generalize personalization, and end to end reinforcement learning or single large language model systems remain difficult to control and audit. We propose the **V**entilator **D**ecision **S**upport **S**ystem (**VDSS**), a human in the loop multi agent framework that coordinates modular decision components through contract driven structured interfaces and produces traceable evidence for review. VDSS performs online preference adaptation with a contextual bandit, updating clinician specific preferences from the final accepted decision at each adjustment cycle and using them to guide subsequent recommendations. Structured rejection feedback triggers targeted replanning to reduce unproductive iterations and improve interaction stability. Retrospective ICU trajectory replay with expert review indicates higher recommendation acceptability and fewer interaction rounds to reach an acceptable plan, supporting clinically deployable human AI collaboration.

Keywords: Ventilator decision support · Multi agent system · Contextual bandit

1 Introduction

Mechanical ventilation is a central life sustaining therapy in the intensive care unit. Its clinical value extends beyond maintaining oxygenation and ventilation,

*Equal contribution. ✉ Corresponding author.

as effective care requires continuous and fine grained adjustment of settings to match a patient's evolving physiology [24, 4]. Ventilator management is inherently dynamic. Bedside teams must repeatedly revise decisions as lung mechanics, gas exchange, hemodynamics, and sedation levels change, while balancing multiple clinical objectives in real time [16]. At the same time, high quality titration depends on frequent assessment and rapid response by respiratory therapists and experienced clinicians [10]. This specialized workforce is persistently scarce in many health systems, and sustained workload together with shift handoffs makes it difficult to deliver consistently individualized management even within the same institution [6]. As patient volume rises and workflow pressure increases, a practical motivation for ventilator decision support is to embed expert practice in a scalable and reusable form that reduces cognitive burden and improves consistency [15].

Current practice is guided by protocols and accumulated experience, yet real ICU care presents persistent challenges for decision support [5]. Patient heterogeneity and nonstationary trajectories cause strategy benefits to vary across individuals and over time, limiting fixed rule reliability [9, 19]. Bedside titration often requires rapid setting changes based on multiple imperfect signals, and limited time to check and document the rationale can make adjustments inconsistent across cycles [21, 7]. Clinical teams further differ in risk tolerance, tuning style, and local conventions, and without explicit modeling and adaptation, decision support is unlikely to achieve sustained trust or uptake [8, 26]. Beyond a recommended setting, clinicians need a transparent link between the recommendation, the patient state, prevailing practice, and team preferences to support efficient repeated use [14]. These factors motivate ventilator decision support that is workflow aligned, coherent, and interpretable, while able to adapt as clinical preferences evolve [20, 1].

Motivated by these needs, we argue that in high risk clinical settings, pointwise model performance is necessary but not sufficient [22, 13]. Successful deployment instead hinges on decision support whose behavior can be inspected at the bedside and refined over time without breaking clinical workflow. We therefore propose the Ventilator Decision Support System, a human in the loop multi agent framework that decomposes ventilator decision making into modular agents and regulates inter agent communication through contract driven structured interfaces, enabling coherent multi round collaboration with a clear, auditable reasoning trace [27, 3]. The system uses layered short term and long term memory to manage trajectory level information and within round context, improving recommendation consistency and interaction stability [17]. It further incorporates a contextual bandit for online preference adaptation, treating the clinician's final decision at the end of each adjustment cycle as feedback to update an individualized preference representation and guide subsequent recommendations. We validate this framework through retrospective ICU trajectory replay and expert review, showing improved recommendation acceptability and fewer interaction rounds to reach an acceptable recommendation [2].

