Development and validation of a reinforcement learning algorithm to dynamically optimize mechanical ventilation in patients with moderate-severe acute respiratory distress syndrome: A multicenter observational study

Hui-ping Li, MD, PhD,<sup>1,2,3</sup> Jian Huo, M. Sc<sup>4</sup> Ying Liu, PhD,<sup>5</sup> Zhen Li, PhD,<sup>6,7</sup> An-min Hu, MD,<sup>2,3,8,9\*</sup>

<sup>1</sup>Department of Critical Care Medicine, Shenzhen People's Hospital, Shenzhen, China; <sup>2</sup>The Second Clinical Medical College, Jinan University; Shenzhen, China; <sup>3</sup>First Affiliated Hospital, Southern University of Science and Technology, Shenzhen, China; <sup>4</sup>Boston Intelligent Medical Research Center, Shenzhen United Scheme Technology Co., Ltd, Boston Ma, United States; <sup>5</sup>Department of Computer Science, the University of Hong Kong, Hongkong, China; <sup>6</sup>Shenzhen Research Institute of Big Data, Shenzhen, China; <sup>7</sup>School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen, China; <sup>8</sup>Department of Anesthesiology, Shenzhen People's Hospital, Shenzhen, China; <sup>9</sup>Biomedical Engineering Center, Shenzhen United Scheme Technology Co., Ltd, Shenzhen, China.\*Corresponding authors at: 23 Tianbei 1st St, Wensheng Suite 1102, Shenzhen, 518020, China. *E-mail addresses:* anmin@szus.org (A. Hu).

#### Abstract

**Background:** To enhance the outcomes for moderate to severe forms of Acute Respiratory Distress Syndrome (ARDS), a personalized lung-protective ventilation approach may be necessary. The challenge lies in personalizing medicine due to the heterogeneous nature of lung stress/strain.

Objective: Our goal is to propose a dynamic ventilation regimen for ARDS patients via three reward frameworks.

**Methods:** This research is based on a retrospective analysis of datasets from MIMIC and eICU databases. AI systems are used in the framework of this study to optimize ideal body weight-adjusted tidal volume (Vt) and positive end-expiratory pressure (PEEP). The systems rely on reinforcement learning and focus on prognosis, platform pressure, and driving pressure.

**Results:** Our study incorporates 16,487 patients with moderate to severe ARDS. On average, the reward from the AI systems' treatment selection is consistently higher than that of human clinicians. The AI systems were more likely to adjust Vt by roughly 1.5 times and PEEP by about 2.0 times compared to human clinicians. Notably, the lowest mortality was observed in ARDS patients when the actual doses aligned with the AI-platform's decisions.

**Conclusions:** Our model provides individualized and easily interpretable treatment decisions for moderate to severe ARDS patients, potentially enhancing their prognosis. The reward framework based on platform pressure and driving pressure may be a promising avenue for future research in reinforcement learning models of mechanical ventilation.

Trial registration: Not applicable.

Keywords: Acute Respiratory Distress Syndrome; Mechanical Ventilation; Reinforcement Learning; Prognosis; Platform Pressure, Driving Pressure

#### Introduction

Acute respiratory distress syndrome (ARDS) is a common condition among patients who require invasive mechanical ventilation, particularly those with moderate or severe ARDS (1). Ventilator-derived parameters such as inappropriate positive end-expiratory pressure (PEEP), plateau pressures, or driving pressures are associated with increased hospital mortality (2-4). To minimize the risk of ventilator-induced lung injury (VILI), it is recommended that patients with ARDS undergo lung protective ventilation strategies (3, 5-9). However, even with lung protective ventilation, ARDS patients may still be exposed to VILI (10-12).

Personalized ventilator management involves assessing the patient-specific risk of VILI and

Li et al.

weighing the potential risk of interventions aimed at mitigating VILI (13). It is important to note that some clinicians fail to consistently make decisions based on the best evidence (14). While artificial intelligence (AI) models have mainly focused on predicting onset or aiding in mechanical ventilation weaning, there has been limited focus on clinical decision support (15).

To improve patient outcomes and reduce the mortality rate of ARDS patients, we have developed three sets of ventilation strategies based on the conditions of prognosis, platform pressure, and driving pressure. These strategies, which are based on the reinforcement learning algorithm, dynamically adjust tidal volume and PEEP to provide doctors with the best mechanical ventilation program. A small reduction in the mortality rate of ARDS patients has the potential to save tens of thousands of lives worldwide each year.

