

<https://doi.org/10.1038/s41746-025-01863-0>

# Personalized and real time hemodynamic management in critical care using Dynamic Cohort Ensemble Learning (DynaCEL)



Lingzhong Meng<sup>1</sup>, Jiangqiong Li<sup>2</sup>, Xiang Liu<sup>2</sup>, Yanhua Sun<sup>1,3</sup>, Zuotian Li<sup>2,4</sup>, Jinjin Cai<sup>2,5</sup>, Ameya D. Parab<sup>6</sup>, George Lu<sup>2</sup>, Aishwarya Budhkar<sup>6</sup>, Saravanan Kanakasabai<sup>7</sup>, David C. Adams<sup>1</sup>, Ziyue Liu<sup>2</sup>, Xuhong Zhang<sup>6</sup> & Jing Su<sup>2</sup> ✉

Effective hemodynamic management in the intensive care unit requires individualized targets that adapt to dynamic clinical conditions. We developed Dynamic Cohort Ensemble Learning (DynaCEL), a real-time framework that recommends personalized heart rate and systolic blood pressure targets by modeling each time point post-intensive care unit admission as a distinct temporal cohort. Trained on eICU data and validated on MIMIC-IV and Indiana University Health datasets, DynaCEL demonstrated robust predictive performance (AUCs 0.83–0.91). In the MIMIC-IV cohort, proximity to DynaCEL-predicted targets was associated with lower 24-hour mortality compared to fixed targets, after adjustment using propensity score matching. Dose-response and comparative analyses revealed that greater deviations from personalized targets were associated with higher mortality. Case studies illustrated temporal and inter-individual variation in optimal targets. DynaCEL offers interpretable and scalable support for exploring precision hemodynamic management, although its clinical utility remains to be established in prospective trials.

Precisely identifying individualized and dynamic clinical management targets is a unique and essential task in daily intensive care unit (ICU) practice. Among ICU management priorities, effective hemodynamic monitoring and control are critical, as circulatory failure is a leading cause of ICU mortality, accounting for over 40% of deaths<sup>1,2</sup>. Among hemodynamic targets, heart rate (HR) and blood pressure (BP) are the most commonly used, as these two complementary, modifiable indicators are closely linked to mortality risk<sup>3–6</sup>. Therefore, setting and adjusting HR and BP targets, including defining “safe zones” and “risk zones,” are routine ICU tasks that guide interventions and provide baselines for assessing the risks of deterioration and acute death. Meanwhile, HR and BP targets vary widely and evolve dynamically among patients, necessitating personalized and adaptive approaches.

The relationship between HR, BP, and mortality in ICU patients is complex and influenced by baseline health, comorbidities, acute illnesses, and disease progression. Prior research has associated tachycardia with increased mortality in septic shock patients<sup>3</sup>, severe bradycardia with impending cardiac arrest<sup>7,8</sup>, and both hypotension<sup>4–6</sup> and hypertension<sup>6</sup> with increased mortality. Yet, there remains a clear gap between population-based evidence and personalized care. Clinical conditions change rapidly in the ICU, requiring continuous, real-time adjustment of hemodynamic targets.

Despite this complexity, current ICU practice largely relies on fixed, population-level targets, such as an HR of 80 beats per minute (bpm) and systolic BP (SBP) of 120 mmHg<sup>9</sup>, which does not account for patient-specific variability in physiology, comorbidities, or illness severity. Recent efforts to establish condition-specific targets have failed to meet the

<sup>1</sup>Department of Anesthesia, Indiana University School of Medicine, Indianapolis, IN, USA. <sup>2</sup>Department of Biostatistics and Health Data Science, Indiana University School of Medicine, Indianapolis, IN, USA. <sup>3</sup>Department of Anesthesiology, Nanjing Drum Tower Hospital, The Affiliated Hospital of Nanjing University Medical School, Nanjing, Jiangsu, China. <sup>4</sup>Department of Computer Graphic Technology, Polytechnic Institute, Purdue University, West Lafayette, IN, USA. <sup>5</sup>Department of Computer and Information Technology, Polytechnic Institute, Purdue University, West Lafayette, IN, USA. <sup>6</sup>Department of Computer Science, Luddy School of Informatics, Computing, and Engineering, Indiana University Bloomington, Bloomington, IN, USA. <sup>7</sup>Clinical Research Systems, Enterprise Analytics, Indiana University Health, Indianapolis, IN, USA. ✉e-mail: [su1@iu.edu](mailto:su1@iu.edu)