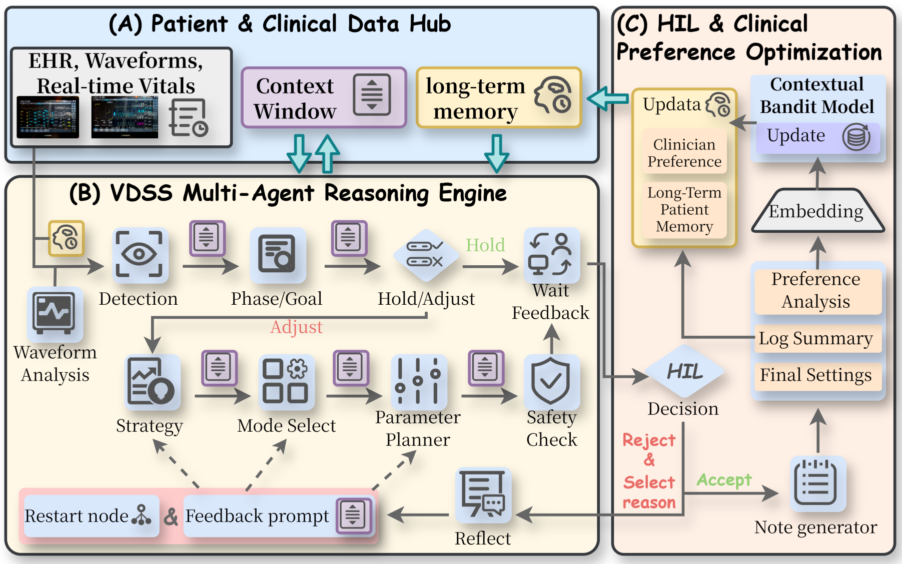

Fig. 1. Overview of VDSS. Panel **A** shows the patient and clinical data hub with a context window and long term memory. Panel **B** depicts the VDSS multi agent reasoning engine from detection and phase goal inference to hold adjust gating, planning, safety checking, and feedback driven revision. Panel **C** summarizes the human in the loop interface and online clinician preference optimization with a contextual bandit.

1. VDSS is introduced as a human in the loop multi agent ventilator decision support framework that enables verifiable and maintainable collaboration through modular decomposition and contract driven interfaces.
2. Layered short term and long term memory is integrated with an auditable evidence trail so that multi round recommendations remain coherent and are accompanied by traceable state summaries and rationale for structured review.
3. An online preference adaptation mechanism is developed based on a contextual bandit that uses cycle level clinician decisions as feedback to adapt to different tuning styles, improving acceptability and reducing interaction rounds.

2 Ventilator Decision Support System

2.1 VDSS Architecture: Multi-Agent Framework

VDSS organizes bedside ventilator titration as a contract governed multi-agent workflow rather than a single model that directly maps clinical inputs to ventilator recommendations. Here, agents refer to LLM or VLM reasoning modules with defined clinical roles, role specific inputs, and schema constrained outputs, while graph routing, mode compatibility validation, and safety checking are implemented as deterministic control components. At each adjustment cycle t, VDSS forms the decision input x_t from the current bedside state, current ventilator settings, short term context, and long term memory. Physiological and ventilatory measurements are treated as state variables, whereas clinician adjustable ventilator settings, including ventilation mode, PEEP, oxygen concentration, pressure support, inspiratory pressure, and respiratory rate, are treated as action variables. This design maintains a traceable path from evidence extraction to decision reasoning and clinician facing communication.

When waveform data are available, the ***Waveform Analyzer*** first converts rendered pressure and flow traces into structured bedside cues, including waveform quality, asynchrony patterns, suspicious events, observed patient state, and uncertainty, thereby complementing tabular signals with dynamic ventilatory evidence [23]. The ***Detection Agent*** integrates these waveform-derived cues with physiological and ventilatory measurements to identify abnormalities and generate a structured state summary. This summary is then interpreted by the ***Phase Goal Manager***, which infers the current treatment phase and specifies primary and secondary goals. Conditioned on the detected abnormalities and phase-consistent objectives, the ***Hold/Adjust Gate*** selects a branch decision $b_t \in \{\text{hold}, \text{adjust}\}$. This gate enables VDSS to avoid unnecessary titration when the current configuration is acceptable or when the available evidence is insufficient for safe adjustment.

