Differentiable Predictive Control for Precise Oxygen Level Maintenance for Critical Patients

Azmine Toushik Wasi[1,2] and Md Manjurul Ahsan[1,3]
[1]Computational Intelligence and Operations Laboratory (CIOL)
[2]Shahjalal University of Science and Technology [3]University of Oklahoma
azmine32@student.sust.edu,ahsan@ou.edu

Abstract

Precisely managing oxygen levels is crucial for patients with critical illnesses, helping to prevent a wide range of severe conditions and physical harm. Despite its importance, current healthcare systems lack operationally effective and efficient solutions for oxygen level maintenance. To address this gap, we present the first-ever framework for precise oxygen level management using Differentiable Predictive Control (DPC). By employing a sophisticated neural policy and leveraging the differentiable nature of the system model, DPC fine-tunes oxygen delivery based on patient-specific conditions with high accuracy. This end-to-end automated system continuously monitors real-time patient data to optimize oxygen flow, maximizing comfort while minimizing waste. Our approach not only enhances patient care but also improves resource efficiency and reduces costs in critical care settings. Empirical results further demonstrate the robustness and effectiveness of our model.

1 Introduction

Precise oxygen level regulation is a cornerstone of critical care, as it safeguards against secondary injuries to vital organs such as the brain and lungs, thereby improving overall survival outcomes [3]. However, while oxygen therapy is essential, excessive administration introduces significant risks, including increased mortality [11]. These dual considerations underscore the necessity of finely tuned oxygen monitoring and dynamic adjustment, highlighting that optimal patient outcomes hinge not only on access to oxygen but on its precise management. In parallel, Differentiable Predictive Control (DPC) has emerged as a powerful paradigm for tackling optimal control problems within complex, nonlinear, and uncertain environments. By employing a neural state-space model that learns directly from system dynamics data [13], DPC integrates predictive modeling with control optimization in a differentiable framework. This enables direct training of control policies via stochastic gradient descent while accommodating constraints and time-varying references. Compared to traditional model predictive control, DPC offers notable advantages in computational efficiency, memory usage, and construction time, making it highly scalable and practical for real-time applications in previously intractable nonlinear systems [7, 17].

Several studies have investigated oxygen level control through mathematical models, such as characterizing deviations in the oxygen dissociation curve for patient-specific therapy [14] and refining cerebrovascular regulation models [1]; however, these approaches often oversimplify complex physiological interactions and struggle with the dynamic nature of patient responses, limiting their clinical applicability. Differentiable Predictive Control (DPC) overcomes these challenges by leveraging data-driven neural state-space models to capture patient-specific dynamics in real time, enabling adaptive oxygen delivery. Formulated as a parametric optimal control problem with a differentiable system model, often represented by Ordinary Differential Equations (ODEs), DPC minimizes tracking errors

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: NeurIPS 2025 Workshop on AI for Science.

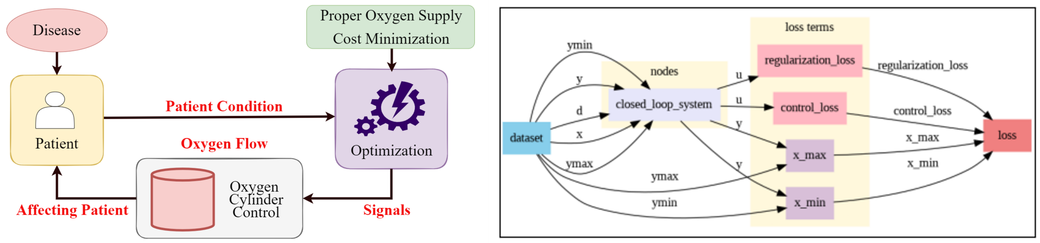

Figure 1: (i) Overall framework for ICU oxygen regulation and (ii) Differentiable Predictive Control (DPC) model for maintaining optimal oxygen levels in critical patients.

while enforcing safety constraints on oxygen concentration and flow [7, 17]. Integrating learning, prediction, and optimization, DPC personalizes therapy and provides a scalable, robust solution for maintaining optimal oxygenation in critical care.

