

Advancing Safe Mechanical Ventilation Using Offline RL With Hybrid Actions and Clinically Aligned Rewards

Muhammad Hamza Yousuf^{*1†}, Jason Li^{*1}, Sahar Vahdati^{1,2,3}, Raphael Theilen⁴, Jakob Wittenstein⁴, Jens Lehmann^{1,5‡}

¹Institut für Angewandte Informatik (InfAI) e. V.

²TIB - Leibniz Information Centre for Science and Technology, Hannover, Germany

³Leibniz Universität Hannover

⁴Department of Anesthesiology and Intensive Care Medicine, University Hospital Carl Gustav Carus, Dresden, Germany

⁵Amazon

{yousuf, li, vahdati}@infai.org, {raphael.theilen, jakob.wittenstein}@ukdd.de, lehmann@infai.org

Abstract

Invasive mechanical ventilation (MV) is a life-sustaining therapy commonly used in the intensive care unit (ICU) for patients with severe and acute conditions. These patients frequently rely on MV for breathing. Given the high risk of death in such cases, optimal MV settings can reduce mortality, minimize ventilator-induced lung injury, shorten ICU stays, and ease the strain on healthcare resources. However, optimizing MV settings remains a complex and error-prone process due to patient-specific variability. While Offline Reinforcement Learning (RL) shows promise for optimizing MV settings, current methods struggle with the hybrid (continuous and discrete) nature of MV settings. Discretizing continuous settings leads to exponential growth in the action space, which limits the number of optimizable settings. Converting the predictions back to continuous can cause a distribution shift, compromising safety and performance. To address this challenge, in the IntelliLung project, we are developing an AI-based approach where we constrain the action space and employ factored action critics. This approach allows us to scale to six optimizable settings compared to 2-3 in previous studies. We adapt SOTA offline RL algorithms to operate directly on hybrid action spaces, avoiding the pitfalls of discretization. We also introduce a clinically grounded reward function based on ventilator-free days and physiological targets. Using multi-objective optimization for reward selection, we show that this leads to a more equitable consideration of all clinically relevant objectives. Notably, we develop a system in close collaboration with healthcare professionals that is aligned with real-world clinical objectives and designed with future deployment in mind.

Project Website — <https://intellilung-project.eu>

1 Introduction

One of the most frequently applied life-saving interventions in the ICU is invasive mechanical ventilation (MV) (Bel-

^{*}These authors contributed equally.

[†]Corresponding author.

[‡]Work done outside of Amazon.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

lani et al. 2016). The importance of MV was especially evident during the COVID-19 pandemic, marked by increased ICU admissions, prolonged ventilation, and early intubations. However, MV is also associated with an increased risk of organ damage, particularly ventilator-induced lung injury (VILI) (Slutsky 1999). To prevent VILI, clinical guidelines recommend the limitation of tidal volumes, respiratory rate and inspiratory pressures (Fichtner et al. 2017). However, these protocols only provide general guidance, leaving the actual choice of ventilator setting to the clinical judgment and expertise of healthcare providers. Furthermore, it was shown that protocols of protective MV are poorly followed worldwide (Bellani et al. 2016). MV also demands a high nurse-to-patient ratio, leading to suboptimal recovery and prolonged ICU stays in times of high workload (Roedl et al. 2019).

AI-based decision support systems (AI-DSS) can provide personalized MV treatment recommendations that reduce the risk of VILI while enhancing accessibility. Offline Reinforcement Learning (RL) algorithms have demonstrated potential to learn interventions from ICU datasets that ensure immediate patient stability and improve long-term outcomes during MV (Kondrup et al. 2023; Peine et al. 2021).

In this work, we focus on developing an AI-DSS for MV in the ICU using Offline RL. To ensure clinical relevance and practical deployability, we engage in close collaboration with clinicians and technical experts from multiple hospitals. This process shaped the identification of appropriate patient cohorts, the problem formulation, and the selection of state representations, action spaces, reward structures, and evaluation methodologies. This paper addresses several critical technical challenges from prior studies and highlights important concerns for real-world deployment. Our contributions are as follows:

C1. We introduce a reward based on ventilator-free days (VFD) combined with physiological parameter safe ranges, which better aligns with the goal of reducing VILI. Previous work focused on optimizing MV using sparse rewards based on mortality. However, prior medical studies (Torrini et al. 2021; Silva, Negrini, and Rocco 2015) indicate that

mortality alone can be a poor endpoint for evaluating MV interventions. We also show how multi-objective optimization can select a reward function that balances the contributions of different clinically relevant objectives.

C2. Previous studies often restrict the number of MV settings because the discrete-action space grows exponentially with the number of settings. We show a simple approach to reduce the action space and adopt an additive factored critic (Tang et al. 2023). This explicitly removes dangerous actions and reduces critic variance improving safety and performance.

C3. MV has both continuous and discrete settings. To avoid the pitfalls of discretizing continuous settings, we demonstrate how to adapt SOTA offline RL algorithms, namely Implicit Q-Learning (IQL) and Ensemble-Diversified Actor-Critic (EDAC), to handle (continuous and discrete) action spaces natively.

C4. Previous methods simplify continuous actions by discretizing them. However, during inference, these discrete outputs are converted back into continuous values using fixed rules or by clinicians selecting a value from the predicted bin based on their expertise. Our experiments show that this reconstruction introduces a distribution shift, potentially causing the learned policy to operate in regions where predictions are highly uncertain.

2 Related Work

In previous studies involving MV and RL, some have focused on binary decisions, such as whether to initiate MV (Lee et al. 2023), while others have addressed complex tasks including sedation and weaning strategies (Prasad et al. 2017; Yu, Ren, and Dong 2020) and optimizing MV for ICU patients (Peine et al. 2021; Kondrup et al. 2023; Liu et al. 2024; Roggeveen et al. 2024; Zhang, Qiu, and Tan 2024; Eghbali, Alhanai, and Ghassemi 2024).

Previous work on optimizing MV settings using offline RL either discretize the actions or purely use continuous actions. Studies that consider discretization (Peine et al. 2021; Kondrup et al. 2023; Liu et al. 2024; Roggeveen et al. 2024; Eghbali, Alhanai, and Ghassemi 2024) restrict the number of settings due to the exponential growth in action space, while studies based on continuous actions (Zhang, Qiu, and Tan 2024) omit categorical actions. For discrete-action spaces, we address the limitations by constraining the action space to the dataset distribution and employing a factored action space. Despite these optimizations, mapping discretized actions back to continuous induces distribution shift and can produce unsafe policies, an issue largely unaddressed in prior work. We introduce hybrid actions for two of the SOTA offline RL algorithms (IQL and EDAC), enabling them to address these issues while capturing the full range of MV settings. (Chen et al. 2022) also uses a hybrid action space for optimizing MV settings. However, they adapt an off-policy RL algorithm that lacks the safety regularization of offline methods, potentially leading to unsafe policies due to overestimation. Instead of modifying the Soft Actor-Critic

(SAC) for hybrid actions using the Gumbel-Softmax reparameterization trick to allow gradient flow, we compute the expectation of the discrete component directly, as the exact distribution is available. This considerably reduces variance in policy updates (Christodoulou 2019).