The overall workflow is illustrated in Fig. 1. If $b_t = \text{hold}$, VDSS proceeds directly to cycle closure by generating a clinician-facing summary and structured record. If $b_t = \text{adjust}$, VDSS enters the active titration branch. The ***Strategy Selector*** determines the adjustment priority, the ***Mode Select*** agent evaluates whether a mode change is required under device- and mode-specific control semantics, and the ***Parameter Planner*** proposes a compact set of executable setting updates over allowable action variables. Before clinician review, proposed updates are passed through deterministic safety and compatibility checks, ensuring that generative planning remains bounded by explicit device, mode, and safety constraints. If the clinician rejects a recommendation, the ***Reflect Agent*** translates structured feedback into revision constraints and identifies the minimal decision layer to revisit, such as strategy selection, mode selection, or parameter planning. Once a recommendation is accepted, the ***Note Generator*** produces the final clinician-facing summary, stores an auditable cycle record, and extracts the preference signal used for subsequent preference adaptation.

2.2 Human in the Loop Interaction and Replanning

VDSS is structured around clinician confirmed adjustment cycles. After the reasoning engine produces candidate setting updates and they pass an initial *Safety* consistency check, the system pauses and presents the proposed ventilator mode and parameter updates together with the safety assessment and the current pref-

Algorithm 1 Cycle end preference update with a contextual bandit in VDSS

Input: clinician d, context x_t, current settings, long term memory
$k \leftarrow 1$, accepted $\leftarrow$ false
while $\neg$accepted and $k \leq K_{\max}$ **do**
  Generate proposal $a_t^{(k)}$, apply *Safety*, present for clinician review, collect $o_t^{(k)}$
  if $o_t^{(k)} =$ reject **then**
    Reflect Agent produces revision constraints and selects the minimal decision stage to revisit
    Replan from the selected stage; $k \leftarrow k + 1$
  else
    $\tilde{a}_t \leftarrow a_t^{(k)}$, $K_t \leftarrow k$, accepted $\leftarrow$ true
  end if
end while
Note Generator outputs log summary s_t and preference signal p_t
Update clinician preference state $\theta_d \leftarrow \text{BanditUpdate}(\theta_d;\, x_t, \tilde{a}_t, h_t, p_t)$
Write s_t and p_t to long term memory

erence context. This cycle level checkpoint matches bedside titration practice and makes clinician oversight an explicit component of the control loop [25].

For cycle t, interaction proceeds through rounds $k = 1, \ldots, K_t$. The proposal at round k is $a_t^{(k)}$, and clinician feedback is $o_t^{(k)} \in \{\text{accept}, \text{reject}\}$. The cycle ends at the first accepted proposal, yielding the final accepted configuration $\tilde{a}_t$:

$$K_t = \min\{k : o_t^{(k)} = \text{accept}\}, \qquad \tilde{a}_t = a_t^{(K_t)}. \tag{1}$$

Rejection is treated as a structured control signal. When a proposal is rejected, the system records the clinician rationale and the specific parameters at issue, and the *Reflect Agent* translates this feedback into revision constraints while localizing the minimal decision level to revisit. Replanning is then applied only to the implicated stage, with intermediate results reused whenever possible rather than rerunning the full pipeline. When waveform input is available, the *Waveform Analyzer* may be re executed to refresh evidence prior to revision so subsequent proposals remain grounded in current bedside cues. This targeted rollback mechanism keeps updates confined to the feedback relevant layer and reduces redundant inference across rounds.

Upon acceptance, the *Note Generator* closes the cycle by producing a structured summary of the final decision and extracting a clinician preference signal from the completed interaction. These outputs are written to long term memory to support trajectory continuity and subsequent preference adaptation, while the within cycle trace is retained for audit and for the cycle end preference update described next. Interaction is bounded by a maximum round budget $K_{\max}$ to maintain a controlled and reviewable process.

2.3 Online Preference Optimization with Contextual Bandit

VDSS adapts to clinician specific tuning styles via an online contextual bandit that is updated only after an adjustment cycle is resolved by clinician acceptance.

For each clinician d, the bandit maintains a preference state θ_d [11]. This state is provided to the reasoning engine as preference context and modulates candidate ranking and selection in subsequent cycles. At cycle closure, *Note Generator* produces two cycle level artifacts that are persisted in long term memory, namely a final patient log summary and a clinician preference signal.