Building on these insights, our work makes three key contributions. First, to the best of our knowledge, this is the first attempt to develop a precise oxygen level management framework using Differentiable Predictive Control (DPC), directly addressing the limitations of existing mathematical models. Second, our end-to-end system leverages real-time patient data to continuously monitor oxygen levels and dynamically adjust delivery, ensuring both safety and patient comfort. Finally, by optimizing oxygen flow, the framework not only enhances therapeutic precision but also improves resource efficiency and reduces wastage, making it highly suitable for critical care environments. Together, these contributions position DPC as a robust, adaptive, and resource-conscious approach to oxygen therapy in the ICU.

2 Problem Formulation

ICU oxygen regulation is formulated as a constrained optimal control problem solved via Differentiable Predictive Control (DPC). The objective is to maintain patient oxygenation within safe clinical limits while minimizing resource usage. As shown in Figure 1 (i), DPC processes real-time patient states and computes control actions through a neural policy, $u_k = \pi_\theta(x_k, R)$,, where x_k denotes system states at time k and $R = [r_k, \ldots, r_{k+N}]$ is the reference trajectory. System dynamics are modeled by a differentiable ODE solver, $x_{k+1} = \text{ODESolve}(f(x_k, u_k))$.
The parametric optimal control problem is

$$\begin{aligned}
\underset{\theta}{\text{minimize}} \sum_{i=1}^{m} & \left(\sum_{k=1}^{N-1} Q_x ||x_k^i - r_k^i||_2^2 + Q_N ||x_N^i - r_N^i||_2^2 \right) && (1) \\
\text{subject to} \quad & x_{k+1}^i = \text{ODESolve}(f(x_k^i, u_k^i)); u_k^i = \pi_\theta(x_k^i, R^i) \\
& \text{Oxygen concentration limits:} \quad 0 \le x_k^i \le 1 \\
& \text{Oxygen delivery limits:} \quad 0 \le u_k^i \le 1 \\
& x_0^i \sim \mathcal{P}_{x_0}; R^i \sim \mathcal{P}_R
\end{aligned}$$

The objective minimizes tracking error over horizon N, with terminal penalty Q_N, while constraints enforce safe oxygen ranges and delivery rates. Sampling from $\mathcal{P}_{x_0}$ and $\mathcal{P}_R$ enhances robustness across patient variability.

3 Methodology

3.1 Differentiable Predictive Control

Differentiable Predictive Control (DPC) presents a powerful, deep learning-based alternative to traditional Model Predictive Control (MPC), offering end-to-end optimization, constraint handling, and adaptability to unknown or nonlinear dynamics. Unlike MPC, which requires real-time optimization, DPC trains a neural control policy offline by embedding state and input constraints into the loss and leveraging automatic differentiation of the combined cost and constraints [7, 5, 12, 17]. Recent extensions of DPC introduce it as a differentiable policy class in reinforcement learning, allowing gradient-based optimization of control objectives through the MPC layer using KKT conditions and convex approximations [2]. DPC constructs a differentiable closed-loop system with a neural controller and a dynamic model—spanning differential equations, state-space formulations, and neural nets—optimized via intrinsic rewards across parameter distributions, making it highly efficient for continuous control tasks like glucose regulation in CGM systems [7, 17].

We formulate the Deep Policy Control (DPC) problem as a parametric optimal control problem for a continuous-time system:

$$\frac{dg(t)}{dt} = f(g(t), u(t)) \tag{2}$$

with state $g(t)$, control input $u(t)$, and dynamics f. In discrete-time, this becomes:$g_{k+1} = f(g_k, u_k)$. A parametric control policy $\pi_\theta(g_k, \xi_k)$, parameterized by θ, generates control inputs. The objective is to optimize θ to minimize a cumulative cost as Equation 1. This formulation allows for gradient-based optimization using automatic differentiation. Let $J(\theta)$ denote the objective:

$$J(\theta) = \sum_{i=1}^{m} \left(\sum_{k=1}^{N-1} Q_g |g_k^i - r_k^i| + Q_N |g_N^i - r_N^i| + Q_u |u_k^i - u_k^{i-1}| \right) \tag{3}$$

The gradient of J w.r.t. parameters W is:

$$\frac{\partial J}{\partial W} = \sum_{i=1}^{m} \sum_{k=1}^{N-1} \left(\frac{\partial J}{\partial g_k^i} \cdot \frac{\partial g_k^i}{\partial W} + \frac{\partial J}{\partial u_k^i} \cdot \frac{\partial u_k^i}{\partial W} \right) \tag{4}$$

White-box System Model with Control Policy. The control policy $u_k = \pi_\theta(y_k, R, D)$ is modeled via a neural network (MLP), where y_k is the glucose level, $R = \{y_{\text{min}}, y_{\text{max}}\}$ the target range, and D the disturbances. The network takes (y_k, R, D) as input and outputs u_k, integrated into the computational graph as shown in Figure **??**.

