LEARNING RESPONSE-AWARE PATIENT DYNAMICS FOR RESPIRATORY SUPPORT

Anonymous authors
Paper under double-blind review

ABSTRACT

Respiratory support can shape the short-term physiological trajectory of critically ill patients, but patients receiving the same intervention may follow different physiological trajectories. Clinical patient dynamics models typically predict future states from recent physiology and recorded interventions, while physiological change is mainly represented through the predicted future state. We propose a response-aware patient dynamics model that explicitly represents physiological change during autoregressive state updating. The model decomposes predicted physiological change into state-dependent baseline dynamics and respiratory-support-associated deviations, with room air providing a reference for the decomposition. We provide a formal analysis of this reference-anchored formulation. A response pathway encodes the predicted physiological change and uses it to update the latent patient state across the forecast horizon. Across ICU cohorts from two independent institutions, the proposed model achieves comparable overall trajectory prediction to patient dynamics baselines, with more consistent improvements when physiological states are changing.

1 INTRODUCTION

Acute respiratory failure is a common reason for hospitalization and admission to the intensive care unit (Villgran et al., 2022). Patients with hypoxemia are often first treated with conventional oxygen therapy. If respiratory failure worsens, support may be escalated to high-flow nasal cannula (HFNC), non-invasive ventilation (NIV), or invasive mechanical ventilation (IMV) (Azevedo et al., 2015; Coudroy et al., 2022). Physiological response to respiratory support varies across patients. Some patients improve after treatment, while others remain unstable or continue to deteriorate and may require further escalation. Changes in physiological measurements over short time intervals can provide useful information about the patient's evolving condition during respiratory support (Perez et al., 2025). Modeling these short-term changes may therefore help characterize how physiology evolves under respiratory support. Figure 1 illustrates this sequential setting and the role of physiological response in describing the transition from one patient state to the next.

Recent patient dynamics and clinical world models predict how patient states evolve over time under recorded interventions (Xu et al., 2026; Wu et al., 2026). A common formulation represents the current patient state from recent clinical measurements and models the next state conditioned on the current intervention, $p(s_{t+1} \mid s_t, a_t)$. The intervention may be represented by a treatment category, dose, or other recorded treatment variables (Yang et al., 2025; Qazi et al., 2025). In these models, physiological progression and intervention-associated change are typically represented together within the same state-transition function.

However, this formulation has two limitations for modeling respiratory support. First, physiological change after treatment may reflect both the patient's underlying progression and variation associated with the current respiratory support, but these components are usually modeled together. Second, action-conditioned dynamics models typically represent the next physiological level, without explicitly representing the change from the previous state during state updating. This change may provide additional information about how the patient's condition is evolving. Clinical studies of HFNC and NIV support this motivation. Patients who later succeed or fail treatment may have similar baseline measurements but show different physiological changes during the first hours of respiratory sup-

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

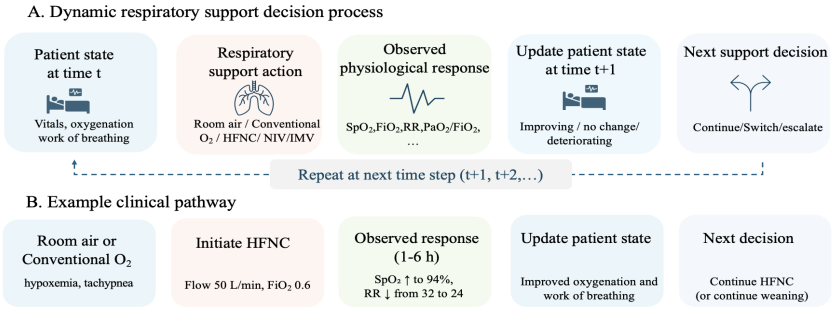

Figure 1: Clinical motivation for response-aware respiratory dynamics. (A) Respiratory support is delivered sequentially as patient physiology evolves. The respiratory support action space includes room air as a no-active-support category, since patients may experience periods without active respiratory support between support modalities in clinical practice. Following a respiratory support action, the observed physiological response provides new information that can update the patient state and inform the next support decision. (B) Example clinical pathway showing how short-term response after HFNC initiation may help determine whether to continue the current support or consider escalation.

port (Mongodi et al., 2026), and the predictive value of response-based indices can increase after treatment initiation (Roca et al., 2019).

In this study, we extend action-conditioned patient dynamics modeling to short-term respiratory-support forecasting with a response-aware formulation. The model separates state-dependent baseline dynamics from respiratory-support-associated deviations and explicitly represents predicted physiological change during autoregressive state updating.

Our main contributions are as follows:

- We formulate short-term respiratory-support forecasting as an action-conditioned patient dynamics problem in which physiological change is explicitly represented during trajectory prediction.
- We develop a response-aware dynamics model that decomposes predicted physiological change into state-dependent baseline dynamics and respiratory-support-associated deviations, and uses the predicted change to update the latent patient state. We further analyze the decomposition using room air as a reference condition.
- We evaluate the model on real-world ICU data from two institutions, including temporal and external validation, and analyze performance across forecast horizons, respiratory-support states, and clinically relevant physiological variables.

2 Related Work

2.1 Clinical Patient Dynamics and World Models

Clinical trajectory models learn how patient states evolve over time from longitudinal EHR data. Clinical world models further condition these dynamics on treatment actions to model future patient trajectories (Xu et al., 2026; Wu et al., 2026; Yang et al., 2025; Qazi et al., 2025). MedDreamer models patient dynamics in a latent state space using posterior state inference from observed measurements and an action-conditioned prior for rollout (Xu et al., 2026). SepsisAgent uses an action-conditioned world model to predict short-term physiological states under different treatment actions (Wu et al., 2026). These models differ in their state representations and transition architectures, but generally model physiological progression and intervention-associated change within the same transition process. Our work instead decomposes predicted physiological change into state-dependent baseline dynamics and respiratory-support-associated deviations, and explicitly represents predicted physiological change during state updating.

2.2 Physiological Response to Respiratory Support

Physiological response during respiratory support can provide information about a patient's evolving respiratory status. Following HFNC or NIV initiation, changes in oxygenation, respiratory rate, and related measurements have been used to assess subsequent treatment success or failure. For example, the predictive value of the ROX index increases during the hours following HFNC initiation (Roca et al., 2019), and other studies have reported different early physiological response patterns between patients who subsequently succeed or fail respiratory support (Mongodi et al., 2026; Perez et al., 2025). These findings suggest that physiological change itself carries information about the patient's evolving state, beyond the physiological level at a single time point, which provides clinical motivation for explicitly representing physiological change in a patient dynamics model. Our model therefore encodes predicted physiological change and uses it during autoregressive state updating.

3 Methodology

3.1 Problem formulation

We denote the patient's recent physiological history over a look-back window of length L as $\boldsymbol{X}_{t-L+1:t}$, where $\boldsymbol{x}_t \in \mathbb{R}^p$ contains the physiological measurements at time t. At each time point, the respiratory intervention is represented by $\boldsymbol{a}_t \in \mathbb{R}^q$, including the respiratory-support modality and available modality-specific settings, such as oxygen flow and FiO_2 for oxygen support and airway pressures for non-invasive ventilation. Static patient characteristics are denoted by $\boldsymbol{c}$.

Given the observed patient history and the recorded respiratory-support sequence, our goal is to predict the physiological trajectory over the next H hours:

$$\hat{\boldsymbol{X}}_{t+1:t+H} = f_\theta\left(\boldsymbol{X}_{t-L+1:t}, \boldsymbol{A}_{t-L+1:t+H-1}, \boldsymbol{c}\right), \tag{1}$$

where the model uses the observed respiratory-support history to initialize the patient state and is conditioned on the recorded respiratory-support action at each prediction step. Measurement-availability indicators and time-since-last-measurement variables are also included in the initial state representation.