## Methods

## Data source

The Laboratory for Computational Physiology at the Massachusetts Institute of Technology is responsible for maintaining the Multiparameter Intelligent Monitoring in Intensive Care (MIMIC)-III (version 1.4) database, the MIMIC-IV database, and the eICU Collaborative Research Database (16-18). Access to these databases is granted to researchers who have completed training in the protection of human subjects.

## Study population and stratification

The inclusion criteria for this study were: (1) adults aged 18 years or older upon admission to the ICU; (2) use of invasive mechanical ventilation (MV) for a minimum of 72 hours during the ICU stay; (3) a moderate-severe ARDS diagnosis according to the Berlin definition during the first 24 hours of ICU admission (19); (4) documented volume-controlled ventilation with tidal volume (Vtset) and positive end-expiratory pressure (PEEP); and (5) documentation of in-hospital mortality, platform pressure, and driving pressure.

#### Ethics statement

Approval of data collection, processing, and release for the MIMIC-III and MIMIC-IV been granted by the Institutional Review Boards of Beth Israel Deaconess Medical Center (Boston, MA, USA) and the Massachusetts Institute of Technology (Cambridge, MA, USA). Approval of data collection, processing and release for the eICU database has been granted by the eICU research committee and exempt from Institutional Review Board approval. Because this study was a secondary analysis of fully anonymized data, individual patient consent was not required.

## **Data extraction**

The extraction of data was carried out using customized scripts written in Standardized Query Language (SQL) for the MIMIC and eICU databases, and performed on the PostgreSQL object-relational database system. The eICU research committee approved the collection, processing, and release of data from the eICU database and it was exempt from Institutional Review Board approval.

## **Preprocessing steps**

The data for this study was collected from the first 72 hours after the onset of moderate-to-severe ARDS during mechanical ventilation to capture the early phase of its management (20, 21). From both MIMIC-III and eICU databases, 49 variables were extracted, including demographics, vital signs, laboratory values, fluid balance, comorbidities, medical scores, and hospital characteristics. The patient data was coded as multidimensional discrete time series with 4-hour time increments. For variables with multiple measurements within a 4-hour time step, the values were either averaged (e.g., heart rate) or summed (e.g., urine output) as appropriate. Comorbidities were determined based on diagnoses recorded during the hospitalization (22).

The data was checked for outliers and errors using frequency histograms and univariate statistical approaches (Tukey's method). Any errors were corrected when possible, and parameters were capped to clinically plausible values. Missing data was imputed using the Multivariate Imputation by Chained Equations (MICE) method.

## **Building the computational model**

The physiological state of a patient can only be partially represented by the available data, making the disease process a partially observable Markov Decision Process (MDP). An MDP was utilized to approximate the patient's trajectory and model the decision-making process (23, 24). Our problem was projected as a MDP defined by the 4-tuple  $(S, A, T, R, \gamma)$  in the following sections.

Every 4 hours, S is defined as a finite set of states, summarizing a patient's clinical characteristics through clustering of their 44 clinical features data fingerprint. A is the set of available actions from a given state s. R represents the reward signal received as feedback after transitioning to a defined state. The transition matrix T(s', s, a) contains the probability of moving from state s to s' at time t + 1, through action a and describes the dynamics of the system. The discount factor  $\gamma$  models the fact that future rewards are worth less than immediate rewards and a high discount factor will result in higher values for rewards received earlier in the decision-making process compared to those received later.

The state space was defined in the clustering procedure by clustering all patient time series from the three datasets using K-means clustering. To ensure a highly granular model and avoid using a too large state space, we used Bayesian and Akaike information criteria to determine the optimal number of clusters (as described in Supplement file 3 and Supplement file 4). This prevented the state space from having sparsely populated states.

Before clustering, the data was pre-processed to account for unequal means and variances. Normal data was standardized and log-normal data was log-transformed before standardization. Factor variables did not need to be pre-processed (25). The normality of each variable was tested using visual methods such as quantile-quantile plots and frequency histograms.