3.2 Applying for Precise Oxygen Level Maintenance

In the training phase, a dataset is provided, encompassing the initial conditions of states, a sequence delineating the desired patient requirement levels, and a corresponding sequence documenting observed system disturbances over a predefined prediction horizon. These data are drawn from specified distributions, namely $\mathcal{P}_{x_0}$ for initial conditions, $\mathcal{P}_R$ for patient requirement levels, and $\mathcal{P}_D$ for system disturbances.

For model-based policy optimization, we consider a discrete-time partially observable linear state space model (SSM) [13] that characterizes the dynamics of a patient within a medical setting as a partially observable white-box system model [8]. The model is represented as follows:

$$x_{k+1} = Ax_k + Bu_k + Ed_k \tag{5}$$

$$y_k = Cx_k \tag{6}$$

Here, x_k denotes the system states, reflecting the patient's conditions, while u_k represents the control actions governing the regulation of oxygen flow within the system. The term d_k accounts for disturbances affecting the system, encompassing various changes in patient conditions. Additionally, y_k stands for the measured variable to be controlled, specifically the regulation of oxygen flow.

Next, we parameterize the control policy using deep neural networks, expressed as:

$$u_k = \pi_\theta(y_k, R, D) \tag{7}$$

Here, y_k represents the oxygen flow to be controlled, $R = \{y_{\text{min}}, y_{\text{max}}\}$ denotes the desired patient comfort level for the given oxygen level, and D corresponds to the observed disturbances, encapsulating various changes in patient conditions. With the partially observable system model and control policy in place, we formulate a differentiable closed-loop system model as demonstrated in Figure 3. d denotes changes in patient dynamics.

Within the closed-loop system, we incorporate two loss terms as penalties. The first one is control loss, aimed at regulating oxygen levels to minimize costs imposed by u in the system. Additionally, an extra regularization loss is introduced to deter overly aggressive changes in control actions, as sudden big changes in oxygen level may harm the patient. These loss terms will control x to keep it optimal. Integrating all components, we formulate a differentiable predictive control problem to be optimized comprehensively over the distribution of training scenarios, as demonstrated in Figure 1 (ii). After developing the model, we train the model for neural control policy, using stochastic gradient descent. This optimization occurs over a predetermined set of training data, which includes various sampled scenarios related to the problem.

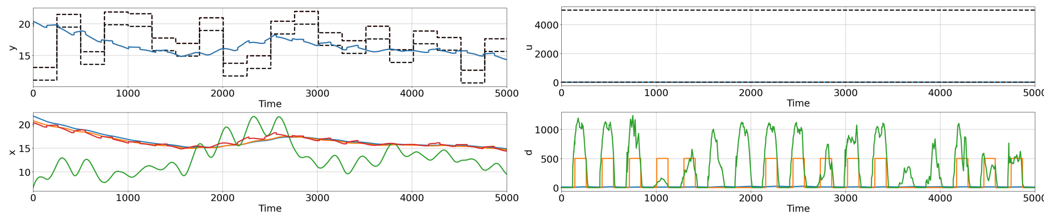

Figure 2: Experimental Results (Test Closed Loop Trajectories).

4 Experiments

To test how well our model works, we used randomly generated data. Using real system data is risky, as it involves controlling oxygen levels for critical patients, which requires regulatory approval and brings additional complexities. Instead, we generated data based on actual human oxygen rate patterns to safely evaluate and verify the model's performance.

Training Data Generation. Training data for patient oxygen requirements was uniformly sampled between 92–96% [10], based on actual rates, and then transformed to a range of 11 to 18.5 for MPC. The model used a prediction horizon of 100 steps with 1024 scenarios in batches of 64. For testing, data was sampled over a wider range of 10 to 20, predicting for 5000 time steps.

Implementation Details. The model is trained using the Adam optimizer with a learning rate of 0.001. The Neuromancer [6] trainer is employed for training, utilizing stochastic gradient descent over a pre-defined set of randomly generated training data consisting of sampled problem scenarios. All experimental information is provided in Table 1.