We model the future trajectory autoregressively. At each prediction step, the predicted physiological change is decomposed into a state-dependent baseline component and a respiratory-support-associated deviation. The predicted change is then used both to obtain the next physiological state and to update the latent patient state before the following prediction step.

3.2 Response-aware physiological dynamics

We represent the patient using a latent state that is updated over the prediction horizon. The model has four main components: (1) initialization of the latent patient state from recent history, (2) state-dependent encoding of respiratory support, (3) decomposition of predicted physiological change into baseline dynamics and respiratory-support-associated deviations, and (4) response-aware latent-state updating using the predicted physiological change. An overview of the model architecture is provided in Appendix A.

Patient state initialization. A sequence encoder summarizes the recent patient trajectory using physiological measurements, time-since-last-measurement information, measurement-availability indicators, and respiratory-support history. Static patient characteristics and comorbidities are encoded separately.

The resulting dynamic and static representations are combined to initialize the latent patient state:

$$\boldsymbol{s}_t = E_{\text{state}}\left(\boldsymbol{X}_{t-L+1:t}, \boldsymbol{A}_{t-L+1:t}, \boldsymbol{T}_{t-L+1:t}, \boldsymbol{M}_{t-L+1:t}, \boldsymbol{c}\right), \tag{2}$$

where $\boldsymbol{T}$ denotes time-since-last-measurement information and $\boldsymbol{M}$ denotes measurement-availability indicators. The initialized state $\boldsymbol{s}_t$ represents the patient's recent physiological trajectory, respiratory-support history, and static characteristics at the start of prediction.

State-dependent respiratory-support representation. Given the current latent patient state and respiratory intervention, we construct a state-dependent respiratory-support representation:

$$\boldsymbol{z}_t^a = E_{\text{action}}\left(\boldsymbol{s}_t, \boldsymbol{a}_t\right). \tag{3}$$

Conditioning the respiratory-support representation on the current latent state allows the same intervention to have different representations for different patient states.

Decomposition of physiological change. The predicted physiological change is decomposed into state-dependent baseline dynamics and a respiratory-support-associated deviation:

$$\Delta\hat{\boldsymbol{x}}_t = F_{\text{base}}(\boldsymbol{s}_t) + \boldsymbol{g}_t \odot F_{\text{treat}}(\boldsymbol{z}_t^a), \tag{4}$$

where

$$\boldsymbol{g}_t = \sigma\left(G_{\text{gate}}\left(\boldsymbol{s}_t, \boldsymbol{a}_t\right)\right) \tag{5}$$

is a feature-wise gate.

The baseline term $F_{\text{base}}(\boldsymbol{s}_t)$ represents the component of predicted physiological change determined by the current latent patient state. The effective respiratory-support-associated deviation is $\boldsymbol{g}_t \odot F_{\text{treat}}(\boldsymbol{z}_t^a)$ and represents additional variation modeled as a function of the recorded respiratory-support condition. Because other concurrent treatments are not explicitly separated in this formulation, the two components should be interpreted as a model-based decomposition rather than as isolated treatment effects.

We use room air as a reference condition for the decomposition. During training, a soft regularization term encourages the respiratory-support-associated deviation to be small for room-air steps. Under the corresponding hard reference constraint, the additive decomposition is uniquely defined relative to the room-air reference; the formal result is provided in Appendix D.2. We use this decomposition to organize the predicted dynamics rather than to attribute physiological change to a causal treatment effect.

The next physiological state is obtained using a residual update:

$$\hat{\boldsymbol{x}}_{t+1} = \tilde{\boldsymbol{x}}_t + \Delta\hat{\boldsymbol{x}}_t, \tag{6}$$

where $\tilde{\boldsymbol{x}}_t = \boldsymbol{x}_t$ at the first prediction step and $\tilde{\boldsymbol{x}}_t = \hat{\boldsymbol{x}}_t$ thereafter.

Physiological response representation. After predicting the next physiological state, we define the short-term physiological response as the change between consecutive states:

$$\hat{\boldsymbol{r}}_t = \hat{\boldsymbol{x}}_{t+1} - \tilde{\boldsymbol{x}}_t. \tag{7}$$

The response is then encoded as

$$\boldsymbol{z}_t^r = E_{\text{response}}\left(\hat{\boldsymbol{r}}_t\right), \tag{8}$$

where E_{response} is a learnable encoder used for latent-state updating. During free-running prediction, $\hat{\boldsymbol{r}}_t = \Delta\hat{\boldsymbol{x}}_t$. The response pathway therefore does not introduce additional information, but provides a separate representation of the predicted change for the state-update module.

Response-aware latent-state updating. The latent patient state is updated using the predicted physiological state, the state-dependent respiratory-support representation, and the encoded physiological response:

$$\boldsymbol{s}_{t+1} = U_{\text{state}}\left(\boldsymbol{s}_t, \hat{\boldsymbol{x}}_{t+1}, \boldsymbol{z}_t^a, \boldsymbol{z}_t^r\right), \tag{9}$$

where U_{state} denotes the state-update module.

The update uses both the predicted physiological level and an explicit representation of its change from the previous state. The resulting latent state is used to compute the baseline dynamics and respiratory-support representation for the next prediction step: $F_{\text{base}}(\boldsymbol{s}_{t+1})$ and $E_{\text{action}}(\boldsymbol{s}_{t+1}, \boldsymbol{a}_{t+1})$, respectively. Appendix D.1 provides additional analysis of the response pathway and its role in latent-state updating.

3.3 TRAINING

Response-guided autoregressive training. The model is trained using multi-step autoregressive trajectory prediction. After initialization from the observed history, future physiological states are generated recursively:

$$\hat{\boldsymbol{x}}_{t+h} = \hat{\boldsymbol{x}}_{t+h-1} + \Delta\hat{\boldsymbol{x}}_{t+h-1}, \qquad h = 2, \ldots, H. \tag{10}$$

At each prediction step, the model-generated physiological state is carried forward to the next step; future ground-truth physiological states are not used as inputs to the trajectory rollout.

To stabilize training of the response pathway, we use a short response-guided schedule during the initial training epochs. For patient i and prediction step h, the observed physiological response is

$$\boldsymbol{r}_{i,h}^{\text{obs}} = \boldsymbol{x}_{i,t+h} - \boldsymbol{x}_{i,t+h-1}, \tag{11}$$

and the model-generated response is

$$\boldsymbol{r}_{i,h}^{\text{pred}} = \hat{\boldsymbol{x}}_{i,t+h} - \tilde{\boldsymbol{x}}_{i,t+h-1}, \tag{12}$$

where $\tilde{\boldsymbol{x}}_{i,t+h-1}$ denotes the state used as the reference for computing the response: it is the final observed physiological state at the first prediction step and the previous model-generated state at all subsequent steps.

During the response-guided period, the response supplied to the latent-state update is sampled independently for each patient and prediction step:

$$\boldsymbol{r}_{i,h}^{\text{use}} = b_{i,h}\boldsymbol{r}_{i,h}^{\text{obs}} + (1 - b_{i,h})\,\boldsymbol{r}_{i,h}^{\text{pred}}, \tag{13}$$

where

$$b_{i,h} \sim \text{Bernoulli}(\rho_e), \tag{14}$$

and ρ_e denotes the probability of using the observed response at epoch e.

The probability ρ_e is reduced to zero over the initial training period. Observed responses are used only as inputs to the response pathway for latent-state updating; the predicted physiological state is always carried forward in the autoregressive rollout. After the response-guided period, training uses only model-generated responses, matching the response input used during validation and testing.