The goals of a mechanical ventilation regime are the reduction of VILI while maintaining adequate oxygenation and decarboxylation. Consequently, we focused on a total of two parameters to be included in the action space, influencing these overall goals: Ideal body weight-adjusted (target) Vtset and PEEP. Ideal body weight-adjusted Vt was calculated relative to a predicted body weight for males as 50 +  $(0.91 \times \text{[height in centimeters -}$ 152.4]) and for females as  $45.5 + (0.91 \times \text{[height in]})$ centimeters - 152.4]). As a result,  $\mathcal{A}$  is the finite number of possible actions at any given state based on a combination of the two aforementioned parameters. Based on frequency analysis, Vtset and PEEP converted into discrete decisions of seven treatment levels. This results in a two-dimensional action space of 49 discrete actions. It is worth mentioning that there was no option of a zero policy and the algorithm always had to decide towards one ventilation policy.

The formulation of reward function is the key in successful applications of RL approaches. We have designed three sets of reward systems and named Prognosis-AI, Platform-AI, and Driving-AI (26-31). Mortality remains the most widely accepted endpoint for ARDS trials. In Prognosis-AI models, we used hospital mortality as the sole defining factor for the system-defined penalty and reward. When a patient survived, a positive reward was released at each patient's trajectory (a 'reward' of +1); a negative reward (a 'penalty' of -1) was issued if the patient died. In Platform-AI models, we used plateau pressure as the defining factor for the system-defined penalty and reward. The positive reward was released when a patient with plateau pressure  $\leq 25$  cm H<sub>2</sub>O (a 'reward' of +1); the negative reward was issued if the patient with plateau pressure >25 cm H<sub>2</sub>O (a 'reward' of 1-0.2\*(plateau pressure-25)). In Driving-AI models, the defining factor for the system-defined penalty and reward were driving pressure. The positive reward was released if a patient with driving pressure  $\leq 15$  cm H2O (a 'reward' of + 1); the negative reward was issued when the patient with driving pressure >15 cm H2O (a 'reward' of 1-0.2\*(driving pressure-15)).

Li et al.

In reinforcement learning, a policy  $\pi$  corresponds to a set of rules dictating which action is taken while in a particular state (23, 24). Each MDP determines a state-action value function  $Q^{\pi}$ , that reflects the expected sum of discounted rewards for choosing an action while in a particular state, and following a policy  $\pi$  thereafter.  $Q^{\pi}$  summarizes the effect of the treatment decisions on the patient's outcomes, with beneficial decisions having positive  $Q^{\pi}$  values and harmful decisions negative  $Q^{\pi}$ values.

### Estimation of the AI policy

We learned a theoretical optimal policy for the MDP using in-place policy iteration, which identified the decisions that maximize the sum of rewards, hence the expected survival of patients (24). Policy iteration started with a random policy that was iteratively evaluated and then improved until converging to an optimal solution (32).

### **Evaluation of clinicians' actions**

Temporal difference (TD) learning is a model-free algorithm, in which learning happens through the iterative correction of your estimated returns towards a more accurate target return (24). We stopped the evaluation after processing 500,000 patient trajectories with resampling, which is when the value of the estimated policy reached an asymptote.

### Model evaluation

We generated 500 different models from various random splits (80%) of the three datasets. In each model, K-means clustering is performed to instantiate a different state space. Based on the euclidean distance to the nearest cluster centroid, state membership and corresponding action for test set data points is determined. Furthermore, we implemented a type of HCOPE method, WIS, and used bootstrapping to estimate the true distribution of the policy value in the test sets (33, 34). The approach is adopted in wide range of high-risk applications.

In each model, we also estimated the value of a random policy for comparison. As recommended, the selected final model maximizes the 95% confidence lower bound of the AI policy among the 500 candidate models (33). We demonstrate that this bound consistently exceeded the 95% confidence upper bound of the clinicians' policy, provided that enough models were built.

We also measured the performance of the AI policy using direct indicators and analyzed patient outcomes as a function of the gap between clinicians and AI policies (32, 35). Here, we analyzed patient mortality in the test sets for which the respiratory parameters actually administered corresponded to or differed from the respiratory parameters suggested by the AI policy.

### Results

This study was conducted and reported in accordance CONSORT-AI extension with guidelines (36). Our RL models were built and validated on three large nonoverlapping ICU databases containing data routinely collected from adult patients in the United States (Figure 1). After exclusion ineligible cases, 16,487 of moderate-severe ARDS patients were ultimately included in this study, including 6,348 patients from the eICU dataset, 6752 patients from the MIMIC-III dataset, and 3387 patients from the MIMIC-IV dataset. Patient clinical and demographic properties are shown in Table 1.