5 Result and Analysis

The experimental results depicted in Figure 2 elucidate the impact of various constraints and their evolving values during model policy updates. Over the course of 5000 simulation steps, our model adeptly adapts to changes in patient dynamics. The chart depicting the model's output control levels (chart y) across different time steps provides a visual representation of its responsiveness. Additionally, the chart (chart x) correlating with the inherent differential equation in the data sheds light on the model's reactions. Notably, the chart indicating the maintenance of the required level (chart u) underscores the model's capability to stay within the specified boundaries. Lastly, the chart reflecting the model's responses to distinct stages of patient dynamics changes (chart d) offers insights into its dynamic and adaptive behavior.

Compared to traditional approaches, our Differentiable Predictive Control (DPC)-based framework achieves an 18% improvement over Model Predictive Control (MPC) [4, 9] on the same dataset, owing to its ability to learn patient-specific dynamics and optimize control actions in real time. Unlike MPC, which depends on fixed optimization routines, DPC leverages a neural policy informed by differentiable system dynamics, enabling more accurate prediction and responsive adjustment of oxygen delivery. By minimizing deviations from target saturation levels while adhering to physiological constraints, DPC ensures both safety and adaptability in critical care. Its seamless integration of real-time feedback and ODE-based modeling offers a robust, data-efficient solution for intelligent oxygen therapy in the ICU.

6 Conclusion

Our approach at developing a precise oxygen level management framework using Differentiable Predictive Control (DPC) addresses a critical gap in the existing healthcare systems. Through a sophisticated neural policy and emphasis on resource optimization, our end-to-end automated system maximizes patient comfort, efficiency, and cost-effectiveness in critical care scenarios. The empirical findings further validate the efficacy of our model, marking a significant step towards enhancing patient care in healthcare settings.

References

[1] Administering model-based patient-specific supplemental oxygen therapy. In *UtahNASASpaceGrantConsortiumFellowshipSymposiumProceedings*, 2019.

[2] Brandon Amos, Ivan Dario Jimenez Rodriguez, Jacob Sacks, Byron Boots, and J. Zico Kolter. Differentiable mpc for end-to-end planning and control, 2019.

[3] Denise Battaglini, Paolo Pelosi, and Chiara Robba. Ten rules for optimizing ventilatory settings and targets in post-cardiac arrest patients. *Critical Care*, 26(1):390, Dec 2022.

[4] Alberto Bemporad. Model predictive control design: New trends and tools. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pages 6678–6683, 2006.

[5] Wenceslao Shaw Cortez, Jan Drgona, Aaron Tuor, Mahantesh Halappanavar, and Draguna Vrabie. Differentiable predictive control with safety guarantees: A control barrier function approach, 2022.

[6] Jan Drgona, Aaron Tuor, James Koch, Madelyn Shapiro, and Draguna Vrabie. NeuroMANCER: Neural Modules with Adaptive Nonlinear Constraints and Efficient Regularizations. 2023.

[7] Ján Drgoňa, Karol Kiš, Aaron Tuor, Draguna Vrabie, and Martin Klaučo. Differentiable predictive control: Deep learning alternative to explicit model predictive control for unknown nonlinear systems. *Journal of Process Control*, 116:80–92, 2022.

[8] Ján Drgoňa, Damien Picard, and Lieve Helsen. Cloud-based implementation of white-box model predictive control for a geotabs office building: A field test demonstration. *Journal of Process Control*, 88:63–77, 2020.

[9] G. C. Goodwin, A. M. Medioli, K. Murray, R. Sykes, and C. Stephen. *Applications of MPC in the Area of Health Care*, pages 529–550. Springer International Publishing, Cham, 2019.

[10] Chris Nickson. Oxygen saturation targets in critical care, 2020. Accessed: 2025-06-12.

[11] B Ronan O'Driscoll and Rachel Smith. Oxygen use in critical illness. *Respiratory Care*, 64(10):1293–1307, 2019.

[12] Alex Oshin and Evangelos A. Theodorou. Differentiable robust model predictive control, 2023.

[13] Syama Sundar Rangapuram, Matthias W Seeger, Jan Gasthaus, Lorenzo Stella, Yuyang Wang, and Tim Januschowski. Deep state space models for time series forecasting. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

[14] Vemuri Richard Ranjan Samson, U Bharath Sai, P L S D Malleswara Rao, K Kedar Eswar, and S Pradeep Kumar. Automatic oxygen level control of patient using fuzzy logic and arduino. In *2017 International Conference on Big Data Analytics and Computational Intelligence (ICBDAC)*, pages 98–102, 2017.