Respiratory-support regularization. Room air is used as the no-active-support reference condition. For prediction steps corresponding to room air, we penalize the magnitude of the effective respiratory-support-associated deviation:

$$\mathcal{L}_{\text{RA}} = \frac{1}{|\mathcal{I}_{\text{RA}}|} \sum_{(i,h)\in\mathcal{I}_{\text{RA}}} \left|\boldsymbol{g}_{i,h} \odot F_{\text{treat}}(\boldsymbol{z}_{i,h}^{a})\right|_2^2, \tag{15}$$

where $\mathcal{I}_{\text{RA}}$ denotes prediction steps corresponding to room air.

This regularization applies only to the respiratory-support-associated deviation. It does not assume that physiology remains unchanged during room-air periods; physiological change can still be represented by $F_{\text{base}}(\boldsymbol{s}_{i,h})$. The penalty therefore provides a soft room-air reference for the decomposition. Under the corresponding hard reference constraint, the additive decomposition is uniquely defined relative to room air; the formal result is provided in Appendix D.2.

Training objective. Let $m_{i,h,j}$ indicate whether physiological variable j is observed for patient i at prediction horizon h. The trajectory-prediction loss is

$$\mathcal{L}_{\text{pred}} = \frac{\sum_{i=1}^{N}\sum_{h=1}^{H}\sum_{j=1}^{p} m_{i,h,j}\left(x_{i,t+h,j} - \hat{x}_{i,t+h,j}\right)^2}{\sum_{i=1}^{N}\sum_{h=1}^{H}\sum_{j=1}^{p} m_{i,h,j}}, \tag{16}$$

where the observation mask restricts the loss to available target measurements.

The complete training objective is

$$\mathcal{L} = \mathcal{L}_{\text{pred}} + \lambda_{\text{RA}}\mathcal{L}_{\text{RA}}, \tag{17}$$

where λ_{RA} controls the strength of the room-air regularization.

Free-running prediction. During validation, testing, and external evaluation, no future physiological measurements or observed responses are provided. Predicted states are recursively carried forward, and the model-generated response satisfies $\hat{\boldsymbol{r}}_t = \Delta\hat{\boldsymbol{x}}_t$. Thus, after initialization, prediction depends only on model-generated states and the specified respiratory-support sequence.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETTING

Datasets. We conducted a retrospective study using de-identified EHR data from adult ICU patients ($\geq$18 years). Data from January 1, 2016, to December 31, 2023, were divided into a development cohort (*Dev*) for model training and an internal validation cohort (*Val*) for model selection. We additionally evaluated the model on a later-period cohort from the same institution (*Temporal*; January 1, 2024–June 30, 2024) and an independent cohort from a separate institution (*External*; January 1, 2023–August 31, 2024). All analyses were conducted under institutional review board approval with waiver of informed consent. Cohort characteristics are summarized in Appendix Table 4, with cohort construction and preprocessing details provided in Appendix B.

Implementation Details. We optimize the model using AdamW (Loshchilov & Hutter, 2017). Model selection is based on the lowest validation RMSE, with early stopping using a patience of 10 epochs. We use a four-epoch response-guided warm-up, with ρ_e linearly decreased from 1 to 0 across the four epochs. After the warm-up period, $\rho_e = 0$ and training uses only model-generated responses. Hyperparameters are tuned on the internal validation cohort using Optuna (Akiba et al., 2019). We search over learning rate, latent-state dimension, action-representation dimension, network depth, dropout, weight decay, and λ_{RA}. All learning-based baselines use the same training and validation splits, preprocessing pipeline, physiological variables, respiratory-support inputs, prediction horizon, and model-selection objective. Hyperparameters are tuned independently for each baseline using the same validation objective and a comparable search space and number of trials.

Evaluation. We evaluate multi-step physiological trajectory prediction in a fully free-running setting. After initialization from a 6-hour observed history, the model recursively predicts the following 6 hours of physiology without access to future ground-truth physiological states. We use a 6-hour prediction horizon to focus on short-term physiological changes during respiratory support, which can be clinically informative for assessing evolving respiratory status and deterioration (Roca et al., 2019). Performance is evaluated only at observed target measurements.

Our primary metric is root mean squared error (RMSE), reported over the full 1–6 h prediction horizon and separately at each forecast hour. Overall RMSE is computed by pooling all observed feature-level targets across the six forecast hours and taking the square root of the mean squared error. We also report mean absolute error (MAE). Because persistence can perform strongly when physiology changes little over short intervals, we separately evaluate targets with non-trivial physiological changes. For each feature and forecast horizon, a target is included when $|x_{t+h} - x_t| > 0.25$ SD and the ground-truth value is observed, where SD is the training-set standard deviation of that variable. This criterion removes near-stable targets while accounting for differences in scale across physiological variables. We report MAE and RMSE on these targets, together with the relative RMSE improvement over persistence. We further stratify performance by respiratory-support modality, including room air, conventional oxygen therapy, HFNC, and NIV.

4.2 ABLATION AND COMPARISON WITH BASELINES

Ablation study. Table 1 evaluates the effects of respiratory-support-associated dynamics and the response-aware state update. On the validation cohort, adding respiratory-support-associated dynamics reduces MSE from 0.270 to 0.268 and MC-MAE from 0.736 to 0.726, while the full model further reduces MSE to 0.266. On the temporal and external cohorts, respiratory-support-associated dynamics alone do not consistently improve performance, whereas the response-aware update recovers or improves several metrics, including MC-MAE. These results suggest that respiratory-support conditioning alone is insufficient, and that feeding predicted physiological change back into the latent state improves the stability of the learned dynamics across cohorts.

Table 1: Ablation study of the proposed patient dynamics model across three evaluation cohorts.

Val					
Model	**MSE ↓**	**1-h RMSE ↓**	**3-h RMSE ↓**	**6-h RMSE ↓**	**MC-MAE ↓**
Base dynamics	0.270	0.369	0.504	0.621	0.736
+ Treatment-associated dynamics	0.268	**0.361**	0.506	**0.619**	**0.726**
+ Response-aware state update (Full)	**0.266**	**0.361**	0.505	**0.619**	**0.726**
Temporal					
Model	**MSE ↓**	**1-h RMSE ↓**	**3-h RMSE ↓**	**6-h RMSE ↓**	**MC-MAE ↓**
Base dynamics	0.304	0.398	0.545	**0.646**	0.728
+ Treatment-associated dynamics	0.304	0.396	0.538	0.653	0.772
+ Response-aware state update (Full)	0.304	**0.390**	**0.537**	0.658	**0.723**
External					
Model	**MSE ↓**	**1-h RMSE ↓**	**3-h RMSE ↓**	**6-h RMSE ↓**	**MC-MAE ↓**
Base dynamics	0.778	0.624	**0.819**	**1.061**	0.911
+ Treatment-associated dynamics	0.814	0.625	0.835	1.103	0.904
+ Response-aware state update (Full)	**0.775**	**0.620**	**0.819**	1.062	**0.901**

MC-MAE denotes mean absolute error on physiologically meaningful-change targets, defined as $|x_{t+h} - x_t| > 0.25$ SD. Lower values indicate better performance.

Baseline models. We compare against representative physiological trajectory and clinical world-model baselines: Persistence, Action-GRU, SepsisAgent-WM (Wu et al., 2026), and MedDreamer-WM (Xu et al., 2026). Persistence carries the most recent observed physiological state forward across the prediction horizon. Action-GRU is a standard action-conditioned recurrent model for autoregressive physiological forecasting. SepsisAgent-WM and MedDreamer-WM are adaptations of recent clinical world models to our respiratory trajectory prediction task. Because both literature-derived dynamics models use GRU-based transition architectures, we use a GRU-based action-conditioned model as the standard recurrent baseline to provide a more direct comparison of the dynamics formulations. For the literature-derived models, we preserve their core transition architectures and modify only the task-specific input and output interfaces required for our setting. Additional baseline details are provided in Appendix C.