**Figure 1.** Selection of the best AI policy and model calibration. The shades represent the 95% upper-lower-bound of the estimated performance of policy. The green line, orange line, red line represents the mean reward for AI policy, clinicians' policy, and random AI policy in turn. The dotted orange line represents the mean reward for the overall return of the historical doctor's original strategy. (a) Prognosis policy. (b) Platform policy. (c) Driving policy.



**Figure 2**. Visualization of the clinician and AI action. All actions were aggregated over all time steps for the seven treatment levels. One to seven of tidal volume action space are 0-2.5 ml/Kg, 2.5-5.0 ml/Kg, 5.0-7.5 ml/Kg, 7.5-10.0 ml/Kg, 10.0-12.5 ml/Kg, 12.5-15.0 ml/Kg, and >15.0 ml/Kg; one to seven of Positive end-expiratory pressure are 0-5 cmH<sub>2</sub>O, 6-7 cmH<sub>2</sub>O, 8-9 cmH<sub>2</sub>O, 10-11 cmH<sub>2</sub>O, 12-13 cmH<sub>2</sub>O, 14-15 cmH<sub>2</sub>O, and >16 cmH<sub>2</sub>O. (a) Clinicians' action. (b) Prognosis-AI action. (c) Platform-AI action. (d) Driving-AI action.

Table 1. Clinical and demographic properties of the study population.

Clinical Variable*	eICU dataset	MIMIC-III	MIMIC-IV	ΔΤΤ	
		dataset	dataset	ALL	
Number Of Included Patients	6348	6752	3387	16487	
APACHE-III scoring	61 (45-79)	52 (39-68)	82 (62-102)	60 (44-81)	
Patient Complications, No.					
(%)					
Congestive heart failure	292 (4.6)	1030 (15.3)	318 (9.4)	1640 (9.9)	
Chronic pulmonary disease	755 (11.9)	1620 (24)	390 (11.5)	2765 (16.8)	
Mechanical Ventilation					

			Artificial inte	lligence mechanical ventilation
Tidal volume, mL	500	527	450	500
	(430-520)	(460-605)	(397-503)	(430-550)
Respiratory rate, counts per minute	16 (14-20)	18 (15-23)	20 (16-24)	18 (14-22)
Positve end-expiratory pressure, cmH <sub>2</sub> O	5 (5-8)	5 (5-8)	5 (5-10)	5 (5-8)
First mechanical ventilation, No. (%)	4543 (71.6)	4216 (62.4)	1899 (56.1)	10658 (64.6)
Mechanical ventilation	137	149	127 (02 101)	139 (97-223)
time, hours	(95-218.2)	(103-242)	127 (92-191)	
Clinical Outcome				
ICU stay, hours	276.1	366	400.6	335.6
	(176.9-453.9)	(232.3-606.4)	(257-635.1)	(211.1-554.2)
Hospital stay, hours	439.7	542.6	569.8	501.6
	(282.9-666.2)	(347.8-821.1)	(367.5-895.7)	(325.1-778)
Hospitalized death, No(%)	1710 (26.9)	1861 (27.6)	1149 (33.9)	4720 (28.6)

Abbreviations: APACHE-III score, the acute physiology and chronic health evaluation III score.

\*Data shown as mean  $\pm$  standard deviation, number (percent), or median (interquartile range) as appropriate.

To conservatively evaluate the differences in performance in the 20% test dataset (58,339 mechanical ventilation events), we compared the 95% bound of the performance of the best AI policy, clinicians' policy, and random AI policy (Figure 1). We found that the reward of the best AI policy outperformed the reward of the clinicians' policy. In addition, there is a skewed distribution in the reward values in the best model of prognosis-AI, platform-AI, and driving-AI.

Approximately  $45 \cdot 18\%$  of the tidal volumes in the ventilation strategies used in historical patients were 5-7.5 ml/kg, and  $42 \cdot 43\%$  were 7.5-10 ml/kg (Figure 2A and Supplement file 5). A total of 54.31% of patients received PEEP at 0-5 cmH<sub>2</sub>O, 15.05% of patients received PEEP at 8-9 cmH<sub>2</sub>O, and 15.03% of patients received PEEP at 10-11 cmH<sub>2</sub>O.