Most prior studies mentioned above that focus on optimizing MV have primarily relied on mortality-based rewards, either sparse or shaped. However, medical studies (see Section 4) indicate that mortality alone is not a reliable indicator of MV treatment quality (Beigmohammadi et al. 2022). Instead, we adopt VFDs as a reward, reflecting that patients who spend less time on MV and avoid mortality received better care. Additionally, we add rewards to maintain MV-related physiological vitals within safe ranges and prevent complications.

A similar setup is used by (Kondrup et al. 2023), where the Apache-II score function is employed as an intermediate reward, and augmented by a terminal reward based on mortality. However, combining an intermediate reward with a terminal reward is non-trivial. Depending on how the terminal reward is scaled, offline RL can end up ignoring one objective and overemphasize the other. Our setup optimally balances the different medical objectives.

3 Preliminaries

Offline RL

The offline RL problem is formulated as an MDP where $\{\mathcal{S}, \mathcal{A}, P, R, \gamma\}$ represents the state space, action space, transition distribution, reward distribution, and discount factor, respectively. The initial state is sampled as $s_0 \sim d_0$.

An RL agent is trained by optimizing a policy $\pi(a|s) : \mathcal{S} \rightarrow \Delta(\mathcal{A})$, guided by Q-values defined as $Q^\pi(s, a) = \mathbb{E}[\sum_{t=0}^{\infty} \gamma^t R(\cdot|s_t, a_t)]$. Value-based RL algorithms use the Bellman error (Sutton and Barto 2018) to update the Q-function. Offline RL learns from a fixed dataset \mathcal{D} collected by a behavior (clinician) policy $\pi_b(a|s)$ and can suffer from overestimation due to out-of-distribution (OOD) actions. Therefore, a regularization term is often added to the standard Bellman error to mitigate this overestimation, e.g.:

Conservative Q-Learning (CQL). The CQL loss function regularizes Q-values on unseen state-action pairs alongside the standard Bellman objective (Kumar et al. 2020).

Implicit Q-Learning (IQL). The IQL objective (Kostrikov, Nair, and Levine 2021) learns high-performing actions by computing action advantages via the expectile of the state value function, thereby updating the policy without querying Q-values for unseen actions.

Ensemble Diversified Actor Critic (EDAC). The EDAC (An et al. 2021) uses an ensemble of critics (Q^π) to estimate the uncertainty of a given state-action pair, effectively lower bounding Q-values for uncertain pairs to prevent overestimation. Additionally, it incorporates a diversity loss among ensemble members to promote varied Q-value estimates, improving uncertainty estimation.

MV Setting	Unit	Setpoint domain
Ventilation Control Mode	-	$\{VCV, PCV\}$
Respiratory rate (RR)	min^{-1}	$[5, 60]$
Tidal Volume (V_T)	ml/kg	$[3, 12]$
Driving Pressure (ΔP)	cmH ₂ O	$[0, 26]$
PEEP	cmH ₂ O	$[0, 20]$
FiO ₂	%	$[21, 100]$

Table 1: Controllable MV Settings, their units and possible setpoints

MDP Formulation

State. The state space consist of 26 clinician-defined variables (see supplementary material C). These include vital signs, respiratory parameters, laboratory values, fluid balance measures, and demographic characteristics.

Actions. We consider six MV settings (control variables) listed in Table 1. These settings include both continuous and discrete variables. For the discrete-actions setup, all the continuous settings are discretized using clinician-defined bins. Thresholds were selected based on clinical guidelines, prior studies, and observed practice ranges, balancing clinical resolution with device limits and safety considerations (see supplementary Material D) and aligning with the literature e.g. (Grasselli et al. 2023). The hybrid actions setup used the settings in their native continuous and categorical forms.

We define a *setpoint* as the value selected for a ventilator setting. An *action* is the tuple of setpoints across all mechanical-ventilation (MV) settings, $a = (a_1, \dots, a_6)$, and the action space is the set of all such tuples, $\mathcal{A} = \mathcal{A}_1 \times \dots \times \mathcal{A}_6$. In our discrete setup with clinician-defined bins, $|\mathcal{A}_{\text{disc}}| = 18,144$.

Rewards. Environment rewards are calculated as defined in Section 4.

4 Clinically-Guided Reward Design (C1)

Prior work (e.g., (Kondrup et al. 2023)) often uses mortality as the endpoint for evaluating MV treatment. However, mortality following a treatment is a highly confounded outcome and can be influenced by various factors, including the underlying disease and comorbidities. Instead, we worked with clinicians to define two main reward objectives aimed at improving survival, reducing the risk of VILI, and lowering healthcare costs.

Primary Objective. The primary objective is to reduce the duration of MV while keeping the patient alive. Prolonged MV increases the risk of complications such as ventilator-induced lung injury, infection (Wu et al. 2019; Slutsky and Ranieri 2013), hypotension, and diaphragm dysfunction due to disuse atrophy (Peñuelas et al. 2019). These complications can hinder successful weaning and are associated with increased mortality. Moreover, effective MV strategies may provide significant clinical benefit even in the absence of mortality reduction, if they facilitate earlier liberation from the ventilator (Schoenfeld, Bernard, and ARDS

Network 2002). Ventilator-free days within the first month after the start of MV are often used to assess the quality of MV. This measure combines mortality within the first month with the duration of MV, and is directly linked to the quality of ventilator settings (Schoenfeld, Bernard, and ARDS Network 2002).

Secondary Objective. The secondary objective is to limit physiological impairments due to MV. Oxygenation levels (e.g., SpO₂, PaO₂) and vital signs (e.g., blood pH, mean arterial pressure (MAP), heart rate) must remain within safe ranges to prevent adverse outcomes. For example PaO₂ demonstrated a U-shaped association with mortality (Boyle et al. 2021), while dangerously in- or decreased blood pH-values are closely linked to organ failure and increased mortality (Kraut and Madias 2001). In collaboration with physicians, we have identified blood pH, mean arterial pressure, PaO₂, SaO₂, PaCO₂, heart rate and peripheral oxygen as key physiological parameters and their safe ranges to guide decision-making. These parameters were selected by clinicians from multiple hospitals within the IntelliLung consortium through a Delphi-like consensus process, ensuring broad expert agreement rather than ad-hoc tuning. Details on the parameters, their safe ranges, and reward weights are listed in supplementary E.