demands of precision hemodynamic care. For example, although BP is a critical target in septic patients<sup>10,11</sup>, a randomized trial found no significant mortality difference when comparing mean arterial pressure targets of 80–85 mmHg versus 65–70 mmHg in septic shock<sup>9</sup>. Our recent randomized controlled trial using individualized baseline tissue oxygenation “safe zone” in off-pump coronary artery bypass grafting also showed no significant reduction in 30-day postoperative complications<sup>12</sup>. These examples suggest the need for data-driven, personalized, and dynamic hemodynamic management in septic shock, cardiac surgery, and other acute care settings<sup>13</sup>.

To date, no approach provides data-driven, personalized, and adaptive HR and BP targets simultaneously, exposing a critical gap in hemodynamic management. As a result, clinicians rely on empirical judgment to tailor fixed targets, a practice that is anecdotal and inconsistent and contributes to variation in care quality, especially in under-resourced areas like Appalachia regions in the United States<sup>14</sup>. This underscores an urgent unmet need: real-time, individualized recommendations for HR and BP management targets.

The emergence of rich, multi-domain, real-world ICU big data, such as the Indiana University Health (IUH) dataset (6918 ICU stays, 11 hospitals, 2019–2024)<sup>15</sup>, the national eICU database (200,859 ICU stays, 335 units across 208 hospitals, 2014–2015)<sup>16</sup>, and the Medical Information Mart for Intensive Care (MIMIC) IV database (73,181 ICU stays, single healthcare system, 2008–2019)<sup>17</sup>, captures the heterogeneity and complexity of ICU patients and presents opportunities to derive transferable insights for hemodynamic optimization.

Still, real-time prediction of optimal HR and BP targets remains a major challenge<sup>18</sup>. Existing temporal models for ICU care include long short-term memory networks<sup>19–21</sup>, temporal convolutional networks<sup>22,23</sup>, recurrent neural networks<sup>24</sup>, other neural network architectures<sup>25,26</sup>, and regression-based approaches such as DYNAMIC-ICU<sup>27</sup>. Outside of ICU settings, models like the Continuous Individualized Risk Index<sup>28</sup>, based on Naïve Bayes, have also been developed. Ensemble models have proven effective in biomedical and clinical applications<sup>29–33</sup>, including in the ICU (e.g., ICU-ISPM<sup>34</sup> and stacked ensemble models<sup>35</sup>). However, these models fall short of addressing the unique complexities of personalized, real-time, and adaptive ICU care.

A defining feature of ICU care is the presence of rapidly evolving clinical conditions, which introduce unique challenges for setting hemodynamic targets. One such challenge is the dynamic shift in patient composition after ICU admission. Most patients are discharged or deceased within two days<sup>16,17,36</sup>, so the ICU population at different time points post-admission varies significantly (Supplementary Fig. 1) in terms of demographics, comorbidities, illness severity, and mortality risk. Thus, using a single model across all ICU time points may not accurately reflect these changing dynamics. Another challenge stems from the continuous evolution of a patient’s clinical state during their ICU course. Therefore, not only are demographics, comorbidities, and initial conditions highly variable at admission<sup>16,17</sup>, but patients’ real-time mortality risk also changes due to disease progression and therapeutic response. These fluctuations are evident in vital signs such as HR and BP, life support interventions like intubation, and measures of organ function such as the Sequential Organ Failure Assessment (SOFA) score. Moreover, outcomes like mortality often violate assumptions of parametric survival models, such as the Cox proportional hazard model<sup>37</sup>.

To address these challenges, we developed a novel modeling framework – Dynamic Cohort Ensemble Learning (DynaCEL) – for personalized, real-time ICU hemodynamic management. DynaCEL treats patients at the same time since the ICU admission as a “temporal cohort” and builds a distinct hemodynamic management model for each such cohort. A patient’s ICU trajectory is modeled as transitions through a series of these independent temporal cohorts, and their evolving hemodynamic needs are captured by aggregating models across temporal cohorts. DynaCEL is a flexible framework compatible with various base models, including artificial intelligence, machine learning, and statistical methods, and can be generalized beyond hemodynamic management to other personalized and dynamic clinical tasks. To support this approach, we first outline the rationale and design principles underlying the DynaCEL framework.