Comparison with baselines. Table 2 shows that the learning-based models have similar overall RMSE, but their performance differs across forecast horizons. On the validation cohort, MedDreamer-WM has lower RMSE during the first 3 hours, whereas our model has lower RMSE from 4 to 6 hours. On the temporal cohort, MedDreamer-WM and our model have the same overall RMSE of 0.550; MedDreamer-WM performs better at the earliest horizons, while our model has lower RMSE at 3–5 h. On the external cohort, our model achieves the lowest RMSE consistently from 2 to 5 h, although SepsisAgent-WM has the lowest overall RMSE.

These results suggest different strengths across models. MedDreamer-WM is stronger at the earliest forecast steps, whereas our model is more competitive over intermediate and later horizons, when predictions depend increasingly on repeated state updates. Together with the ablation results, this pattern is consistent with the proposed formulation: decomposing physiological change and feeding the predicted change back into the latent state appears most useful during multi-step autoregressive rollout, rather than at the first prediction step.

Meaningful physiological changes. We further evaluate targets with $|x_{t+h} - x_t| > 0.25$ SD (Table 3). On the validation cohort, our model achieves an RMSE of 1.016, corresponding to a 10.1% improvement over persistence, compared with 9.2% for MedDreamer-WM. On the temporal cohort, the corresponding improvements are 9.7% and 9.6%, respectively. On the external cohort, our model and Action-GRU both achieve an RMSE of 1.771, corresponding to a 4.2% improvement over persistence. MedDreamer-WM has the lowest external-cohort MAE (0.897), while our model achieves 0.906.

The larger gains over persistence on these changing targets show that the differences between models become more apparent when physiology is evolving. Our model remains among the best-performing methods across all three cohorts, which is consistent with the intended role of explicitly representing physiological change and feeding the predicted change back into the latent state during rollout.

Table 2: Comparison with clinical world models for multi-step physiological trajectory prediction (values are RMSE; lower is better).

Model	Overall	1h	2h	3h	4h	5h	6h
Val							
Persistence	0.549	0.370	0.469	0.535	0.591	0.619	0.659
Action-GRU	0.520	0.355	0.445	0.507	0.559	0.585	0.624
SepsisAgent-WM	0.519	0.355	0.445	0.508	0.559	0.585	0.624
MedDreamer-WM	**0.516**	**0.344**	**0.435**	**0.502**	0.556	0.584	0.624
Ours	**0.516**	0.361	0.440	0.505	**0.550**	**0.578**	**0.619**
Temporal							
Persistence	0.587	0.400	0.515	0.583	0.627	0.658	0.694
Action-GRU	0.552	0.383	0.492	0.546	0.585	0.617	**0.652**
SepsisAgent-WM	0.552	0.383	0.495	0.547	0.585	**0.615**	0.649
MedDreamer-WM	**0.550**	**0.369**	**0.485**	0.543	0.584	0.618	0.654
Ours	**0.550**	0.389	0.487	**0.539**	**0.579**	**0.615**	0.653
External							
Persistence	0.901	0.624	0.729	0.840	0.979	1.057	1.084
Action-GRU	0.880	**0.613**	0.707	0.822	0.958	1.028	1.061
SepsisAgent-WM	**0.878**	**0.613**	0.708	0.821	0.959	1.028	**1.055**
MedDreamer-WM	0.889	0.610	0.708	0.829	0.970	1.044	1.076
Ours	0.880	0.620	**0.706**	**0.819**	**0.957**	**1.027**	1.062

Table 3: Performance on physiologically meaningful changes ($|x_{t+h} - x_t| > 0.25$ SD). Δ denotes relative RMSE improvement over persistence.

	Val			Temporal			External		
Model	**MAE**	**RMSE**	**Δ (%)**	**MAE**	**RMSE**	**Δ (%)**	**MAE**	**RMSE**	**Δ (%)**
Persistence	0.831	1.131	–	0.828	1.224	–	1.002	1.848	–
Action-GRU	0.738	1.036	+8.4	0.733	1.118	+8.7	0.911	**1.771**	**+4.2**
SepsisAgent-WM	0.742	1.041	+7.9	0.735	1.124	+8.2	0.911	1.777	+3.8
MedDreamer-WM	**0.726**	1.027	+9.2	0.718	1.106	+9.6	**0.897**	1.783	+3.5
Ours	**0.726**	**1.016**	**+10.1**	**0.710**	**1.105**	**+9.7**	0.906	**1.771**	**+4.2**

4.3 Clinical Stratified Analysis

Performance across respiratory-support states. We stratify performance by respiratory-support state to assess whether the overall results are consistent across support conditions or driven by a single subgroup. Figure 2 shows the relative RMSE improvement over persistence for room air, conventional oxygen, HFNC, and NIV. Our model remains competitive across all support states and achieves the lowest RMSE in several subgroups, including conventional oxygen and NIV on the validation cohort, room air and HFNC on the temporal cohort, and conventional oxygen and NIV on the external cohort. Differences between learning-based models are generally small, particularly for HFNC and the smaller NIV subgroups. Overall, these results suggest that the model's performance is not driven by a single respiratory-support category. Full subgroup results are provided in Appendix E.

Performance across physiological variables. We further analyze performance at the feature level to determine whether the improvements on changing physiological states are driven by a small number of variables or are distributed across different aspects of patient physiology. We focus on selected variables related to respiratory status and hemodynamics, including oxygen saturation (O_2Sat), respiratory rate (Resp), fraction of inspired oxygen (FiO_2), arterial oxygen and carbon dioxide partial pressures (PaO_2 and $PaCO_2$), pH, heart rate (HR), mean arterial pressure (MAP), and systolic and diastolic blood pressure (SBP and DBP). For each feature, evaluation is restricted to targets satisfying $|x_{t+h} - x_t| > 0.25$ SD.

Figure 3 shows that the improvement over persistence is distributed across multiple variables rather than being driven by a single measurement. Our model achieves the lowest RMSE for 6 of 10 variables on the validation cohort, 5 of 10 on the temporal cohort, with a tie for HR, and 8 of 10 on the external cohort. The largest gains are observed for several respiratory variables, particularly FiO_2,

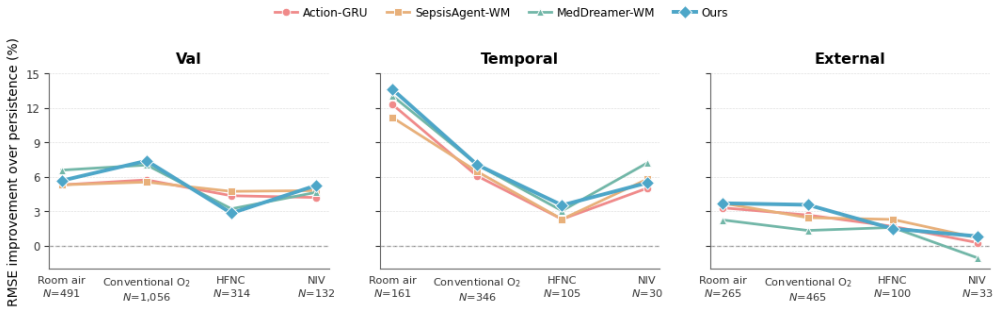

Figure 2: Relative RMSE improvement over persistence across respiratory-support modalities in the validation, temporal, and external cohorts. Positive values indicate lower RMSE than persistence. N denotes the number of evaluation windows in each subgroup.