Preference learning is defined over a fixed set of twelve long term preference categories, which correspond to the twelve bandit arms used for clinician modeling. These categories span preferences over adjustment granularity and adjustment style, as well as emphasis across clinical objectives and pragmatic workflow choices commonly observed in bedside titration. Specifically, the categories include preferences for mode level changes versus staying within the current mode, conservative small step updates versus target driven assertive adjustments, and prioritization of oxygenation, ventilation and acid base control, lung protection, hemodynamics, synchrony and comfort, and weaning. They further encode a tendency to focus on a single key parameter first and a tendency to defer action when available data or context are insufficient.

Let x_t denote the cycle context assembled from the data hub and memory, and let $h_t = \{(a_t^{(k)}, o_t^{(k)})\}_{k=1}^{K_t}$ denote the full interaction trace for cycle t. After a proposal is accepted, the final accepted settings $\tilde{a}_t$ and the preference signal p_t produced at cycle closure define a single cycle end update:

$$\theta_d \leftarrow \mathrm{BanditUpdate}(\theta_d;\, x_t, \tilde{a}_t, h_t, p_t)\,. \tag{2}$$

This update is driven by the clinician approved outcome and the full cycle trace rather than intermediate rejected proposals, aligning preference adaptation with bedside decision making.

3 Experiments

We evaluate VDSS using retrospective ICU trajectory replay with cycle based clinician review. The dataset is collected from a multi center ICU cohort and includes 1309 structured records and 7447 ventilator setting entries covering 13 ventilation modes across two ventilator brands, with the four most frequent modes accounting for 83.4% of all entries. We quantify next step replay by predicting the next clinically recorded setting update from the current bedside state and settings, reporting normalized error metrics and averaged R^2.

Table 1 summarizes performance across replay and expert review. Across backbones, deploying models within VDSS yields substantially better replay accuracy and clinician ratings than direct single model generation. For example, under GPT-5.2, VDSS reduces MSE from 0.343 to 0.102 while improving the overall rating from 2.63 to 4.11. The ablations show consistent degradation when removing waveform evidence or preference context, with MSE increasing to 0.145 (NoImg) or 0.128 (NoPref) and overall ratings dropping to 3.62 and 3.89, respectively, indicating that both components contribute to more stable recommendations and clearer clinician facing justification. Similar gains are ob-

Table 1. Main results. Replay metrics are computed on the full dataset and clinician ratings (1–5) are from expert review of 100 cycles for single model baselines, VDSS, and VDSS ablations.

	Replay metrics			Clinician ratings (1–5)			
Method	**MSE** $\downarrow$	**MAE** $\downarrow$	R^2 $\uparrow$	**Accept.**	**Safety**	**Clarity**	**Overall**
Single model end to end							
Qwen3-VL-8B	0.486	0.595	-0.132	2.45	2.11	2.26	2.27
Qwen3-VL-235B	0.390	0.426	0.019	2.53	2.17	2.19	2.30
Gemini 3 Pro	0.362	0.386	0.130	2.80	2.31	2.24	2.45
GPT-5.2	0.343	0.381	0.153	2.75	2.43	2.71	2.63
VDSS (same backbone, multi agent)							
Qwen3(VL)-8B (VDSS)	0.175	0.210	0.354	3.30	4.12	3.56	3.22
Qwen3-VL-235B (VDSS)	0.130	0.183	0.661	3.51	4.20	3.79	3.59
Gemini 3 Pro (VDSS)	0.124	0.166	0.693	3.80	4.30	3.91	3.82
GPT-5.2 (VDSS)	**0.102**	**0.162**	**0.743**	**4.09**	**4.46**	**4.15**	**4.11**
Ablations (GPT-5.2 within VDSS)							
NoImg	0.145	0.194	0.660	3.69	4.26	3.75	3.62
NoPref	0.128	0.183	0.707	3.81	4.32	3.95	3.89
NoImgNoPref	0.158	0.202	0.620	3.50	4.15	3.61	3.46