## Results

### DynaCEL: Rationale and framework

Predicting individualized hemodynamic management targets, such as heart rate (HR) and systolic blood pressure (SBP), in real time remains a major challenge in ICU practice. Existing temporal dynamic models—such as long short-term memory networks<sup>19–21</sup>, temporal convolutional neural networks<sup>22,23</sup>, recurrent neural networks<sup>24</sup>, other neural network models<sup>25,26</sup>, and regression-based approaches like DYNAMIC-ICU<sup>27</sup>, and ensemble frameworks<sup>34,35</sup>—have advanced outcome prediction for ICU patients. However, these models primarily forecast general risks (e.g., mortality or sepsis) rather than generating actionable, dynamic hemodynamic targets tailored to the patient’s evolving physiological state.

A key limitation is the profound temporal heterogeneity in ICU populations. Most ICU patients are either discharged or deceased within two days of admission<sup>16,17,36</sup>, leading to rapid shifts in demographic composition, comorbidity burden, illness severity, and outcome risks over time. Single-model strategies, which assume relatively stable cohort characteristics, may fail to capture this dynamic evolution. Moreover, real-time changes in patient status, reflected in dynamic variables such as HR, BP, respiratory support requirements, and SOFA scores, require modeling approaches that can accommodate instantaneous physiological changes. Standard survival models, such as Cox regression, also impose assumptions like proportional hazards<sup>37</sup> that may not hold in this volatile clinical context.

Recent availability of large, diverse ICU datasets—including the eICU Collaborative Research Database<sup>16</sup>, MIMIC-IV<sup>17</sup>, and IUH dataset<sup>15</sup>—provides new opportunities to learn patient-specific hemodynamic management strategies from real-world data. However, leveraging these resources effectively demands approaches that account for rapid physiological shifts and changes in cohort characteristics.