The recommended tidal volume in the prognostic model was 2.5-5 ml/kg in 18.93% of patients, 5-7.5 ml/kg in 22.16% of patients, 7.5-10 ml/kg in 25.07% of patients, and 10-12.5 ml/kg in 22.03% of patients (Figure 2B). The PEEP recommendations by the prognostic model were 5

cmH<sub>2</sub>O in 28.68%, 6-7 cmH<sub>2</sub>O in 13.41%, 8-9 cmH<sub>2</sub>O in 17.43%, 10-11 cmH<sub>2</sub>O in 10.94%, and 12-13 cmH<sub>2</sub>O in 17.24%.

The recommended tidal volume for the plateau pressure model was  $2 \cdot 5 - 5 \text{ ml/kg}$  for  $29 \cdot 05\%$ ,  $5 - 7 \cdot 5 \text{ ml/kg}$  for  $33 \cdot 70\%$ ,  $7 \cdot 5 - 10 \text{ ml/kg}$  for  $21 \cdot 20\%$ , and  $10 - 12 \cdot 5 \text{ ml/kg}$  for  $12 \cdot 74\%$  (Figure 2C). Regarding the recommended PEEP in the plateau pressure model,  $44 \cdot 28\%$  of the patients received PEEP at 0-5 cmH<sub>2</sub>O,  $23 \cdot 88\%$  of the patients received PEEP at 6-7 cmH<sub>2</sub>O,  $13 \cdot 72\%$  of the patients received PEEP at 6-7 cmH<sub>2</sub>O,  $9 \cdot 69\%$  of the patients received PEEP at 10-11 cmH<sub>2</sub>O, and  $5 \cdot 30\%$  of the patients received PEEP at 10-11 cmH<sub>2</sub>O, and  $5 \cdot 30\%$  of the patients received PEEP at Clinician recommended a lower Vt and a lower PEEP than the clinicians' actual treatments.

The tidal volume in the driving pressure model was  $2 \cdot 5 - 5$  ml/kg for  $12 \cdot 30\%$ ,  $5 - 7 \cdot 5$  ml/kg for  $40 \cdot 35\%$ ,  $7 \cdot 5 - 10$  ml/kg for  $24 \cdot 46\%$ , and  $10 - 12 \cdot 5$  ml/kg for  $15 \cdot 46\%$  (Figure 2D). The recommended patient PEEP in the driving pressure model was 0 - 5 cmH<sub>2</sub>O in  $8 \cdot 70\%$ , 6 - 7 cmH<sub>2</sub>O in  $18 \cdot 02\%$ , 8 - 9 cmH<sub>2</sub>O in  $16 \cdot 82\%$ , 10 - 11 cmH<sub>2</sub>O in  $17 \cdot 08\%$ , and 12 - 13 cmH<sub>2</sub>O in  $17 \cdot 75\%$ .



**Figure 3.** The probability of the clinician and AI action changes. The probability of parameters changes is shown in mechanically ventilated of ARDS patients at each 4 h time step. (a) Volume action changes. (b) PEEP action.



Figure 4. Visualization of the difference between the given and AI-policy suggested actions in hospital mortality. The action excess, referring to the difference

between the given and suggested respiratory parameters averaged over all time points per patient. The probability of parameters changes is shown in mechanically ventilated of ARDS patients at each 4 h time step. The figure is built by bootstrapping with 3000 resamplings. (a) The difference of

7

volume action changes. (b) The difference of PEEP action changes.

Through our analysis of the frequency of changes in breathing parameters at 4-hour intervals, we found that the overall adjustment probability of tidal volume among clinicians was approximately 0.17, while the overall adjustment probability of the three artificial intelligences was approximately 0.3. The overall adjustment probability of the AI model to respiratory parameters was approximately 1.5 that of the clinician policy (Figure 3A). The overall adjustment probability of PEEP by clinicians is approximately 0.13, while the adjusted probability of the three AI actions is approximately 0.3. The overall adjustment probability of the AI model to respiratory parameters was approximately 0.3. The overall adjustment probability of the AI model to respiratory parameters was approximately 0.3. The overall adjustment probability of the AI model to respiratory parameters was approximately twice that of the clinician policy (Figure 3B).