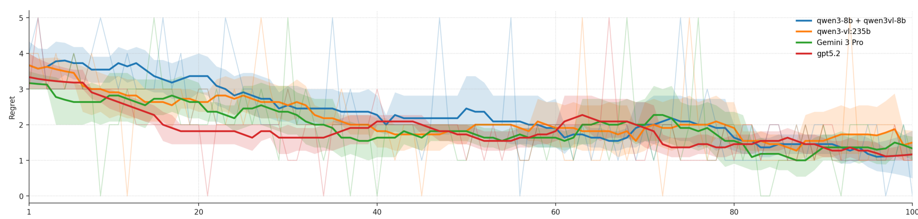

Fig. 2. Regret over 100 cycles from a single clinician, defined as $K_t - 1$ and capped by $K_{\max}$, shows an overall downward trend.

served across other backbones, suggesting that the improvement is driven by the VDSS workflow rather than any single model.

We do not emphasize conventional reinforcement learning because ventilation modes differ substantially in parameter sets and semantics, and per mode data are comparatively limited, which degrades generalization [12, 18]. The high risk setting further motivates a reviewable interaction loop with explicit safety checks and an auditable evidence trail, which is difficult to guarantee with end to end policies across heterogeneous modes.

We assess interaction efficiency through cycle level rejection dynamics on the same 100 cycle study. For each cycle t, interaction proceeds across rounds $k = 1, \dots, K_t$ and ends at the first accepted proposal. Regret within a cycle is the number of rejected proposals, $K_t - 1$, capped by the maximum round budget

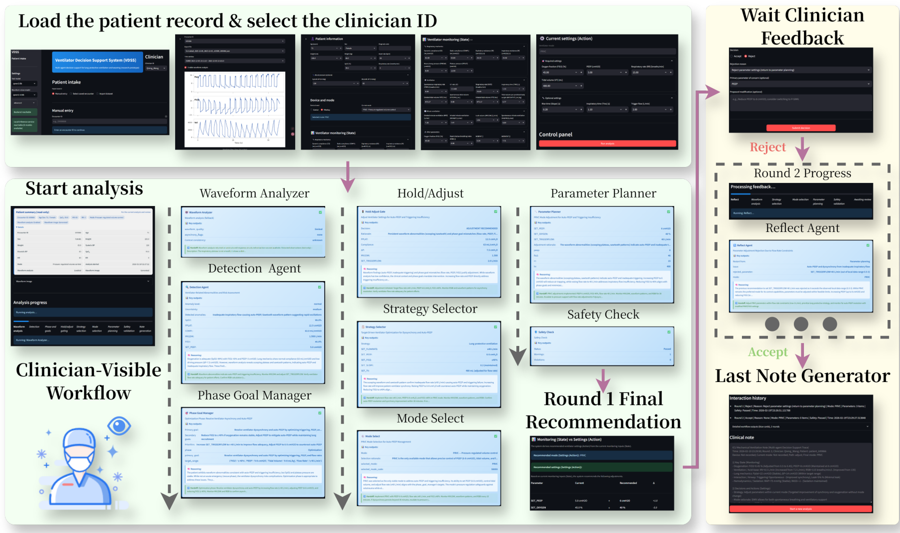

Fig. 3. Case study of one VDSS adjustment cycle.

$K_{\max}$. Fig. 2 shows an overall downward regret trend across model variants, consistent with improved interaction efficiency as preference context is updated.

We report runtime on an RTX 5070 Ti for local inference with Qwen3-8B and Qwen3-VL-8B. The workflow averages 305.3 s per cycle, and waveform analysis is the main bottleneck at 56.46 s. With external APIs, latency can vary with network conditions and server load. The framework maintains a completion failure rate below **7%**.

4 Case Study

Fig. 3 illustrates a complete adjustment cycle in the prototype interface. The clinician selects a clinician ID, loads a patient encounter, specifies the analysis time window, and enables waveform analysis when available. The patient is managed in PRVC with stable oxygenation and lung mechanics, while pressure and flow waveforms show a scooped inspiratory plateau and sawtooth patterns consistent with insufficient inspiratory flow and elevated auto PEEP risk. The system summarizes abnormalities and evidence, defines phase consistent objectives to improve synchrony and reduce auto PEEP risk while de escalating FiO2 when appropriate, and maintains PRVC while proposing a small set of safety checked parameter updates for review. The first proposal is rejected due to local feasibility constraints, which are converted into explicit revision constraints to replan from the implicated stage, yielding an updated recommendation accepted in the second round. The cycle ends with a structured note and an auditable interaction log written to long term memory, and the accepted outcome updates clinician specific preferences for subsequent cycles.