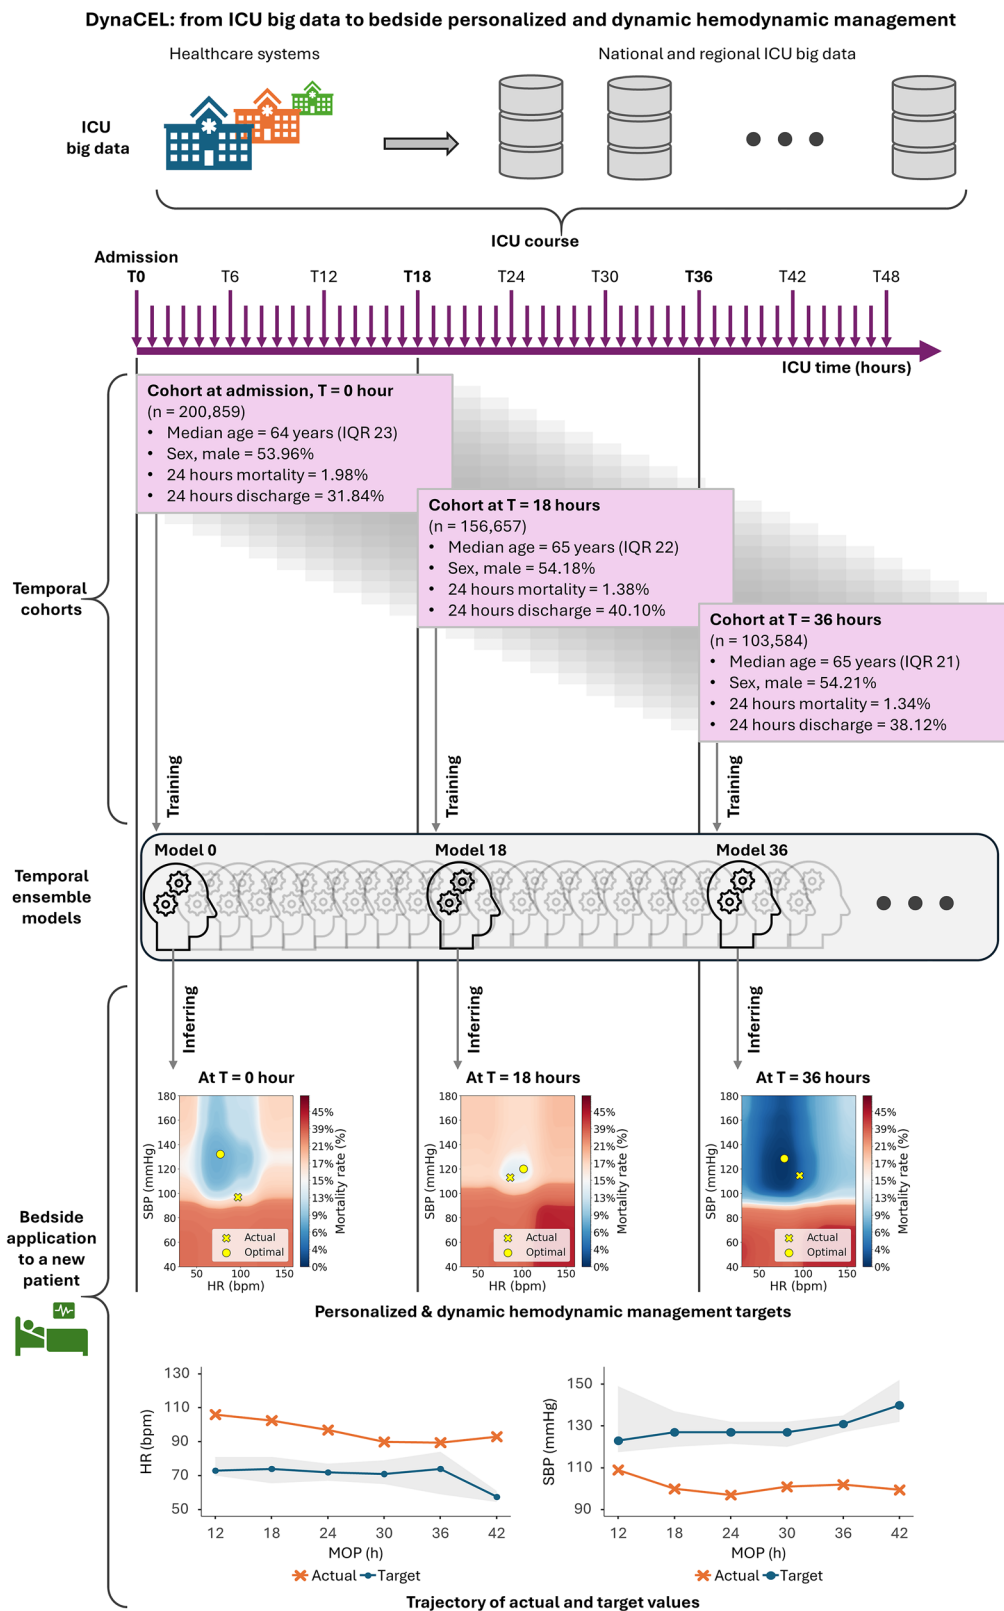
To address these challenges, we developed DynaCEL. DynaCEL temporally decomposes the overall heterogeneous modeling problem into multiple homogeneous subproblems. Specifically, it treats each moment after ICU admission—known as a moment of prediction (MOP)—as a distinct temporal cohort<sup>38</sup>. Separate base models are trained for each MOP using demographic characteristics, baseline clinical information, and dynamic variables within a designated predictor window (PW). These models estimate acute mortality risks within a defined outcome window (OW) without requiring assumptions about the underlying hazard function. The individually trained models are then temporally assembled into an ensemble that defines personalized HR and SBP targets associated with the lowest predicted mortality risk at any given time.

By modeling dynamic cohorts independently and aggregating them longitudinally, DynaCEL addresses the shifting trajectories of patient physiology, illness progression, and treatment response in the ICU setting. This approach leverages the depth of ICU big data to maintain sufficient sample sizes for training each sub-model, even after temporal partitioning. The framework is designed to be extensible, allowing new MOPs, PWs, and OWs to be incorporated flexibly without retraining the entire model system. By enabling real-time, precision hemodynamic management tailored to individual patients and time points, DynaCEL aims to address critical challenges in ICU care and lays the foundation for broader applications in dynamic clinical prediction.