Reward Design

The total reward at each step is the sum of range reward r_{range} , time penalty r_{tp} : $r = r_{\text{range}} + r_{\text{tp}} + r_{\text{vfd}}$, and VFD reward r_{vfd} :

Range Reward r_{range} guides the agent toward learning the secondary objective. It is calculated as follows:

$$r_{\text{range}} = \frac{\sum_{i=1}^N w_i \cdot \mathbf{1}_{[a_i, b_i]}(p_i)}{\sum_{i=1}^N w_i}, \quad r_{\text{range}} \in [0, 1] \quad (1)$$

where N is the total number of physiological parameters, p_i is the value of the i -th parameter, w_i is its assigned weight and $\mathbf{1}_{[a_i, b_i]}(p_i)$ is the indicator function that activates when p_i is within its defined safe range $[a_i, b_i]$.

Time Penalty r_{tp} encourages the policy to prefer actions that shorten time on MV and to not prolong it by simply collecting positive r_{range} step rewards. We set the time penalty to the negative of the maximum per-step range reward, i.e. $r_{\text{tp}} = -1$.

Ventilator Free Days (VFDs) (Schoenfeld, Bernard, and ARDS Network 2002) are commonly used in clinical trials to evaluate MV interventions. We can include VFDs in the reward function to check how well MV met the primary objective during an episode. VFDs are calculated as follows:

$$VFD = \text{alive} \times \max(0, \min(\Delta t_{\text{re}}, T_{\text{max}}) - \Delta t_{\text{mv}}) \quad (2)$$

where Δt_{re} is the re-intubation time, T_{max} the fixed observation window (typically around 30 days) and Δt_{mv} the days on the ventilator.

Effectively, VFDs are the number of days a patient is both alive and free from MV within T_{max} . If the patient dies, then

VFD = 0. If the patient survives, the VFDs equal the number of days spent off the ventilator, $T_{\max} - \Delta t_{mv}$. However, if the patient is re-intubated before the end of the observation period T_{\max} , the VFDs are reduced to $\Delta t_{re} - \Delta t_{mv}$ (Schoenfeld, Bernard, and ARDS Network 2002). The final VFD reward r_{vfd} is:

$$r_{vfd} = w_{vfd} \cdot \frac{VFD}{T_{\max}} \quad (3)$$

where w_{vfd} is a hyper-parameter that controls the contribution of r_{vfd} .

r_{vfd} can be applied for each episode in two ways: (i) we can apply r_{vfd} at the terminal time step and 0 otherwise (*VFD@TerminalStep*) or (ii) at each time step (*VFD@EachStep*). Each option interacts differently with the weight w_{vfd} and changes the relative contribution of the intermediate range reward r_{range} . We analyze the effect of w_{vfd} and the two options in Section 8.

5 Discrete-Action Optimizations (C2)

Using clinician-defined bins for each MV setting (see supplementary material D), results in $|\mathcal{A}_{disc}| = 18,144$ distinct actions. However, this large action space introduces challenges, including increased computational complexity and Q-value overestimation for rarely observed actions. To address these issues, we introduce the following optimizations:

Restrict Action Space

The action space is restricted to only the clinician-observed actions present in the dataset. Beyond efficiency, this constraint eliminates unsafe actions, as they do not appear in the dataset because clinicians avoid them in practice. For example, setting both V_T and FiO_2 too low could cause severe hypoxia, leading to organ damage or worse. RL algorithms cannot estimate the effect (Q-values) of actions absent from the dataset, and may overestimate and still select them at inference despite offline-RL regularizations. By removing these unseen actions from the policy’s action space, we completely avoid the risk of choosing them. This reduces the action space to 53.6% of $|\mathcal{A}_{disc}|$.

Based on clinical insights, we found further potential to optimize the action space. In volume-controlled ventilation (VCV) the control variable is V_T and ΔP is masked. In pressure-controlled ventilation (PCV) the control variable is ΔP and V_T is masked. We implemented this by assigning a dedicated (null) bin for the masked setting. This reduces the action space to just 6.9% of $|\mathcal{A}_{disc}|$. However, this was not used in our experiments to maintain comparability with hybrid action algorithms.

Linear Decomposition of the Q-function

Even after reducing the discrete-action space, the critic must evaluate Q-values over a large number of actions. In offline settings with limited coverage, many state-action pairs are rare or unseen, which increases estimator variance and can introduce extrapolation bias. When the action is factored as $a = (a_1, \dots, a_K)$, we model the critic additively,

$Q(s, a) \approx \sum_{k=1}^K q_k(s, a_k)$, following prior work on factored action spaces (Tang et al. 2023). This reduces the output dimensionality from $O(\prod_k |A_k|)$ to $O(\sum_k |A_k|)$, enabling faster training and better sample efficiency because each gradient step updates multiple per-factor values. An additive critic introduces bias by omitting cross-factor interactions. However, in low-coverage regimes, the reduction in estimator variance often outweighs this bias, yielding a more favorable bias–variance trade-off. In practice, this leads to improved policies, as shown in Section 8. Section H and J of the supplementary materials provide an illustration and the code implementation.

6 Hybrid Action Space for Offline RL (C3)

The default implementations of IQL and EDAC operate in continuous action space. We modified their CORL implementations (Tarasov et al. 2022) to support hybrid actions as follows:

IQL. For IQL, the critic function stays the same except that both continuous and one-hot encoded discrete-actions are input into the network. IQL uses Advantage-Weighted Regression (AWR) (Peng et al. 2019) for policy optimization. The adapted $\log \pi_\phi(a|s)$ for AWR is calculated as $\log \pi_\phi(a|s) = \log \pi_\phi^d(a^d|s) + \log \pi_\phi^c(a^c|s)$ and $(a^c, a^d) \sim \mathcal{D}$.

EDAC. EDAC is a combination of SAC (Haarnoja et al. 2018) with ensemble of critics and a diversity loss. We follow the approach described in (Delalleau et al. 2019) for SAC adaptation, but adapted the critic to accept both discrete and continuous actions as inputs rather than outputting Q-values for each discrete-action combination. We opted for this design because empirical results showed reduced critic-loss variance and more stable training. It is important to note that the diversity loss requires $\Delta_a Q(s, a)$, which in the hybrid case becomes $\Delta_{(a^c, a^d)} Q(s, a^c, a^d)$.