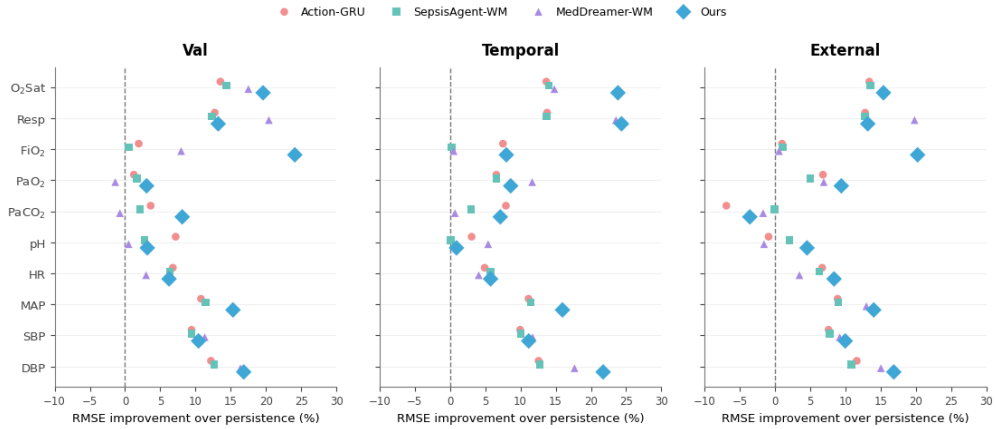

Figure 3: Relative RMSE improvement over persistence across respiratory-support modalities in the validation, temporal, and external cohorts. Positive values indicate lower RMSE than persistence. N denotes the number of evaluation windows in each subgroup.

while improvements are also seen for hemodynamic variables such as MAP and DBP. These results suggest that the benefit of the response-aware formulation is not limited to a single physiological variable or measurement type. Full absolute RMSE values and the number of evaluated targets for each feature are reported in Appendix Table 6.

5 Conclusion

We introduced a response-aware patient dynamics model for short-term physiological forecasting under respiratory support. The model decomposes predicted physiological change into state-dependent baseline dynamics and respiratory-support-associated deviations, and explicitly represents the predicted change during latent-state updating. Across internal, temporal, and external cohorts, the model achieved overall performance comparable to strong patient dynamics baselines, with more consistent improvements when physiological states were changing.

The current study is retrospective and focuses on observational trajectory prediction rather than causal treatment-effect estimation or treatment recommendation. Future work will investigate whether the learned dynamics can support counterfactual trajectory analysis and respiratory-support decision modeling, with explicit treatment of confounding and prospective validation.

AI USE DISCLOSURE

Generative AI tools were used to assist with manuscript drafting and editing. All AI-assisted text is independently reviewed and verified by the authors. The authors take full responsibility for the final content of the paper.

ETHICS STATEMENT

This study uses retrospective, de-identified electronic health record data from multiple clinical institutions under applicable institutional review, data-use, and data-governance requirements. Patient consent was waived where permitted for retrospective analysis of de-identified data. All analyses were conducted in accordance with institutional privacy and security policies, and no attempt was made to re-identify individual patients. The underlying clinical data cannot be publicly released where restricted by institutional policies or data-use agreements.

The proposed model is intended solely as a research framework for modeling short-term physiological trajectories under respiratory support and is not designed for autonomous clinical decision-making. Model predictions should not be interpreted as guaranteed patient responses or as recommendations for respiratory-support selection. Any prospective clinical use would require additional external and prospective validation, assessment of failure modes, clinical oversight, and appropriate regulatory review.

Clinical practice patterns, patient populations, measurement processes, and respiratory-support use may differ across institutions. Although we evaluate the model across multiple cohorts to assess robustness under distribution shift, residual bias and under-representation may remain, and performance may not generalize to all patient groups or care settings.

Finally, the decomposition of predicted physiological change into underlying progression and respiratory-support-associated components is a modeling construct and should not be interpreted as identifying a causal treatment effect. Respiratory interventions are observed rather than randomized and may be influenced by measured and unmeasured clinical factors. Counterfactual claims regarding the effectiveness of alternative respiratory-support strategies are therefore outside the scope of this work.

REPRODUCIBILITY STATEMENT

We provide detailed descriptions of cohort construction, trajectory generation, preprocessing, respiratory-support representation, model architecture, training, and evaluation in the main text and appendix. Each modeling example consists of a fixed observed history followed by a fixed prediction horizon, and all reported trajectory results are obtained using free-running autoregressive prediction without access to future ground-truth physiology.

The architecture specification includes latent-state initialization, state-dependent respiratory-support encoding, decomposition of physiological change into underlying and treatment-associated components, feature-wise treatment gating, recurrent latent-state updating, and all training objectives and regularization terms. We also document the optimization settings, hyperparameters, horizon weighting scheme, model-selection procedure, and random seeds used in the reported experiments.

We evaluate the proposed method using predefined trajectory-level and horizon-specific error metrics across internal and external cohorts, together with baseline comparisons, ablation studies, and sensitivity analyses. Subject to institutional and data-use restrictions, we will release the implementation of preprocessing, trajectory construction, model training, free-running evaluation, baseline models, and scripts used to generate the reported tables and figures, together with configuration files sufficient to reproduce the experiments on appropriately formatted data. The underlying clinical datasets cannot be publicly released where restricted by institutional policies or data-use agreements.

REFERENCES

Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework. In *Proceedings of the 25th ACM*

SIGKDD international conference on knowledge discovery & data mining, pp. 2623–2631, 2019.

JR Azevedo, WS Montenegro, AL Leitao, MM Silva, JS Prazeres, and JP Maranhao. High flow nasal cannula oxygen (hfnc) versus non-invasive positive pressure ventilation (nippv) in acute hypoxemic respiratory failure. a pilot randomized controlled trial. *Intensive Care Medicine Experimental*, 3(Suppl 1):A166, 2015.

Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In Alessandro Moschitti, Bo Pang, and Walter Daelemans (eds.), *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734, Doha, Qatar, October 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1179. URL https://aclanthology.org/D14-1179/.

Remi Coudroy, Jean-Pierre Frat, Stephan Ehrmann, Frederic Pene, Maxens Decavèle, Nicolas Terzi, Gwenael Prat, Charlotte Garret, Damien Contou, Arnaud Gacouin, et al. High-flow nasal oxygen alone or alternating with non-invasive ventilation in critically ill immunocompromised patients with acute respiratory failure: a randomised controlled trial. *The Lancet Respiratory Medicine*, 10(7):641–649, 2022.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

Silvia Mongodi, Erminio Santangelo, Domenico Luca Grieco, Valeria Musella, Nello De Vita, Rosanna Vaschetto, Bélaïd Bouhemad, and Francesco Mojoli. Early lung ultrasound score changes predict the failure of non-invasive respiratory supports in acute hypoxemic patients: a multicenter prospective observational study. *Critical Care*, 30(1):58, 2026.

Joaquín Perez, Luciano Brandan, and Irene Telias. Monitoring patients with acute respiratory failure during non-invasive respiratory support to minimize harm and identify treatment failure. *Critical Care*, 29(1):147, 2025.

Mohammad Areeb Qazi, Maryam Nadeem, and Mohammad Yaqub. Beyond generative ai: World models for clinical prediction, counterfactuals, and planning. *arXiv preprint arXiv:2511.16333*, 2025.

Oriol Roca, Berta Caralt, Jonathan Messika, Manuel Samper, Benjamin Sztrymf, Gonzalo Hernández, Marina García-de Acilu, Jean-Pierre Frat, Joan R Masclans, and Jean-Damien Ricard. An index combining respiratory rate and oxygenation to predict outcome of nasal high-flow therapy. *American journal of respiratory and critical care medicine*, 199(11):1368–1376, 2019.