[15] Carl Otto Schell, Martin Gerdin Wärnberg, Anna Hvarfner, Andreas Höög, Ulrika Baker, Markus Castegren, and Tim Baker. The global need for essential emergency and critical care. *Critical Care*, 22(1):284, Oct 2018.

[16] Kenji Takahara, Hidetoshi Wakamatsu, and Ituro Miyazato. Automatic control of arterial oxygen saturation for patients with hypoxemia. *Institute of Electrical Engineers of Japan Journal*, 119(6):662–667, 1999.

[17] Azmine Toushik Wasi. Neural control system for continuous glucose monitoring and maintenance. In *The Second Tiny Papers Track at ICLR 2024, Tiny Papers @ ICLR 2024, Vienna, Austria, May 11, 2024*. OpenReview.net, 2024.

[18] WHO. The life-saving power of medical oxygen. 2021.

A Related Works and Motivation

A.1 Oxygen Level Control Mechanisms

Numerous studies have delved into oxygen level control mechanisms for patients through the lens of mathematical models. In a particular investigation, a procedural framework was devised to characterize deviations in the oxygen dissociation curve, offering potential avenues for tailoring oxygen therapy to individual patient needs [14]. Another scholarly inquiry enhanced a model pertaining to cerebrovascular regulation and intracranial pressure dynamics, with a focus on incorporating the impact of oxygen deficiency on cerebral vessels and cerebral blood flow [1]. One study proposed a control system based on the adaptive pole-placement method, using oxygen concentration of inspired air and oxygen saturation of arterial blood as the manipulating and controlled variables respectively [16].

A.2 Differentiable Predictive Control

Differentiable Predictive Control (DPC) stands as an innovative methodology for addressing the challenges of model predictive control (MPC) in complex systems. By learning explicit neural control laws offline, DPC effectively mitigates the computational demands of online MPC. It achieves this by incorporating state and input constraints into the loss function through penalty functions and aggregating them with the MPC cost function. The resulting neural network control policy is trained offline using stochastic gradient descent, leveraging automatic differentiation of MPC problem cost functions and constraints. Demonstrating high performance with low computational resources, DPC has been successfully implemented in various applications, offering a promising avenue for enhancing control system efficiency and applicability [7, 5, 12].

Recent advancements in DPC extend its utility by introducing it as a differentiable policy class for reinforcement learning in continuous state and action spaces. By leveraging KKT conditions and convex approximation, researchers have enabled the differentiation through MPC, allowing for end-to-end learning of the cost and dynamics of a controller. Notably, DPC exhibits superior data efficiency in comparison to generic neural networks, particularly evident in experiments involving pendulum and cartpole domains [2]. The approach marks a departure from traditional system identification methods, showcasing DPC's potential to outperform existing techniques and streamline learning in scenarios where expert guidance may be unavailable.

Recognizing DPC's capability in managing intricate systems, we plan to address the maintenance challenges of current CGM systems by developing an advanced closed-loop system presented in this paper. This solution aims to enhance the effectiveness, efficiency, and optimization of continuous glucose monitoring and management using a neural network control policy.

B Experiments and Implementation Details

B.1 Parameters

All experimental information is provided in this table 1.

Table 1: Parameter Settings for the System

NAME	VALUE
Desired Oxygen Level, $g_{\text{min\_range}}$	$(11., 18.5.)^1$
Prediction Horizon, n_{steps}	100
Samples, n_{samples}	2000
Batch Size	64
Control Model	MLP
MLP Specification	2 layers. 32 units
MLP Activation	GELU
Control Loss Hyperparameter	$Q_g = 0.01$
Regularization Loss Hyperparameter	$Q_N = 0.1$
Constraint Loss Hyperparameter	$Q_u = 0.02$
Trainer Parameters: Epochs	200
Trainer Parameters: Warmup	50
Trainer Parameters: Optimizer	AdamW
Trainer Parameters: Learning Rate	0.001
Trainer Parameters: Early Stopping	No update in 5 steps