7 Experimental Setup

Study Datasets

For our experiments, we use data from three publicly available clinical databases on PhysioNet (Goldberger et al. 2000): MIMIC IV (Johnson et al. 2020), eICU (Pollard et al. 2018) and HiRID (Faltys et al. 2021). The datasets include patients from different hospitals across Europe and US, ensuring broad representativeness of patient characteristics and treatment regimes. The experiment cohort includes patients aged 18 years or older who underwent at least 4 hours of MV in the ICU. Identical pre-processing steps were applied to each database. These steps included data cleaning, filtering, episode construction with 1 hour time windows, and computation & imputation. Detailed descriptions of each step are provided in supplementary materials A.

The resulting datasets for each database (see supplementary materials B) were then combined, forming a final dataset containing 12,572 patients and 1,252,505 hours of MV. The dataset was split into 80% training and 20% testing. Stratified splitting was performed based on episode length and mortality while ensuring that no patient appeared in both

the training and test sets. To ensure comparability, the data splits remained unchanged across all experiments.

RL Training

Discrete-Actions. We compare four CQL variants and IQL for the discrete-action setup:

1. **CQL:** Unmodified CQL as in prior work, e.g., (Kondrup et al. 2023).
2. **C-CQL:** Constrained-Action Space CQL.
3. **F-CQL:** Factored Critic CQL.
4. **CF-CQL:** Constrained-Action Space with a Factored Critic CQL(proposed discrete variant).
5. **DiscreteIQL:** IQL adapted for discrete-action space.

All setups use an MLP with four hidden layers of 256 units and a discount factor $\gamma = 0.99$. We set the CQL regularization to $\alpha = 0.1$ and the IQL parameters to $\tau = 0.8$ and $\beta = 5$. Gradients are clipped by global norm to 1.0, and we use a learning rate of 10^{-6} for CQL and $5e^{-5}$ for IQL. Each model is trained for 100,000 gradient steps, with the target critic updated after every step via Polyak averaging (coefficient $\tau_{polyak} = 0.005$).

Hybrid Actions. We evaluate two algorithms for the hybrid-action setup:

1. **HybridIQL:** IQL adapted for hybrid actions. It uses learning rates of $lr = 1e^{-4}$, inverse temperature $\beta = 5$, expectile $\tau = 0.8$, and a MLP with four hidden layers of 256 units.
2. **HybridEDAC:** EDAC modified for hybrid actions. We also apply automatic entropy adjustment (Haarnoja et al. 2019) with a target entropy of $\mathcal{H}_c = -0.3$ for continuous actions and $\mathcal{H}_d = 0.3$ for discrete-actions, $lr = 3e^{-5}$, diversity term $\eta = 1.0$ and an ensemble of 25 critics.

Evaluation

Fitted Q-Evaluation. FQE (Le, Voloshin, and Yue 2019) is an Off-Policy Evaluation (OPE) method used to estimate the performance of a policy π using previously collected dataset \mathcal{D} . The FQE method fits a Q^π using $y = r + \gamma Q(s', \pi(s'))$ as target. The policy performance metric V^π is defined as estimated returns of policy π on the initial state distribution:

$$V^\pi = \mathbb{E}_{s_0 \sim d_0} [Q^\pi(s_0, \pi(s_0))] \quad (4)$$

Since traditional FQE only captures expected returns, the distributional FQE (DistFQE) is implemented following the Quantile Regression DQN (QR-DQN) approach (Dabney et al. 2018). The performance of the behavior (or clinicians) policy, V^{π_b} , is evaluated by replacing the target with $y = r + \gamma Q(s', a')$, where a' is drawn from \mathcal{D} , and then applying Eq. (4).

Policy Coverage. To quantify policy coverage with respect to the behavior policy at states in \mathcal{D} , we first estimate the behavior policy π_b . We do so by fitting a parametric model $p_{\hat{\theta}}(a|s)$ using negative log likelihood (NLL) loss:

$$\theta^* = \arg \min_{\hat{\theta}} \sum_{(s,a) \in \mathcal{D}} -\log p_{\hat{\theta}}(a | s)$$

We then score π by the expected log-density that the fitted behavior model assigns to actions drawn from π at dataset states:

$$d^\pi = \mathbb{E}_{s \sim \hat{d}_{\mathcal{D}}, a \sim \pi(\cdot|s)} [\log p_{\theta^*}(a | s)], \quad (5)$$

where $\hat{d}_{\mathcal{D}}$ denotes the empirical marginal over states in \mathcal{D} . Smaller values of d^π indicate that π tends to select OOD actions, potentially leading to poor performance and overestimation of V^π .

Reward Evaluation Metrics. To evaluate how well different reward functions align with the primary and secondary objectives, we compute two metrics. The first is the Spearman correlation between the learned Q-values and the episode VFD , denoted $Corr@VFD$. The second is the Spearman correlation between the learned Q-values and the episode-mean r_{range} , denoted $Corr@RangeReward$. The Q-values are estimated from Q^π learned using FQE.

Tchebycheff Scalarization for Reward Selection. We frame the problem of selecting a reward function as a multi-objective optimization task over objectives: $Corr@VFD$ & $Corr@RangeReward$. We apply a range-normalized Tchebycheff scalarization (Wierzbicki 1980) to convert the objective vector into a single scalar score. For each objective, we compute ideal and nadir values to normalize the scales, then evaluate each candidate’s normalized distance from the ideal point. The Tchebycheff value is defined as the maximum of these distances. We select the reward function with the smallest Tchebycheff value, yielding a Pareto-efficient compromise.

8 Results & Discussion

Reward Functions Analysis

A high $Corr@VFD$ suggests that episodes in which the patient survives and spends fewer days on MV receive higher mean Q-values for their state–action pairs. Similarly, a high $Corr@RangeReward$ indicates that episodes where the patient vital signs stay within safe limits are associated with higher mean Q-values.

Fig. 1 plots correlation metrics for both $VFD@EachStep$ and $VFD@TerminalStep$ across different w_{vfd} values. In order to compare with mortality-based baselines (Kondrup et al. 2023; Eghbali, Alhanai, and Ghassemi 2024), we include $r_{morta} \in [-w_{morta}, w_{morta}]$ and set r_{range} as the intermediate reward to directly benchmark both objectives.

Fig. 1a shows that the time penalty r_{tp} already aligns Q-values with the primary objective, as $Corr@VFD$ is already relatively high at $w_{vfd} = 0$. Adding VFD strengthens this, as $Corr@VFD$ rises from 0.48 to 0.56 for $VFD@EachStep$ at $w_{vfd} = 0.5$, and to 0.66 for $VFD@TerminalStep$ at $w_{vfd} = 100$. In contrast, alignment with the secondary objective ($Corr@RangeReward$) is highest at small w_{vfd} ; as w_{vfd} increases, intermediate rewards are effectively ignored and the correlation drops (Fig. 1b).