Vipin Das Villgran, Caitlan Lyons, Adeel Nasrullah, Charmaine Clarisse Abalos, Eric Bihler, and Ahmad Alhajhusain. Acute respiratory failure. *Critical care nursing quarterly*, 45(3):233–247, 2022.

Minghao Wu, Yuting Yan, Zhenyang Cai, Ke Ji, Chuangsen Fang, Ziying Sheng, Xidong Wang, Rongsheng Wang, Hejia Zhang, Shuang Li, et al. Agentifying patient dynamics within llms through interacting with clinical world model. *arXiv preprint arXiv:2605.14723*, 2026.

Qianyi Xu, Gousia Habib, Feng Wu, Dilruk Perera, and Mengling Feng. Meddreamer: Model-based reinforcement learning with latent imagination on complex ehrs for clinical decision support. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*, pp. 1693–1704, 2026.

Yijun Yang, Zhao-Yang Wang, Qiuping Liu, Shuwen Sun, Kang Wang, Rama Chellappa, Zongwei Zhou, Alan Yuille, Lei Zhu, Yu-Dong Zhang, et al. Medical world model: Generative simulation of tumor evolution for treatment planning. *arXiv preprint arXiv:2506.02327*, 2025.

A Model architecture overview

Figure 4 summarizes the response-aware patient dynamics framework and the information flow across the autoregressive prediction horizon.

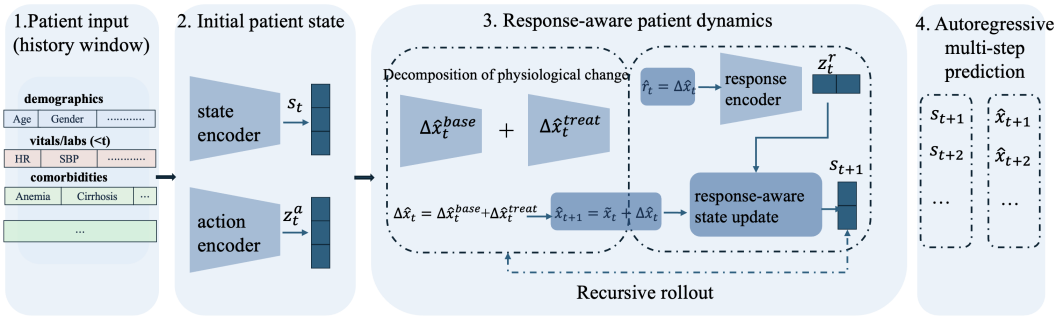

Figure 4: Overview of the response-aware patient dynamics model. The model initializes a latent patient state from recent clinical history and predicts physiological change as the sum of state-dependent baseline dynamics and a respiratory-support-associated deviation. The predicted change produces the next physiological state through a residual update and is also explicitly encoded as the physiological response for latent-state updating. The resulting state is recursively propagated over the prediction horizon.

Table 4: Characteristics of the study cohorts.

Characteristic	**Dev**	**Val**	**Temporal**	**External**
Patients, N	6,762	1,912	564	791
Age, years, mean ± SD	62.9 ± 16.4	62.1 ± 16.6	64.3 ± 16.5	64.3 ± 17.2
Female, N (%)	3,171 (42.5%)	820 (41.7%)	257 (44.4%)	360 (44.4%)
Race, N (%)				
White	3,884 (52.1%)	963 (49.0%)	274 (47.3%)	393 (48.5%)
Black or African American	609 (8.2%)	146 (7.4%)	49 (8.5%)	19 (2.3%)
Asian	563 (7.5%)	152 (7.7%)	50 (8.6%)	182 (22.5%)
Other/Unknown	2,402 (32.2%)	704 (35.8%)	206 (35.6%)	216 (26.7%)
Prediction windows, N	7,972	1,993	642	863
Available target observations, N	1,046,277	264,366	85,590	116,805
Meaningful-change targets, N (%)	27,981 (2.7%)	7,316 (2.8%)	2,073 (2.4%)	2,173 (1.9%)
Respiratory support at prediction start, N (%)				
Room air	1,955 (24.5%)	510 (25.6%)	158 (24.6%)	262 (30.4%)
Conventional O_2	4,362 (54.7%)	1,046 (52.5%)	348 (54.2%)	466 (54.0%)
HFNC	1,188 (14.9%)	312 (15.7%)	107 (16.7%)	104 (12.1%)
NIV	467 (5.9%)	125 (6.3%)	29 (4.5%)	31 (3.6%)

B Data Preprocessing and Trajectory Construction

Table 4 summarizes cohort demographics, prediction-window counts, target availability, and respiratory-support distributions across the development, internal validation, temporal, and external cohorts.

Trajectory representation. Each patient encounter was discretized into hourly bins and represented as a longitudinal trajectory containing physiological measurements, respiratory-support information, time-since-last-measurement variables, and static patient characteristics. Trajectories extended from ICU admission until transfer out of the ICU or initiation of invasive mechanical ventilation (IMV), whichever occurred first. Respiratory-support actions before IMV included room air, conventional oxygen therapy, high-flow nasal cannula (HFNC), and non-invasive ventilation (NIV).

At each hour t, the physiological state $\boldsymbol{x}_t \in \mathbb{R}^{50}$ contained 50 dynamic clinical variables. Each variable was paired with a corresponding time-since-last-measurement feature, forming a 50-dimensional TSLM vector. Static patient information consisted of 36 demographic and clinical variables and 62 comorbidity indicators, which were fixed for each trajectory and remained unchanged during autoregressive prediction.

Physiological feature preprocessing. Missing values in the 50 physiological variables were handled using feature-wise median imputation. To avoid information leakage, all imputation parameters were estimated using the training cohort only. For feature j, let m_j^{tr} denote the median of the observed training values. A missing observation was replaced by

$$\tilde{x}_{t,j} = \begin{cases} x_{t,j}, & x_{t,j} \text{ observed}, \\ m_j^{\mathrm{tr}}, & x_{t,j} \text{ missing}. \end{cases} \tag{18}$$

After imputation, each physiological variable was standardized using the mean and standard deviation estimated from the imputed training data. Let μ_j^{tr} and σ_j^{tr} denote these statistics. The normalized physiological variable was computed as

$$x_{t,j}^{\mathrm{norm}} = \frac{\tilde{x}_{t,j} - \mu_j^{\mathrm{tr}}}{\sigma_j^{\mathrm{tr}}}. \tag{19}$$

The same training-derived median, mean, and standard deviation were applied without re-estimation to the validation and external cohorts. We additionally retained a binary missingness indicator for each physiological variable, denoting whether the corresponding value had originally been observed or imputed. The time-since-last-measurement variables were retained separately from the normalized physiological measurements and were provided to the initial state encoder.

Respiratory intervention representation. The respiratory support administered at each time point was represented by a treatment vector $\boldsymbol{a}_t$ containing the respiratory support modality, device subtype, available device settings, and indicators describing whether individual settings were observed. The support modality and device subtype were categorical variables and were encoded using one-hot vocabularies constructed from the training cohort.

Continuous respiratory support settings consisted of $\mathrm{FiO_2}$, oxygen flow, IPAP, and EPAP. These settings are not applicable to every respiratory support modality. For example, flow may be available for high-flow oxygen support, whereas IPAP and EPAP are primarily associated with non-invasive positive-pressure ventilation. We therefore paired each continuous setting with a binary availability mask. For a treatment setting q_t with corresponding mask r_t, only training observations satisfying $r_t = 1$ were used to estimate its normalization statistics. When the setting was available, it was normalized as

$$q_t^{\mathrm{norm}} = \frac{q_t - \mu_q^{\mathrm{tr}}}{\sigma_q^{\mathrm{tr}}}, \qquad r_t = 1, \tag{20}$$

where μ_q^{tr} and σ_q^{tr} were estimated from valid training observations. When $r_t = 0$, the normalized setting was assigned a value of zero while the availability mask remained zero, which allows the model to distinguish an unavailable treatment setting from an observed value whose standardized value happens to be close to zero.