### A temporal ensemble learning framework for real-time, personalized hemodynamic management

The DynaCEL model is a big data-driven framework for defining personalized, dynamic hemodynamic targets—specifically, HR and BP safe zones—that better address individual patient needs and reduce health disparities across health systems<sup>13</sup>. It infers real-time, patient-specific HR-BP targets and visualizes associated safe and risk zones as contour-based mortality risk maps, enabling precise computation of HR/BP values linked to the lowest mortality risk.

The framework architecture is illustrated in Fig. 1. DynaCEL translates ICU big data into bedside decision support for real-time, personalized hemodynamic management. Using national-scale datasets such as eICU,



**Fig. 1 | The DynaCEL framework: translating ICU big data into personalized, dynamic bedside hemodynamic management.** DynaCEL identifies real-time, individualized hemodynamic targets associated with the lowest acute mortality risk. It leverages national ICU big data (e.g., eICU) to construct temporal cohorts representing patient populations at specific post-admission time points (e.g., MOP = 0, 18, 36 h), each with distinct clinical characteristics and mortality profiles. For each cohort, a separate model is trained to learn HR and SBP targets linked to 24-hour mortality risk. These models form a temporal ensemble, enabling bedside deployment in external hospitals

(e.g., Indiana University Health, Beth Israel Deaconess Medical Center). For a new patient, DynaCEL recommends HR/SBP targets based on current clinical data, visualizes optimal targets on a mortality contour map, and displays the patient's actual vs. target trajectory. It also provides alerts for deviations from recommended targets. DynaCEL is plug-and-play, supporting real-time, interpretable, and personalized hemodynamic management at the bedside. DynaCEL Dynamic Cohort Ensemble Learning, ICU intensive care unit, HR heart rate, SBP systolic blood pressure, MOP moment of prediction (time since ICU admission).

DynaCEL applies a temporal ensemble learning strategy to recommend individualized HR and SBP targets. Temporal cohorts were first constructed based on post-admission time, each representing distinct patient characteristics and mortality risks. A separate model was trained for each cohort to capture the relationship between HR, SBP, and acute mortality. These temporally stratified models were then integrated into an ensemble ready for bedside deployment in external settings (e.g., IU Health and Beth Israel Deaconess Medical Center).

In real-time, DynaCEL recommends HR/SBP target combinations and projects them onto patient-specific mortality contour maps (Fig. 2). It simultaneously visualizes the patient's actual HR and SBP, highlighting deviations from target values and alerting clinicians to high-risk deviations. DynaCEL's plug-and-play architecture allows implementation without site-specific retraining, facilitating equitable care delivery across hospitals.

As detailed in the Methods, DynaCEL is trained to minimize the risk of acute mortality, defined as death within a 24-hour period (i.e., OW), starting from a given MOP. The MOP was defined as the time after ICU admission. The model incorporates demographic and admission data, as well as dynamic clinical features from the 12-hour period prior to the MOP (i.e., PW), to infer an acute mortality risk map over a clinically relevant HR/SBP space. It then identifies the HR and SBP pair associated with the lowest predicted mortality and defines "safe" and "risk" zones accordingly.

To validate the DynaCEL targets, we compared outcomes in patients whose HR/SBP values were near versus distant from the model-recommended targets and contrasted personalized versus population-based targets in propensity score-matched analyses (Figs. 3–5). Model performance was evaluated using its ability to predict acute mortality in IUH and MIMIC cohorts across various MOPs and hyperparameter settings (Figs. 6–9).

DynaCEL's generalizability across modeling approaches was demonstrated using multilayer perceptron (AI), random forest and support vector machine (machine learning), and logistic regression (statistical models) (Figs. 6a and 7). Compared to single-model approaches that included time since admission as a covariate, DynaCEL's cohort-specific temporal models showed significantly improved performance (Figs. 6b and 8). Models trained on the national eICU dataset generalized well to the IUH and the MIMIC IV populations (Figs. 6c and 9). Model fairness was assessed across subpopulations (Fig. 6d and 6e), and versatility was shown via consistent performance across different subpopulations. Two ICU case studies further illustrate real-world applications.

By shifting from fixed, population-based targets to data-driven, patient-specific recommendations, DynaCEL offers a scalable solution for precision hemodynamic management. It is ready for bedside deployment and has the potential to improve outcomes by delivering personalized, dynamic decision support in the ICU.