Mortality-based rewards, such as those used in the prior work mentioned above, lack the time penalty and achieve a high $Corr@VFD$ only at $w_{morta} = 100$; at that point




Figure 1: Panels (a)–(b) report the correlations (higher is better) $Corr@VFD$ and $Corr@RangeReward$, respectively, across reward scales (w_{vfd} for VFD and w_{morta} for Mortality). Panel (c) shows the Tchebycheff value (lower is better) over the same scales. For every (reward, scale) pair, we trained five independent HybridIQL policies and report the mean of their final values.

$Corr@RangeReward = 0.13$, meaning range rewards are largely ignored.

Optimal Reward Function. Maximizing only $Corr@VFD$ or only $Corr@RangeReward$ can drive the algorithm to excel on one objective at the expense of the other. Instead, we seek a Pareto-balanced compromise. By casting reward selection as a multi-objective problem and applying range-normalized Tchebycheff scalarization. We find that the $VFD@EachStep$ reward with $w_{vfd} \in [0.25, 1]$ and in particular $w_{vfd} = 0.5$ provides the best overall trade-off (see Fig. 1c). If the reward definition or its weighting parameters change, we rerun the same multi-objective selection to produce an updated Pareto-optimal reward function.

Impact of Action Discretization (C4)

Prior work (e.g. (Kondrup et al. 2023; Eghbali, Alhanai, and Ghassemi 2024)) discretizes continuous MV settings into bins to enable training with discrete-action offline RL algorithms. At deployment, however, the selected bin must be mapped back to a real-valued setting. This bin-to-value *reconstruction* can induce distribution shift, with potential performance and safety implications.

We evaluate several reconstruction functions, each mapping a chosen bin to real valued MV setting: (i) the bin mode (estimated from the dataset), (ii) a normal distribution centered at the bin mode, (iii) the bin mean, and (iv) a uniform distribution over the bin interval. We quantify the effect on policy coverage d^π using a policy trained with CF-CQL. As shown in Table 2, uniform sampling yields the lowest coverage (largest distribution mismatch), whereas the bin-mode attains the highest coverage (smallest mismatch). To ensure comparability with hybrid-action methods that operate directly in the original action space, we therefore use bin-mode reconstruction for discrete-action policies when estimating d^π .

Policy Analysis

In this section, we evaluate the performance and reliability of discrete and hybrid-action offline RL algorithms using FQE estimates (V^π) and policy coverage (d^π).

Reconstruction Method	d^π
Bin Mode	-0.62
Gaussian Sampling	-0.70
Bin Mean	-0.89
Uniform Sampling	-1.26

Table 2: Policy coverage d^π of a CF-CQL policy with different reconstruction methods.

Policy	$V^\pi(\pm\sigma)$	$d^\pi(\pm\sigma)$
Clinician policy	-9.51	-0.59
HybridIQL	-5.92 ± 0.12	-0.43 ± 0.01
DiscreteIQL	-6.88 ± 0.07	-0.40 ± 0.01
HybridEDAC	-3.03 ± 1.08	-1.80 ± 0.07
CQL	-5.30 ± 0.24	-1.84 ± 0.04
CQL($\alpha = 0.5$)	-8.11 ± 0.58	-1.05 ± 0.19
C-CQL	-5.26 ± 0.17	-1.81 ± 0.09
F-CQL	-6.99 ± 0.19	-0.52 ± 0.01
CF-CQL	-6.82 ± 0.38	-0.55 ± 0.01

Table 3: Initial-state return (V^π) and policy coverage (d^π) for policies trained using discrete- and hybrid-action algorithms, reported relative to the clinician policy. Results are reported as $\mu \pm \sigma$ over five runs with random initializations.

Discrete-Action Optimizations. In Table 3 we compare the unmodified CQL from previous work, e.g. (Kondrup et al. 2023) with our proposed variant CF-CQL. While unmodified CQL achieves the highest estimated value (V^π), it suffers from very low policy coverage (d^π), indicating over-reliance on OOD actions. This is confirmed by Fig. 2, where most CQL-selected actions lie beyond the OOD threshold. In contrast, CF-CQL maintains most actions within the in-distribution (ID) region and achieves a more reliable value estimate, albeit a bit lower. To disentangle which component drives the gains, we also test variants with only constrained-action (C-CQL) and only factored-critic (F-CQL). C-CQL provides marginal improvement over CQL, while F-CQL performs similarly to CF-CQL, confirming that the improvement in coverage (d^π) stems from the factored critic,

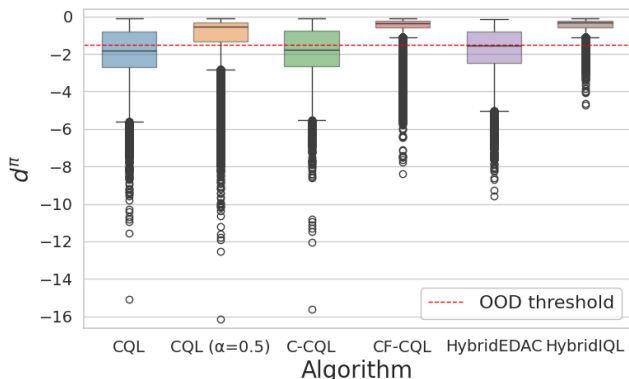


Figure 2: Distribution of policy coverage d^π (across states in test set) for each algorithm. The red dashed line represents the OOD threshold, defined as lower Tukey fence (at least 75% samples lie above it) of d^π distribution under the clinician policy. We classify actions with d^π below threshold as OOD.

Policy	$V^\pi(\pm\sigma)$	$d^\pi(\pm\sigma)$
Clinician policy	-12.50	-0.74
HybridIQL	-6.40 ± 0.18	-0.55 ± 0.01
HybridEDAC	-2.56 ± 1.23	-1.65 ± 0.10
CQL	-6.85 ± 0.88	-1.46 ± 0.10
CF-CQL	-7.57 ± 0.21	-0.80 ± 0.05

Table 4: Initial-state return (V^π) and policy coverage (d^π), trained on eICU and HiRID and evaluated on MIMIC-IV for external generalizability, reported as $\mu \pm \sigma$ over five random initializations.

whereas constraining actions mainly serves to eliminate unsafe actions.

Higher Critic Regularization. We also experimented with a higher CQL critic regularization ($\alpha = 0.5$) to increase policy coverage. Increasing α leads to higher d^π , but at the cost of significantly reduced V^π , while d^π still remained below that of CF-CQL. This suggests that regularization alone is insufficient to match both the performance and distributional robustness achieved through factored critics.