The resulting treatment representation can be written as

$$\boldsymbol{a}_t = \left[\boldsymbol{a}_t^{\mathrm{mod}}, \boldsymbol{a}_t^{\mathrm{sub}}, \tilde{f}_t^{\mathrm{FiO_2}}, \tilde{f}_t^{\mathrm{flow}}, \tilde{f}_t^{\mathrm{IPAP}}, \tilde{f}_t^{\mathrm{EPAP}}, r_t^{\mathrm{FiO_2}}, r_t^{\mathrm{flow}}, r_t^{\mathrm{IPAP}}, r_t^{\mathrm{EPAP}} \right], \tag{21}$$

where $\boldsymbol{a}_t^{\mathrm{mod}}$ and $\boldsymbol{a}_t^{\mathrm{sub}}$ denote the one-hot modality and subtype representations, respectively. Encounters containing an unknown respiratory support modality were excluded from model development and evaluation.

Trajectory construction. We constructed fixed-length trajectories consisting of a 6-hour observed history followed by a 6-hour prediction horizon. The 6-hour horizon was chosen to focus on short-term physiological changes during respiratory support, which can provide clinically relevant information about evolving respiratory status and deterioration (Roca et al., 2019).

For an encounter with a selected starting index k, the observed physiological history was defined as

$$\boldsymbol{X}_{k:k+5} = \{\boldsymbol{x}_k, \boldsymbol{x}_{k+1}, \ldots, \boldsymbol{x}_{k+5}\}, \tag{22}$$

with corresponding time-since-last-measurement sequence

$$\boldsymbol{L}_{k:k+5} = \{\boldsymbol{\ell}_k, \boldsymbol{\ell}_{k+1}, \ldots, \boldsymbol{\ell}_{k+5}\}. \tag{23}$$

The subsequent 6-hour physiological trajectory was

$$\boldsymbol{X}_{k+6:k+11} = \{\boldsymbol{x}_{k+6}, \boldsymbol{x}_{k+7}, \ldots, \boldsymbol{x}_{k+11}\}, \tag{24}$$

and the respiratory support sequence over the prediction horizon was

$$\boldsymbol{A}_{k+6:k+11} = \{\boldsymbol{a}_{k+6}, \boldsymbol{a}_{k+7}, \ldots, \boldsymbol{a}_{k+11}\}. \tag{25}$$

Encounters containing fewer than 12 hourly observations were excluded because they could not provide both the complete history and prediction horizon. Encounters containing exactly 12 observations were used in full. For encounters with more than 12 observations, one valid contiguous 12-hour window was randomly selected from the available trajectory. Window selection was performed once during dataset construction, after which the resulting training, validation, and external cohorts were fixed and reused across all experiments.

Each resulting sample therefore consists of

$$\left(\boldsymbol{X}_{k:k+5}, \boldsymbol{L}_{k:k+5}, \boldsymbol{A}_{k:k+5}, \boldsymbol{z}^{\text{static}}, \boldsymbol{X}_{k+6:k+11}, \boldsymbol{A}_{k+6:k+11}\right), \tag{26}$$

where $\boldsymbol{z}^{\text{static}}$ contains the static demographic, clinical, and comorbidity variables. The observed history and static information are used to construct the initial latent patient state, whereas the respiratory interventions over the prediction horizon condition the subsequent autoregressive physiological rollout.

External cohort preprocessing. External cohorts were transformed using exactly the same preprocessing parameters estimated from the internal training cohort. In particular, missing physiological observations were imputed using the internal training medians, physiological variables were normalized using the internal training means and standard deviations after imputation, and respiratory support variables were encoded using the treatment vocabulary and continuous-setting normalization parameters learned from the internal training cohort. No preprocessing parameters were re-estimated on external data, which ensured that external evaluation reflected deployment of a fixed model and fixed preprocessing pipeline to previously unseen patient populations.

C BASELINES

We compare against a set of representative physiological trajectory and clinical world-models. GRU-based recurrent models are commonly used for longitudinal EHR dynamics, and recent clinical world models such as SepsisAgent (Wu et al., 2026) and medDreamer (Xu et al., 2026) also adopt GRU-based transition architectures. We therefore include both a standard action-conditioned GRU baseline and adaptations of these recent clinical world models to the respiratory trajectory prediction setting.

To ensure a fair comparison, all learning-based methods use the same training and validation splits, preprocessing pipeline, physiological input variables, treatment representations, prediction horizons, and observation masks. Hyperparameters are tuned independently for each method using Optuna (Akiba et al., 2019) with the same number of trials and the same validation objective. We search over learning rate, hidden dimension, number of layers, dropout, and weight decay, and select the configuration achieving the lowest validation RMSE. The same early-stopping criterion is applied across methods. For literature-derived world models, we preserve their core transition architectures and modify only the task-specific input and output interfaces required for respiratory trajectory prediction.

- **Persistence**: A non-parametric baseline that assumes the patient's physiological state remains unchanged over the prediction horizon. Specifically, the most recent observed state is carried forward as the prediction for all future time points.
- **Action-GRU**: A standard action-conditioned recurrent dynamics baseline that encodes the observed physiological history with a GRU (Cho et al., 2014) and autoregressively predicts future states conditioned on respiratory support interventions.
- **SepsisAgent-WM (Wu et al., 2026)**: Following the Clinical World Model used in SepsisAgent, we use a GRU-based world model that encodes the full patient trajectory, including both physiological states and prior treatment actions, and predicts the next physiological state conditioned on the current respiratory support intervention.
- **medDreamer-WM (Xu et al., 2026)**: Following the latent world model in medDreamer, we model stochastic patient dynamics in a latent state space using GRU-based posterior and prior transition models. The posterior infers the latent state from observed physiology and treatment history, while the prior predicts future latent states conditioned on the previous latent state and respiratory support action. Future physiological states are reconstructed from the predicted latent trajectory.

D Additional Analysis of Response-Aware Dynamics

This section clarifies two design choices in the proposed model. We first describe the role of the explicit response pathway in latent-state updating. We then analyze the reference constraint used to separate baseline dynamics from respiratory-support-associated deviations. The analysis concerns the model parameterization and observational transition dynamics and does not establish causal treatment effects.

D.1 Explicit response representation

Let X_t denote the physiological state at time t, and define the short-term physiological response as

$$R_t = X_{t+1} - X_t. \tag{27}$$

In the proposed model, the response pathway is used to represent physiological change explicitly during latent-state updating. Under free-running prediction, the next physiological state is generated by

$$\hat{X}_{t+1} = \tilde{X}_t + \Delta\hat{X}_t, \tag{28}$$

and therefore

$$\hat{R}_t = \hat{X}_{t+1} - \tilde{X}_t = \Delta\hat{X}_t. \tag{29}$$

Thus, the model-generated response does not introduce an additional observation during inference. Instead, the predicted change is passed through a separate response encoder,

$$Z_t^R = E_{\text{response}}(\hat{R}_t), \tag{30}$$

and supplied to the state-update function together with the predicted physiological level:

$$S_{t+1} = U_{\text{state}}\left(S_t, \hat{X}_{t+1}, Z_t^A, Z_t^R\right). \tag{31}$$

The response pathway therefore provides an explicit representation of physiological change rather than an additional source of information at inference time. This parameterization allows the state-update module to process the predicted physiological level and its change through separate learned representations.