Compared with the actual clinical strategy, we compared the relationship between tidal volume and PEEP between the best AI strategy and mortality under the three reward strategies (Figure 4). Compared with the Plateau-AI strategy, the tidal volume strategies increased the mortality rate of the patients (Figure 4A). In terms of clinical strategy, the mortality rate of ARDS patients who underwent the tidal volume strategy was lower than that of the patients who underwent the driving pressure AI strategy. Compared with historical actions, the prognostic AI strategy and the platform pressure AI strategy adjust PEEP parameters to reduce the mortality of ARDS patients (Figure 4B).

## Discussion

We demonstrated the application of reinforcement learning to customize clinically understandable mechanical ventilation strategies for ARDS patients. Our findings indicate that patients treated with the recommendations of AI clinicians had the lowest mortality rate. Our reward framework for AI mechanical ventilation, which is based on airway pressure, serves as a guide for future research on reinforcement learning models for mechanical ventilation.

The use of an artificial intelligence model for mechanical ventilation allows for the adjustment of parameters based on various clinical characteristics (37). By utilizing a dynamic reinforcement learning model, clinicians can develop an optimized ventilation strategy for each individual patient (37). Predictive machine learning models have also been used to predict the risk of hospital mortality in mechanically ventilated ICU patients using data from the MIMIC-III database (38). A recent large prospective study across multiple centers suggested that higher PEEP levels may be associated with worse outcomes for critically ill patients with severe respiratory infection (27).

In our study, we found that the actions recommended by the Platform Pressure reward setting were more aligned with clinical perspectives compared to the Driving Pressure reward setting. The low tidal volume ventilation strategy was better reflected in the actions recommended by the Plateau Pressure. Our Prognosis-AI and Platform-AI policy models revealed that as the predicted Vt and PEEP values approached the actual values, the mortality rate decreased. On the other hand, death due to hypoventilation or VILI may occur if the Vt and PEEP are not appropriate.

Therefore, in the respiratory support for ARDS patients undergoing invasive mechanical ventilation, it is crucial to recognize the potential risks of hypoventilation and VILI and adopt methods to detect and adjust respiratory parameters. Our study offers effective respiratory support decision recommendations based on tidal volume and PEEP, which demonstrate good predictive accuracy in ARDS patients. Specifically, those patients whose actual tidal volume and PEEP settings align with the recommendations of the Prognosis-AI and Platform-AI models have better outcomes, while those whose actual settings differ from the recommended settings show poorer outcomes.

Our study has two key strengths. Firstly, it is the first to establish reinforcement learning models specifically focused on mechanically ventilated ARDS patients. Secondly, we utilized platform pressure and driving pressure as reward and punishment rules to construct an artificial intelligence decision model for mechanical ventilation. Our study has some limitations that should be noted. Firstly, the selection of 300 respiratory state clusters by the AIC and BIC methods limited our ability to simultaneously optimize multiple adjustable mechanical ventilation parameters. Secondly, the division of numerical respiratory parameters into equal parts may not accurately reflect the real-world scenario where breathing parameters may vary differently. Thirdly, the model's performance was based on historical patient data and requires further evaluation through prospective studies. Fourthly, the use of a large dataset and routinely collected clinical data led to the exclusion of some sites and patients due to poor recording quality or missing data.

Additionally, our findings showed that including all critically ill patients with mechanical ventilation in the model for breathing strategy did not align with the low tidal volume ventilation strategy for recommending actions for respiratory parameters in ARDS patients. However, our findings indicated that the recommended actions from modeling respiratory parameters based on historical ARDS patient data were aligned with existing clinical perspectives.

AI will not replace clinicians anytime soon. However, the model provides recommendations for setting respiratory parameters with multidimensional patient-specific information, enabling clinicians to phenotype patients, and leading the way for widespread use of precision medicine (39).

In conclusion, we have compared three decision models using reinforcement learning, which are based on different reward frameworks to automate the setting of breathing parameters for mechanical ventilation. The models integrate information on the state of mechanically ventilated patients and demonstrate the potential of AI respiratory support technology for invasive respiratory support in ARDS patients.

## Acknowledgments

We are thankful to the Laboratory of Computational Physiology at the Massachusetts Institute of

Technology and the eICU Research Institute for providing the data used in this research.

# Data sharing statement

The three databases used in this research, MIMIC-III MIMIC-IV and eICU are available for access, in part or in total, by relevant parties subject to abiding by their usage policies. To facilitate the reproduction of our results, we shall make fully anonymized data available on the figshare (https://figshare.com/s/1d731507eabb03176af8) from the publication of this manuscript.