5 Conclusion

We present a human in the loop multi agent ventilator decision support framework that combines contract governed modules with traceable evidence and safety checked recommendations, enabling interpretable and reviewable bedside titration. Retrospective ICU trajectory replay and expert review show improved acceptability and fewer interaction rounds, indicating that the framework is effective in practice. An online contextual bandit updates clinician specific preferences at cycle closure, which improves alignment with individual tuning styles and supports consistent, preference aware recommendations under the same safety boundaries.

Acknowledgments. This work was supported by the Shanghai Municipal Science and Technology Commission (Grant No. 25692114700).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Bion, J.F., Abrusci, T., Hibbert, P.: Human factors in the management of the critically ill patient. British journal of anaesthesia **105**(1), 26–33 (2010)
2. Cox, C.E., White, D.B., Hough, C.L., Jones, D.M., Kahn, J.M., Olsen, M.K., Lewis, C.L., Hanson, L.C., Carson, S.S.: Effects of a personalized web-based decision aid for surrogate decision makers of patients with prolonged mechanical ventilation: a randomized clinical trial. Annals of internal medicine **170**(5), 285–297 (2019)
3. Dewes, R., Dimitrova, R.: Contract-based design and verification of multi-agent systems with quantitative temporal requirements. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 23152–23159 (2025)
4. Fan, E., Del Sorbo, L., Goligher, E.C., Hodgson, C.L., Munshi, L., Walkey, A.J., Adhikari, N.K., Amato, M.B., Branson, R., Brower, R.G., et al.: An official american thoracic society/european society of intensive care medicine/society of critical care medicine clinical practice guideline: mechanical ventilation in adult patients with acute respiratory distress syndrome. American journal of respiratory and critical care medicine **195**(9), 1253–1263 (2017)
5. Goodfellow, L.T., Miller, A.G., Varekojis, S.M., LaVita, C.J., Glogowski, J.T., Hess, D.R.: Aarc clinical practice guideline: patient-ventilator assessment. Respiratory Care **69**(8), 1042–1054 (2024)
6. Gray, B.M., Vandergrift, J.L., Barnhart, B.J., Reddy, S.G., Chesluk, B.J., Stevens, J.S., Lipner, R.S., Lynn, L.A., Barnett, M.L., Landon, B.E.: Changes in stress and workplace shortages reported by us critical care physicians treating coronavirus disease 2019 patients. Critical Care Medicine **49**(7), 1068–1082 (2021)