During the initial response-guided training period, the response encoder may instead receive the observed change

$$R_t^{\text{obs}} = X_{t+1} - X_t. \tag{32}$$

This supervision is used only for the response input to the latent-state update; the physiological trajectory itself remains autoregressive. The probability of using the observed response is reduced to zero after the initial training period, after which both training and evaluation use model-generated responses.

D.2 REFERENCE-ANCHORED DECOMPOSITION OF PHYSIOLOGICAL DYNAMICS

We consider the decomposition of the conditional physiological transition into baseline dynamics and a respiratory-support-associated deviation.

Let

$$F(s,a) = \mathbb{E}\left[\Delta X_t \mid S_t = s, A_t = a\right] \tag{33}$$

denote the conditional mean physiological change for latent state s and respiratory-support action a.

For the analysis, we write the model decomposition as

$$F(s,a) = f_{\text{base}}(s) + d(s,a), \tag{34}$$

where $f_{\text{base}}(s)$ denotes the shared state-dependent dynamics and $d(s,a)$ denotes the effective respiratory-support-associated deviation. In the implemented model, this deviation corresponds to

$$d(s,a) = g(s,a) \odot F_{\text{treat}}\left(E_{\text{action}}(s,a)\right). \tag{35}$$

Without a reference constraint, the decomposition in Eq. 34 is not unique. For any function $c(s)$,

$$f'_{\text{base}}(s) = f_{\text{base}}(s) + c(s), \tag{36}$$

$$d'(s,a) = d(s,a) - c(s), \tag{37}$$

and therefore

$$f'_{\text{base}}(s) + d'(s,a) = F(s,a). \tag{38}$$

We use room air as the reference respiratory-support condition. Let

$$a_0 = \text{room air}, \tag{39}$$

and consider the reference constraint

$$d(s,a_0) = 0. \tag{40}$$

Proposition 1 (Reference-anchored decomposition)**.** *Let $F(s,a)$ be a fixed conditional mean transition and let a_0 be a reference action. Among additive decompositions satisfying*

$$F(s,a) = f_{\text{base}}(s) + d(s,a) \tag{41}$$

and

$$d(s,a_0) = 0, \tag{42}$$

the decomposition is uniquely given by

$$f_{\text{base}}(s) = F(s,a_0), \tag{43}$$

$$d(s,a) = F(s,a) - F(s,a_0). \tag{44}$$

Proof. Evaluating the additive decomposition at a_0 gives

$$F(s,a_0) = f_{\text{base}}(s) + d(s,a_0). \tag{45}$$

Using $d(s,a_0) = 0$ yields

$$f_{\text{base}}(s) = F(s,a_0). \tag{46}$$

Substituting this expression back into Eq. 34 gives

$$d(s,a) = F(s,a) - F(s,a_0), \tag{47}$$

which uniquely determines both components.

Proposition 1 gives the decomposition a reference-relative interpretation. Under the hard reference constraint, the baseline component equals the conditional transition associated with room air, and the support-associated component represents the deviation from that reference transition.

The implemented model uses a soft rather than hard reference constraint. Specifically, during room-air prediction steps we penalize

$$\mathcal{L}_{\text{RA}} = \mathbb{E}\left[\|d(S_t,A_t)\|_2^2 \middle| A_t = a_0\right], \tag{48}$$

Table 5: Trajectory prediction performance stratified by respiratory support modality. RMSE is computed over all observed targets across the 1–6 h prediction horizon within each subgroup; lower is better. N denotes the number of evaluation windows.

Cohort	Support	N	Persistence	Action-GRU	SepsisAgent-WM	MedDreamer-WM	Ours
Val	Room air	491	0.546	0.517	0.517	0.510	0.515
	Conventional O_2	1,056	0.540	0.509	0.510	0.502	0.500
	HFNC	314	0.527	0.504	0.502	0.510	0.512
	NIV	132	0.665	0.637	0.633	0.634	0.630
Temporal	Room air	161	0.528	0.463	0.469	0.459	0.456
	Conventional O_2	346	0.510	0.479	0.477	0.474	0.474
	HFNC	105	0.819	0.800	0.800	0.794	0.790
	NIV	30	0.639	0.607	0.602	0.593	0.604
External	Room air	265	0.755	0.730	0.727	0.738	0.727
	Conventional O_2	465	0.895	0.871	0.873	0.883	0.863
	HFNC	100	1.000	0.983	0.977	0.984	0.985
	NIV	33	1.437	1.433	1.427	1.452	1.425

which corresponds to

$$\mathcal{L}_{\mathrm{RA}} = \mathbb{E}\left[\|\boldsymbol{g}_t \odot F_{\mathrm{treat}}(\boldsymbol{z}_t^a)\|_2^2 \middle| A_t = a_0\right]. \tag{49}$$

The learned model is therefore encouraged toward the reference-anchored decomposition but is not constrained to satisfy $d(s, a_0) = 0$ exactly. In particular, the proposition establishes uniqueness only under the corresponding hard reference constraint.

Finally, the decomposition is defined in terms of observational conditional transitions. The support-associated deviation should therefore not be interpreted as an identified causal effect of respiratory support. Respiratory support is not randomized in the retrospective data, and treatment assignment may depend on measured and unmeasured patient characteristics.

E Additional Stratified Results

Table 5 reports the corresponding absolute RMSE values for each respiratory-support subgroup.

Table 6: Feature-level RMSE on physiologically meaningful changes ($|x_{t+h} - x_t| > 0.25$ SD) across internal and external cohorts. N denotes the number of observed meaningful-change targets. Lower RMSE indicates better performance.

Feature	N	**Persistence**	**Action-GRU**	**SepsisAgent-WM**	**MedDreamer-WM**	**Ours**
Val						
O_2Sat	8,275	1.057	0.914	0.905	0.872	0.850
Resp	8,524	1.029	0.898	0.902	0.819	0.893
FiO_2	151	2.474	2.426	2.462	2.277	1.878
PaO_2	272	1.440	1.422	1.416	1.460	1.396
$PaCO_2$	278	1.652	1.592	1.617	1.664	1.518
pH	377	1.614	1.498	1.569	1.606	1.563
HR	6,107	0.737	0.687	0.690	0.715	0.691
MAP	7,161	1.012	0.903	0.896	0.858	0.857
SBP	6,966	0.911	0.825	0.825	0.808	0.816
DBP	7,345	1.075	0.944	0.939	0.899	0.894
Temporal						
O_2Sat	2,777	1.053	0.909	0.905	0.897	0.802
Resp	2,727	1.023	0.882	0.883	0.782	0.774
FiO_2	48	2.287	2.115	2.281	2.275	2.104
PaO_2	120	1.474	1.377	1.377	1.302	1.347
$PaCO_2$	104	1.134	1.044	1.100	1.126	1.053
pH	145	2.049	1.986	2.047	1.938	2.030
HR	1,969	0.711	0.676	0.670	0.682	0.670
MAP	2,211	0.996	0.885	0.882	0.838	0.837
SBP	2,086	0.904	0.814	0.813	0.798	0.803
DBP	2,206	1.058	0.925	0.923	0.871	0.828
External						
O_2Sat	3,649	1.194	1.034	1.032	1.011	1.010
Resp	3,766	1.312	1.144	1.144	1.052	1.139
FiO_2	85	2.365	2.341	2.338	2.351	1.886
PaO_2	58	1.729	1.611	1.642	1.609	1.566
$PaCO_2$	67	1.380	1.475	1.381	1.403	1.429
pH	82	2.145	2.164	2.101	2.178	2.047
HR	2,607	0.774	0.722	0.725	0.747	0.709
MAP	3,498	0.955	0.870	0.869	0.831	0.821
SBP	3,415	1.011	0.934	0.932	0.918	0.910
DBP	3,135	0.877	0.775	0.782	0.745	0.729