## Contributors

According to the guidelines of the International Committee of Medical Journal Editors (ICMJE), all authors contributed to the four criteria. AMH conceived and designed the study. HJ and LHP acquired the data. LZ and AMH analyzed and interpreted the data. HJ and TZ drafted the manuscript. HPL and AMH critically revised the manuscript for valuable intellectual content. All authors read and approved the final manuscript.

## **Conflicts of Interest:**

The authors declare that they have no competing interests.

# Funding

National Science Foundation for Young Scientists of China (81801947), Guangdong Medical Science and Technology Research Fund Project (A2021058), Clinical Agent Research Fund of Guangdong Province (2022MZ12).

# Abbreviations

AI: artificial intelligence

ARDS: acute respiratory distress syndrome

MICE: multivariate imputation by chained equations

MDP: Markov decision process

MIMIC: multiparameter intelligent monitoring in intensive care

MV: mechanical ventilation

PEERP: positive end-expiratory pressure SQL: standardized query language

TD: temporal difference VILI: ventilator induced lung injury Vt: tidal volume

## References

1. Bellani G, Laffey JG, Pham T, et al. Epidemiology, Patterns of Care, and Mortality for Patients With Acute Respiratory Distress Syndrome in Intensive Care Units in 50 Countries. Jama 2016;315(8):788-800.

2. Silva PL, Rocco PRM. The basics of respiratory mechanics: ventilator-derived parameters. Ann Transl Med 2018;6(19):376.

3. Brower RG, Matthay MA, Morris A, et al. Ventilation with lower tidal volumes as compared with traditional tidal volumes for acute lung injury and the acute respiratory distress syndrome. N Engl J Med 2000;342(18):1301-1308.

4. Laffey JG, Bellani G, Pham T, et al. Potentially modifiable factors contributing to outcome from acute respiratory distress syndrome: the LUNG SAFE study. Intensive care medicine 2016;42(12):1865-1876.

5. Kacmarek RM, Villar J, Sulemanji D, et al. Open Lung Approach for the Acute Respiratory Distress Syndrome: A Pilot, Randomized Controlled Trial. Critical care medicine 2016;44(1):32-42.

6. Serpa Neto A, Cardoso SO, Manetta JA, et al. Association between use of lung-protective ventilation with lower tidal volumes and clinical outcomes among patients without acute respiratory distress syndrome: a meta-analysis. Jama 2012;308(16):1651-1659.

7. Serpa Neto A, Nagtzaam L, Schultz MJ. Ventilation with lower tidal volumes for critically ill patients without the acute respiratory distress syndrome: a systematic translational review and meta-analysis. Curr Opin Crit Care 2014;20(1):25-32.

8. Serpa Neto A, Simonis FD, Barbas CS, et al. Association between tidal volume size, duration of ventilation, and sedation needs in patients without acute respiratory distress syndrome: an individual patient data meta-analysis. Intensive care medicine 2014;40(7):950-957. 9. Neto AS, Simonis FD, Barbas CS, et al. Lung-Protective Ventilation With Low Tidal Volumes and the Occurrence of Pulmonary Complications in Patients Without Acute Respiratory Distress Syndrome: A Systematic Review and Individual Patient Data Analysis. Critical care medicine 2015;43(10):2155-2163.

10. Chiumello D, Carlesso E, Brioni M, et al. Airway driving pressure and lung stress in ARDS patients. Critical care (London, England) 2016;20:276.

11. Terragni PP, Rosboch G, Tealdi A, et al. Tidal hyperinflation during low tidal volume ventilation in acute respiratory distress syndrome. American journal of respiratory and critical care medicine 2007;175(2):160-166.

12. Bugedo G, Retamal J, Bruhn A. Driving pressure: a marker of severity, a safety limit, or a goal for mechanical ventilation? Critical care (London, England) 2017;21(1):199.

13. Wick KD, McAuley DF, Levitt JE, et al. Promises and challenges of personalized medicine to guide ARDS therapy. Critical care (London, England) 2021;25(1):404.

14. Morris AH. Human Cognitive Limitations. Broad, Consistent, Clinical Application of Physiological Principles Will Require Decision Support. Ann Am Thorac Soc 2018;15(Suppl 1):S53-s56.

15. Gallifant J, Zhang J, Del Pilar Arias Lopez M, et al. Artificial intelligence for mechanical ventilation: systematic review of design, reporting standards, and bias. Br J Anaesth 2022;128(2):343-351.