Hybrid-Action Algorithms. Among hybrid-action methods, HybridEDAC achieves the highest V^π , but like CQL, its low d^π suggests heavy reliance on OOD actions, raising concerns about potential overestimation. In contrast, HybridIQL strikes a more favorable trade-off: it achieves competitive value estimates while maintaining significantly higher policy coverage across both discrete and hybrid methods, highlighting the advantage of the hybrid-action setup.

Are hybrid-actions needed? To assess whether the gains of HybridIQL are due to the hybrid action formulation or IQL itself, we also compare it with DiscreteIQL in Table 3. HybridIQL achieves higher V^π while maintaining the most favorable trade-off between performance and distribution

Policy	$V^\pi(\pm\sigma)$	$d^\pi(\pm\sigma)$
Clinician policy	-9.51	-0.59
HybridIQL	-6.19 ± 1.23	-0.38 ± 0.04
HybridEDAC	-7.14 ± 3.63	-1.61 ± 0.54
CQL	-8.16 ± 1.15	-0.91 ± 0.37
CF-CQL	-9.11 ± 1.99	-0.62 ± 0.30

Table 5: Initial-state return (V^π) and policy coverage (d^π), reported as $\mu \pm \sigma$ over a hyperparameter sweep (learning rate, parameter count, offline RL regularization) and two random initializations. We report results from 162 runs for each of HybridIQL and HybridEDAC, and 54 runs for each of CQL and CF-CQL.

coverage d^π . As discussed previously, the discrete-action setups are also affected by distribution shifts at inference time due to bin-to-value reconstruction, which may pose additional safety concerns.

Generalization. To test cross-dataset generalization, we train only on eICU and HiRID and test on MIMIC-IV. Results are given in Table 4 and closely follow Table 3. HybridIQL shows the best generalization by achieving highest performance V^π while having policy coverage d^π higher than the clinician policy.

Robustness to hyper-parameters. Table 5 reports a sweep over learning rates $\{10^{-4}, 10^{-5}, 10^{-6}\}$, network capacities $\{[64, 64], [128, 128, 128], [256, 256, 256]\}$, offline RL regularization coefficients for CQL variants $\alpha \in \{0.1, 0.5, 2.5\}$, for HybridIQL $\tau \in \{0.7, 0.8, 0.9\}$ and $\beta \in \{5, 10, 25\}$, for HybridEDAC $\eta \in \{0.1, 1, 10\}$ and total number of critics $\in \{5, 10, 25\}$, as well as two random seeds. HybridIQL is the most stable, having high coverage and strong initial-state value V^π across the sweep, followed by CF-CQL. By contrast, HybridEDAC and CQL exhibit higher variance and consistently poorer coverage.

Action Distribution. Fig. 3 presents the policy distribution of FiO₂ in the hybrid setting, and Fig. 4 shows the discrete setting. Clinician baseline puts little probability mass above 80%. However, the low d^π policies (HybridEDAC, CQL, C-CQL) overestimate in these low-density regions (high variance). In contrast, CQL with a higher α (0.5) is overly conservative, placing most mass where the clinician density is highest and leading to a lower V^π (high bias). Algorithms with higher d^π (CF-CQL, HybridIQL) better balance the bias–variance trade-off and achieve competitive performance. A complete set of action-distributions plot is provided in the supplementary material G.

Key Takeaways. Focusing solely on high performance (V^π) is insufficient for safe deployment when a policy frequently selects OOD actions (i.e., has low d^π). Prioritizing algorithms that minimize distributional shift yields more trustworthy performance and increases the likelihood of real-world adoption. We find that our discrete-action optimizations remove unsafe actions and improve distributional coverage, enabling CQL to scale to six actions




Figure 3: Action distribution of FiO_2 for hybrid-action setup vs clinician.




Figure 4: Action distribution of FiO_2 for discrete-action setup vs clinician.

(vs. two–three in prior work). In parallel, the hybrid-action method, specifically HybridIQL, avoids inference-time distribution shift inherent to discretization and delivers higher performance, better coverage, stronger cross-dataset generalization, and greater hyperparameter robustness than discrete baselines. Taken together, these results indicate that enforcing coverage and leveraging hybrid actions provide a principled path toward AI-DSS that reduce ventilator-induced lung injuries and enable safer mechanical ventilation.

Limitations & Future Research

While offline RL methods inherently avoid OOD actions, the agent may still propose unsafe actions if such actions are present in the dataset. In future work, we aim to address this by collaborating with clinicians to identify and avoid potentially unsafe action patterns by incorporating national and international ventilation guidelines (Grasselli et al. 2023; Qadir et al. 2024). Furthermore, our study relies on publicly available ICU datasets that suffer from noise, limited resolution, and require assumptions about patient states and interventions. Although our datasets span multiple US and European hospitals, exceeding prior single-center or US-only studies, we acknowledge no dataset captures all ICU scenarios. To mitigate this, we are currently collaborating with partner hospitals in Germany, Spain, Poland, Italy and the US to collect high-quality, prospectively curated data. This system is designed for recommendation only, as full agency and responsibility reside with the attending physician. Still, the potential of recommendations to bias clinical decision-making warrants further consideration. Lastly, all

evaluations were retrospective and offline. We used FQE, coverage, and OOD thresholding to estimate performance and provide empirical safety guarantees supporting initial retrospective claims. However, prospective validation of the observed benefits is required. Ongoing observational ICU trials are assessing real-world performance and safety without altering care, with a randomized controlled trial planned to evaluate efficacy and safety against standard care in accordance with regulatory best practices.

9 Conclusion

This research tackles the critical challenge of suboptimal MV settings, a persistent issue in ICUs that contributes to preventable harm and mortality worldwide. By advancing offline RL methods in this context, our work supports the development of safer, more reliable AI-DSS for MV. Specifically, we addressed key technical challenges by incorporating constrained and factored action space optimization for discretized setup to increase MV settings recommended to clinicians. Our work also identified that discretizing actions can introduce distribution shifts, for which we studied the impact of different reconstruction methods. In parallel, we extended SOTA offline RL algorithms (IQL & EDAC) to incorporate hybrid actions to overcome the limitations of discretization. CF-CQL and HybridIQL demonstrated improvements over clinician policies while offering trade-offs between enhanced safety and performance. Additionally, in cooperation with clinical experts, we introduced rewards based on VFD and safe ranges, which are better aligned with clinical objectives. Overall, the methods presented in this study contribute to a more robust adaptation of offline RL in MV by paving the way for a prospective multi-hospital study in Germany, Spain and Poland scheduled to commence later this year. While our focus was on MV, the developed methodologies may also have broader applicability beyond this specific clinical task.