7. Henry, K.E., Kornfield, R., Sridharan, A., Linton, R.C., Groh, C., Wang, T., Wu, A., Mutlu, B., Saria, S.: Human–machine teaming is key to ai adoption: clinicians' experiences with a deployed machine learning system. NPJ digital medicine **5**(1), 97 (2022)
8. Jivraj, N.K., Hill, A.D., Shieh, M.S., Hua, M., Gershengorn, H.B., Ferrando-Vivas, P., Harrison, D., Rowan, K., Lindenauer, P.K., Wunsch, H.: Use of mechanical ventilation across 3 countries. JAMA internal medicine **183**(8), 824–831 (2023)
9. Jouvet, P., Hernert, P., Wysocki, M.: Development and implementation of explicit computerized protocols for mechanical ventilation in children. Annals of intensive care **1**(1), 51 (2011)
10. Kollef, M.H., Shapiro, S.D., Clinkscale, D., Cracchiolo, L., Clayton, D., Wilner, R., Hossin, L.: The effect of respiratory therapist-initiated treatment protocols on patient outcomes and resource utilization. Chest **117**(2), 467–475 (2000)
11. Lei, H., Lu, Y., Tewari, A., Murphy, S.A.: An actor-critic contextual bandit algorithm for personalized mobile health interventions. arXiv preprint arXiv:1706.09090 (2017)
12. Liu, S., See, K.C., Ngiam, K.Y., Celi, L.A., Sun, X., Feng, M.: Reinforcement learning for clinical decision support in critical care: comprehensive review. Journal of medical Internet research **22**(7), e18477 (2020)
13. Liu, X., Rivera, S.C., Moher, D., Calvert, M.J., Denniston, A.K., Ashrafian, H., Beam, A.L., Chan, A.W., Collins, G.S., Deeks, A.D.J., et al.: Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the consort-ai extension. The Lancet Digital Health **2**(10), e537–e548 (2020)
14. Muñoz, J., Ruíz-Cacho, R., Fernández-Araujo, N.J., Candela, A., Visedo, L.C., Muñoz-Visedo, J.: Artificial intelligence in the management of patient-ventilator asynchronies: A scoping review. Heart & Lung **73**, 139–152 (2025)
15. Murali, M., Ni, M., Karbing, D.S., Rees, S.E., Komorowski, M., Marshall, D., Ramnarayan, P., Patel, B.V.: Clinical practice, decision-making, and use of clinical decision support systems in invasive mechanical ventilation: a narrative review. British Journal of Anaesthesia **133**(1), 164–177 (2024)
16. Network, A.R.D.S.: 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 (2000)
17. Packer, C., Fang, V., Patil, S., Lin, K., Wooders, S., Gonzalez, J.: Memgpt: towards llms as operating systems. (2023)
18. Peine, A., Hallawa, A., Bickenbach, J., Dartmann, G., Fazlic, L.B., Schmeink, A., Ascheid, G., Thiemermann, C., Schuppert, A., Kindle, R., et al.: Development and validation of a reinforcement learning algorithm to dynamically optimize mechanical ventilation in critical care. NPJ digital medicine **4**(1), 32 (2021)
19. Ranieri, V.M., Rubenfeld, G.D., Taylor Thompson, B., Ferguson, N.D., Caldwell, E., Fan, E., Camporota, L., Slutsky, A.S.: Acute respiratory distress syndrome: the berlin definition. JAMA: Journal of the American Medical Association **307**(23) (2012)
20. Rosenbacke, R., Melhus, Å., McKee, M., Stuckler, D.: How explainable artificial intelligence can increase or decrease clinicians' trust in ai applications in health care: systematic review. Jmir Ai **3**, e53207 (2024)
21. Shi, F., Li, Y., Zhao, Y.: How do nurses manage their work under time pressure? occurrence of implicit rationing of nursing care in the intensive care unit: A qualitative study. Intensive and Critical Care Nursing **75**, 103367 (2023)
22. Slutsky, A.S., Ranieri, V.M.: Ventilator-induced lung injury. New England Journal of Medicine **369**(22), 2126–2136 (2013)

23. Thille, A.W., Rodriguez, P., Cabello, B., Lellouche, F., Brochard, L.: Patient-ventilator asynchrony during assisted mechanical ventilation. Intensive care medicine **32**(10), 1515–1522 (2006)
24. Topol, E.J.: High-performance medicine: the convergence of human and artificial intelligence. Nature medicine **25**(1), 44–56 (2019)
25. Vasey, B., Nagendran, M., Campbell, B., Clifton, D.A., Collins, G.S., Denaxas, S., Denniston, A.K., Faes, L., Geerts, B., Ibrahim, M., et al.: Reporting guideline for the early stage clinical evaluation of decision support systems driven by artificial intelligence: Decide-ai. bmj **377** (2022)
26. von Wedel, D., Redaelli, S., Suleiman, A., Wachtendorf, L.J., Fosset, M., Santer, P., Shay, D., Munoz-Acuna, R., Chen, G., Talmor, D., et al.: Adjustments of ventilator parameters during operating room–to–icu transition and 28-day mortality. American Journal of Respiratory and Critical Care Medicine **209**(5), 553–562 (2024)
27. Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Liu, J., et al.: Autogen: Enabling next-gen llm applications via multi-agent conversations. In: First conference on language modeling (2024)