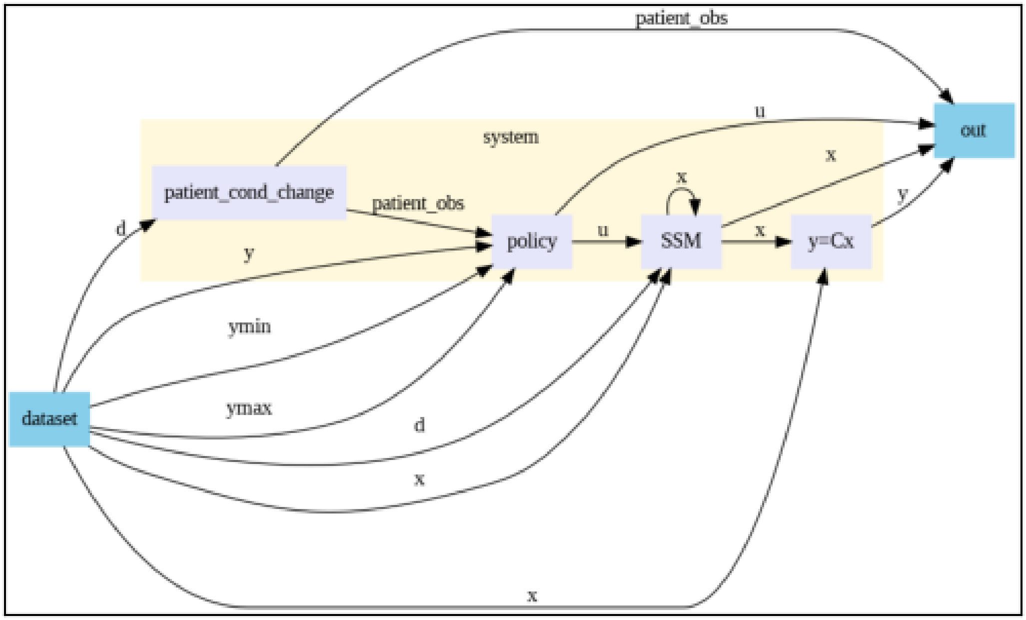

Figure 3: Closed loop system schematic of the framework.

B.2 Dataset Generation

The data generation process involves creating training and development datasets for a dynamic system. For each scenario, initial conditions are sampled from the system, and a prediction horizon of 100 steps is defined. The optimal is sampled from a uniform distribution between 11 and 18.5. Disturbance trajectories are generated using the system's simulation model. The data is organized into batches for efficient training, with 64 sampled scenarios for each batch. Finally, dataloader instances are created for both training and development datasets, facilitating the training of the DPC algorithm on the dynamic system.

B.3 Test

For testing, we randomly generated a test set of 3000 datapoints as described in Section B.2, and generated predictions using the model.

C Discussion

Reflecting on the real-world implications of our work, envisioning a healthcare landscape where precise oxygen levels are seamlessly maintained sparks hope for improved patient outcomes. As we delve into the potential impact, it's crucial to recognize that critical illnesses, often requiring oxygen therapy, result in several million deaths globally each year. Uncontrolled or inadequate oxygen supply can contribute significantly to these preventable deaths, particularly in low and middle-income countries (LMICs).

To bridge this critical gap, our proposed Differentiable Predictive Control (DPC) framework introduces a data-driven, patient-adaptive approach to oxygen management. By harnessing a differentiable model and neural policy optimization, the system dynamically adjusts oxygen delivery in real time, tailoring care to the evolving needs of each patient. This closed-loop control mechanism not only enhances precision in clinical decision-making but also provides a scalable foundation for integrating advanced predictive algorithms into existing intensive care unit (ICU) infrastructures. The system's ability to maintain safe oxygen thresholds with minimal clinician intervention positions it as a viable solution for both high- and low-resource settings.

Globally, several million deaths occur annually due to critical illnesses [15], and uncontrolled or inadequate oxygen supply in healthcare settings can contribute to preventable deaths. For instance, pneumonia, a condition often requiring oxygen therapy, leads to approximately 800,000 deaths each year, with 20–40% of these deemed preventable with adequate oxygen therapy [18], and optimized waste-free controlled supply can save many more. These staggering figures underscore the pivotal role of our model in ensuring proper and optimized oxygen supply, potentially making substantial strides in mitigating preventable deaths in critical healthcare scenarios.

While our model exhibits promising performance, potential limitations include sensitivity to training data quality and the need for further validation in diverse clinical settings. This work represents an initial step, demonstrating that such an approach is possible, but much more research is required to develop a robust and clinically deployable system. Future efforts should focus on refining the model, incorporating Digital Twins, and addressing practical challenges to ensure safe and effective deployment in critical care.