Ethical Statement

We used publicly available (MIMIC IV, eICU and HiRID), de-identified data in compliance with relevant usage agreements and applicable data protection laws. This retrospective analysis did not involve any direct patient interventions or modifications in patient care. We recognize that our reward function is specified at the cohort level and may not fully capture the needs, risks, or preferences of specific individuals or minority subgroups, which could lead to systematic harm. To mitigate this, we constrain the learned policy to remain close to clinician-like behavior, ensure that any deployment would occur only under active clinician oversight, and plan explicit subgroup- and individual-level safety analyses in future prospective studies.

Acknowledgments

Our work is supported by the European Union project IntelliLung with Grant Agreement No. 101057434. The authors gratefully acknowledge the support of the IntelliLung Consortium¹.

¹<https://intellilung-project.eu/consortium-members>

References

- An, G.; Moon, S.; Kim, J.-H.; and Song, H. O. 2021. Uncertainty-based offline reinforcement learning with diversified q-ensemble. *Advances in neural information processing systems*, 34: 7436–7447.
- Beigmohammadi, M. T.; Amoozadeh, L.; Rezaei Motlagh, F.; Rahimi, M.; Maghsoudloo, M.; Jafarnejad, B.; Eslami, B.; Salehi, M. R.; and Zendejdel, K. 2022. Mortality predictive value of APACHE II and SOFA scores in COVID-19 patients in the intensive care unit. *Canadian Respiratory Journal*, 2022(1): 5129314.
- Bellani, G.; Laffey, J. G.; Pham, T.; Fan, E.; Brochard, L.; Esteban, A.; Gattinoni, L.; Van Haren, F.; Larsson, A.; McAuley, D. F.; et al. 2016. Epidemiology, patterns of care, and mortality for patients with acute respiratory distress syndrome in intensive care units in 50 countries. *Jama*, 315(8): 788–800.
- Boyle, A. J.; Holmes, D. N.; Hackett, J.; Gilliland, S.; McCloskey, M.; O’Kane, C. M.; Young, P.; Di Gangi, S.; and McAuley, D. F. 2021. Hyperoxaemia and hypoxaemia are associated with harm in patients with ARDS. *BMC Pulmonary Medicine*, 21: 1–12.
- Chen, S.; Qiu, X.; Tan, X.; Fang, Z.; and Jin, Y. 2022. A model-based hybrid soft actor-critic deep reinforcement learning algorithm for optimal ventilator settings. *Information Sciences*, 611: 47–64.
- Christodoulou, P. 2019. Soft Actor-Critic for Discrete Action Settings. arXiv:1910.07207.
- Dabney, W.; Rowland, M.; Bellemare, M.; and Munos, R. 2018. Distributional reinforcement learning with quantile regression. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.
- Delalleau, O.; Peter, M.; Alonso, E.; and Logut, A. 2019. Discrete and Continuous Action Representation for Practical RL in Video Games. arXiv:1912.11077.
- Eghbali, N.; Alhanai, T.; and Ghassemi, M. M. 2024. Distribution-Free Uncertainty Quantification in Mechanical Ventilation Treatment: A Conformal Deep Q-Learning Framework. arXiv:2412.12597.
- Faltys, M.; Zimmermann, M.; Lyu, X.; Hüser, M.; Hyland, S.; Rätsch, G.; and Merz, T. 2021. HiRID, a high time-resolution ICU dataset (version 1.1. 1). *Physio. Net*, 10.
- Fichtner, F.; Mörer, O.; Laudi, S.; Weber-Carstens, S.; and Kaisers, U. 2017. Invasive Beatmung und Einsatz extrakorporaler Verfahren bei akuter respiratorischer Insuffizienz. *DIVI*, 4: 154–163.
- Goldberger, A. L.; Amaral, L. A.; Glass, L.; Hausdorff, J. M.; Ivanov, P. C.; Mark, R. G.; Mietus, J. E.; Moody, G. B.; Peng, C.-K.; and Stanley, H. E. 2000. PhysioBank, PhysioToolkit, and PhysioNet: components of a new research resource for complex physiologic signals. *circulation*, 101(23): e215–e220.
- Grasselli, G.; Calfee, C. S.; Camporota, L.; Poole, D.; Amato, M. B. P.; Antonelli, M.; Arabi, Y. M.; Baroncelli, F.; Beitler, J. R.; Bellani, G.; Bellingan, G.; Blackwood, B.; Bos, L. D. J.; Brochard, L.; Brodie, D.; Burns, K. E. A.; Combes, A.; D’Arrigo, S.; De Backer, D.; Demoule, A.; Einav, S.; Fan, E.; Ferguson, N. D.; Frat, J.-P.; Gattinoni, L.; Guérin, C.; Herridge, M. S.; Hodgson, C.; Hough, C. L.; Jaber, S.; Juffermans, N. P.; Karagiannidis, C.; Kesecioglu, J.; Kwizera, A.; Laffey, J. G.; Mancebo, J.; Matthay, M. A.; McAuley, D. F.; Mercat, A.; Meyer, N. J.; Moss, M.; Munsch, L.; Myatra, S. N.; Ng Gong, M.; Papazian, L.; Patel, B. K.; Pellegrini, M.; Perner, A.; Pesenti, A.; Piquilloud, L.; Qiu, H.; Ranieri, M. V.; Riviello, E.; Slutsky, A. S.; Stapleton, R. D.; Summers, C.; Thompson, T. B.; Valente Barbas, C. S.; Villar, J.; Ware, L. B.; Weiss, B.; Zampieri, F. G.; Azoulay, E.; and Cecconi, M. 2023. ESICM guidelines on acute respiratory distress syndrome: definition, phenotyping and respiratory support strategies. *Intensive Care Medicine*, 49(7): 727–759.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. arXiv:1801.01290.
- Haarnoja, T.; Zhou, A.; Hartikainen, K.; Tucker, G.; Ha, S.; Tan, J.; Kumar, V.; Zhu, H.; Gupta, A.; Abbeel, P.; and Levine, S. 2019. Soft Actor-Critic Algorithms and Applications. arXiv:1812.05905.
- Johnson, A.; Bulgarelli, L.; Pollard, T.; Horng, S.; Celi, L. A.; and Mark, R. 2020. MIMIC-IV. *PhysioNet*. Available online at: <https://physionet.org/content/mimiciv/>, 49–55.
- Kondrup, F.; Jiralerspong, T.; Lau, E.; de Lara, N.; Shkrob, J.; Tran, M. D.; Precup, D.; and Basu, S. 2023. Towards safe mechanical ventilation treatment using deep offline reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 15696–15702.
- Kostrikov, I.; Nair, A.; and Levine, S. 2021. Offline reinforcement learning with implicit q-learning. *arXiv preprint arXiv:2110.06169*.
- Kraut, J. A.; and Madias, N. E. 2001. Approach to patients with acid-base disorders. *Respiratory care*, 46(4): 392–403.
- Kumar, A.; Zhou, A.; Tucker, G.; and Levine, S. 2020. Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 33: 1179–1191.
- Le, H.; Voloshin, C.; and Yue, Y. 2019. Batch policy learning under constraints. In *International Conference on Machine Learning*, 3703–3712. PMLR.
- Lee, H.; Yoon, H.-K.; Kim, J.; Park, J. S.; Koo, C.-H.; Won, D.; and Lee, H.-C. 2023. Development and validation of a reinforcement learning model for ventilation control during emergence from general anesthesia. *npj Digital Medicine*, 6(1): 145.
- Liu, S.; Xu, Q.; Xu, Z.; Liu, Z.; Sun, X.; Xie, G.; Feng, M.; and See, K. C. 2024. Reinforcement Learning to Optimize Ventilator Settings for Patients on Invasive Mechanical Ventilation: Retrospective Study. *Journal of Medical Internet Research*, 26: e44494.
- Peine, A.; Hallawa, A.; Bickenbach, J.; Dartmann, G.; Fazlic, L. B.; Schmeink, A.; Ascheid, G.; Thiemermann, C.; Schuppert, A.; Kindle, R.; et al. 2021. Development and