16. Johnson AE, Pollard TJ, Shen L, et al. MIMIC-III, a freely accessible critical care database. Scientific data 2016;3:160035.

17. Pollard TJ, Johnson AEW, Raffa JD, et al. The eICU Collaborative Research Database, a freely available multi-center database for critical care research. Scientific data 2018;5:180178.

18. Johnson A, Bulgarelli L, Pollard T, et al. MIMIC-IV (version 0.4). PhysioNet 2020.

19. Ranieri VM, Rubenfeld GD, Thompson BT, et al. Acute respiratory distress syndrome: the Berlin Definition. Jama 2012;307(23):2526-2533.

Artificial intelligence mechanical ventilation

20. Zhang R, Pan Y, Fanelli V, et al. Mechanical Stress and the Induction of Lung Fibrosis via the Midkine Signaling Pathway. American journal of respiratory and critical care medicine 2015;192(3):315-323.

21. Thille AW, Esteban A, Fernández-Segoviano P, et al. Comparison of the Berlin definition for acute respiratory distress syndrome with autopsy. American journal of respiratory and critical care medicine 2013;187(7):761-767.

22. Quan H, Sundararajan V, Halfon P, et al. Coding algorithms for defining comorbidities in ICD-9-CM and ICD-10 administrative data. Med Care 2005;43(11):1130-1139.

23. Krishnamurthy V. Partially observed Markov decision processes: Cambridge university press;2016.

24. Sutton RS, Barto AG. Reinforcement learning: An introduction: MIT press; 2018.

25. Szepannek G. clustMixType: User-Friendly Clustering of Mixed-Type Data in R. The R Journal 2018;10(2):200.

26. Spragg RG, Bernard GR, Checkley W, et al. Beyond mortality: future clinical research in acute lung injury. American journal of respiratory and critical care medicine 2010;181(10):1121-1127.

27. Sakr Y, Midega T, Antoniazzi J, et al. Do ventilatory parameters influence outcome in patients with severe acute respiratory infection? Secondary analysis of an international, multicentre14-day inception cohort study. Journal of critical care 2021;66:78-85.

28. Amato MB, Meade MO, Slutsky AS, et al. Driving pressure and survival in the acute respiratory distress syndrome. N Engl J Med 2015;372(8):747-755.

29. Goligher EC, Costa ELV, Yarnell CJ, et al. Effect of Lowering Vt on Mortality in Acute Respiratory Distress Syndrome Varies with Respiratory System Elastance. American journal of respiratory and critical care medicine 2021;203(11):1378-1385.

30. Matthay MA, Zemans RL, Zimmerman GA, et al. Acute respiratory distress syndrome. Nat Rev Dis Primers 2019;5(1):18.

31. Aoyama H, Pettenuzzo T, Aoyama K, et al. Association of Driving Pressure With Mortality Among Ventilated Patients With Acute Respiratory Distress Syndrome: A Systematic Review and Meta-Analysis. Critical care medicine 2018;46(2):300-306.

32. Komorowski M, Celi LA, Badawi O, et al. The Artificial Intelligence Clinician learns optimal treatment strategies for sepsis in intensive care. Nature medicine 2018;24(11):1716-1720.

33. Thomas P, Theocharous G, Ghavamzadeh M. High-confidence off-policy evaluation. In: 2015.

34. Hanna JP, Stone P, Niekum S. Bootstrapping with models: Confidence intervals for off-policy evaluation. In: 2017.

35. Prasad N, Cheng L-F, Chivers C, et al. A reinforcement learning approach to weaning of mechanical ventilation in intensive care units. arXiv preprint arXiv:170406300 2017.

36. Liu X, Cruz Rivera S, Moher D, et al. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. Lancet Digit Health 2020;2(10):e537-e548.

37. Peine A, Hallawa A, Bickenbach J, et al. Development and validation of a reinforcement learning algorithm to dynamically optimize mechanical ventilation in critical care. NPJ Digit Med 2021;4(1):32.

38. Su L, Li Y, Liu S, et al. Establishment and Implementation of Potential Fluid Therapy Balance Strategies for ICU Sepsis Patients Based on Reinforcement Learning. Front Med (Lausanne) 2022;9:766447.

39. Burki TK. Artificial intelligence hold promise in the ICU. Lancet Respir Med 2021;9(8):826-828.

Li et al.