- validation of a reinforcement learning algorithm to dynamically optimize mechanical ventilation in critical care. *NPJ digital medicine*, 4(1): 32.
- Peng, X. B.; Kumar, A.; Zhang, G.; and Levine, S. 2019. Advantage-Weighted Regression: Simple and Scalable Off-Policy Reinforcement Learning. arXiv:1910.00177.
- Peñuelas, O.; Keough, E.; López-Rodríguez, L.; Carriedo, D.; Gonçalves, G.; Barreiro, E.; and Lorente, J. Á. 2019. Ventilator-induced diaphragm dysfunction: translational mechanisms lead to therapeutic alternatives in the critically ill. *Intensive care medicine experimental*, 7(Suppl 1): 48.
- Pollard, T. J.; Johnson, A. E.; Raffa, J. D.; Celi, L. A.; Mark, R. G.; and Badawi, O. 2018. The eICU Collaborative Research Database, a freely available multi-center database for critical care research. *Scientific data*, 5(1): 1–13.
- Prasad, N.; Cheng, L.-F.; Chivers, C.; Draugelis, M.; and Engelhardt, B. E. 2017. A reinforcement learning approach to weaning of mechanical ventilation in intensive care units. *arXiv preprint arXiv:1704.06300*.
- Qadir, N.; Sahetya, S.; Munshi, L.; Summers, C.; Abrams, D.; Beitler, J.; Bellani, G.; Brower, R. G.; Burry, L.; Chen, J.-T.; Hodgson, C.; Hough, C. L.; Lamontagne, F.; Law, A.; Papazian, L.; Pham, T.; Rubin, E.; Siuba, M.; Telias, I.; Patolia, S.; Chaudhuri, D.; Walkey, A.; Rochweg, B.; and Fan, E. 2024. An Update on Management of Adult Patients with Acute Respiratory Distress Syndrome: An Official American Thoracic Society Clinical Practice Guideline. *American Journal of Respiratory and Critical Care Medicine*, 209(1): 24–36.
- Roedl, K.; Amann, D.; Eichler, L.; Fuhrmann, V.; Kluge, S.; and Müller, J. 2019. The chronic ICU patient: Is intensive care worthwhile for patients with very prolonged ICU-stay (≥ 90 days)? *European journal of internal medicine*, 69: 71–76.
- Roggeveen, L. F.; Hassouni, A. e.; de Grooth, H.-J.; Girbes, A. R.; Hoogendoorn, M.; Elbers, P. W.; and Collaborators, D. I. D. S. A. C.-. 2024. Reinforcement learning for intensive care medicine: actionable clinical insights from novel approaches to reward shaping and off-policy model evaluation. *Intensive Care Medicine Experimental*, 12(1): 32.
- Schoenfeld, D. A.; Bernard, G. R.; and ARDS Network. 2002. Statistical evaluation of ventilator-free days as an efficacy measure in clinical trials of treatments for acute respiratory distress syndrome. *Crit. Care Med.*, 30(8): 1772–1777.
- Silva, P. L.; Negrini, D.; and Rocco, P. R. M. 2015. Mechanisms of ventilator-induced lung injury in healthy lungs. *Best Practice & Research Clinical Anaesthesiology*, 29(3): 301–313.
- Slutsky, A. S. 1999. Lung injury caused by mechanical ventilation. *Chest*, 116: 9S–15S.
- Slutsky, A. S.; and Ranieri, V. M. 2013. Ventilator-Induced Lung Injury. *New England Journal of Medicine*, 369(22): 2126–2136.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press, 2nd edition.
- Tang, S.; Makar, M.; Sjoding, M. W.; Doshi-Velez, F.; and Wiens, J. 2023. Leveraging Factored Action Spaces for Efficient Offline Reinforcement Learning in Healthcare. arXiv:2305.01738.
- Tarasov, D.; Nikulin, A.; Akimov, D.; Kurenkov, V.; and Kolesnikov, S. 2022. CORL: Research-oriented Deep Offline Reinforcement Learning Library. In *3rd Offline RL Workshop: Offline RL as a "Launchpad"*.
- Torrini, F.; Gendreau, S.; Morel, J.; Carteaux, G.; Thille, A. W.; Antonelli, M.; and Mekontso Dessap, A. 2021. Prediction of extubation outcome in critically ill patients: a systematic review and meta-analysis. *Critical Care*, 25: 1–12.
- Wierzbicki, A. P. 1980. The Use of Reference Objectives in Multiobjective Optimization. In Fandel, G.; and Gal, T., eds., *Multiple Criteria Decision Making Theory and Application*, volume 177 of *Lecture Notes in Economics and Mathematical Systems*, 468–486. Springer.
- Wu, D.; Wu, C.; Zhang, S.; and Zhong, Y. 2019. Risk factors of ventilator-associated pneumonia in critically ill patients. *Frontiers in pharmacology*, 10: 482.
- Yu, C.; Ren, G.; and Dong, Y. 2020. Supervised-actor-critic reinforcement learning for intelligent mechanical ventilation and sedative dosing in intensive care units. *BMC medical informatics and decision making*, 20: 1–8.
- Zhang, B.; Qiu, X.; and Tan, X. 2024. Balancing therapeutic effect and safety in ventilator parameter recommendation: An offline reinforcement learning approach. *Eng. Appl. Artif. Intell.*, 131: 107784.