### Dose-response relationships highlight the sensitivity of DynaCEL targets

Deviation analysis provides a key metric for evaluating clinical management targets, using real-world observations to quantitatively assess the dose-response relationship between deviations from targets and negative outcomes. In this study, we hypothesized that greater deviation from model-recommended hemodynamic targets correlates with higher 24-hour mortality. Deviations from model-recommended HR or SBP values were treated as the "dose," and the observed 24-hour mortality as the "response." We used this framework to compare DynaCEL's personalized, dynamic targets with fixed, population-based targets (HR of 80 bpm and SBP of 120 mmHg).

As shown in Fig. 3, we used the MIMIC-IV cohort at 18 h post-admission (MOP = 18 h,  $n = 63,310$  ICU stays) as an example, dividing it into subgroups based on deviation levels from DynaCEL or fixed targets (Fig. 3a), ranging from within 20% to beyond 50%. Figures 3b and 3c present subgroup distributions.

Mortality rates (Fig. 3d and e) and odds ratios (Fig. 3f and g) relative to the within-target subgroup (deviation <20%) indicated that greater deviations were consistently associated with higher mortality. While both target

strategies showed this trend, the dose-response gradient was more pronounced under DynaCEL's personalized targets. Specifically, mortality rose from 0.04% (within 20%) to 18.8% (beyond 50%) under DynaCEL (Fig. 3d), compared with 0.8% (within 20%) to 7.3% (beyond 50%) under fixed targets (Fig. 3e). Corresponding odds ratios increased from 1 (within 20%, reference) to 421.6 (beyond 50%) for DynaCEL (Fig. 3f), and from 1 (within 20%, reference) to 6.5 (beyond 50%) for population-based targets (Fig. 3g).

To further clarify, each patient was assigned to exactly one mutually exclusive subgroup based on their degree of deviation from the personalized or fixed targets. Mortality rates were lower for patients closer to either set of targets, particularly within 20–30% deviations. As deviation increased, mortality rose progressively in both groups. Importantly, the dose-response gradient was steeper with DynaCEL-predicted targets than with fixed targets, highlighting the stronger sensitivity of personalized targets in discriminating mortality risk. The apparent "crossing" of mortality rates beyond 30% deviation does not imply that fixed targets are superior, but rather reflects the greater sensitivity of personalized targets to deviations from the optimal range.

These results suggest that DynaCEL's personalized, dynamic targets are more strongly associated with mortality risk than static benchmarks, suggesting potential utility in guiding future clinical trials. DynaCEL's improved performance may be attributed to its ability to learn transferable, generalizable knowledge from successful clinician-led deviations in the eICU dataset—instances in which providers intentionally adjusted standard guidelines based on clinical judgment and patient-specific needs—and effectively apply these insights to patients at Beth Israel Deaconess Medical Center.

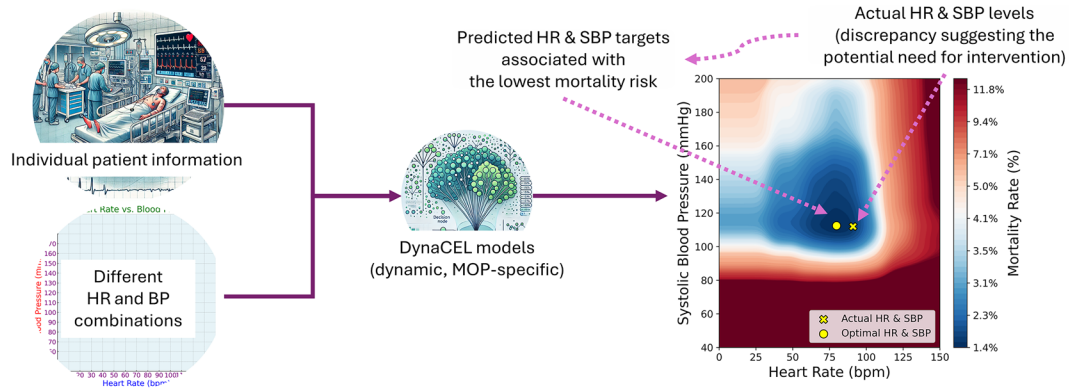
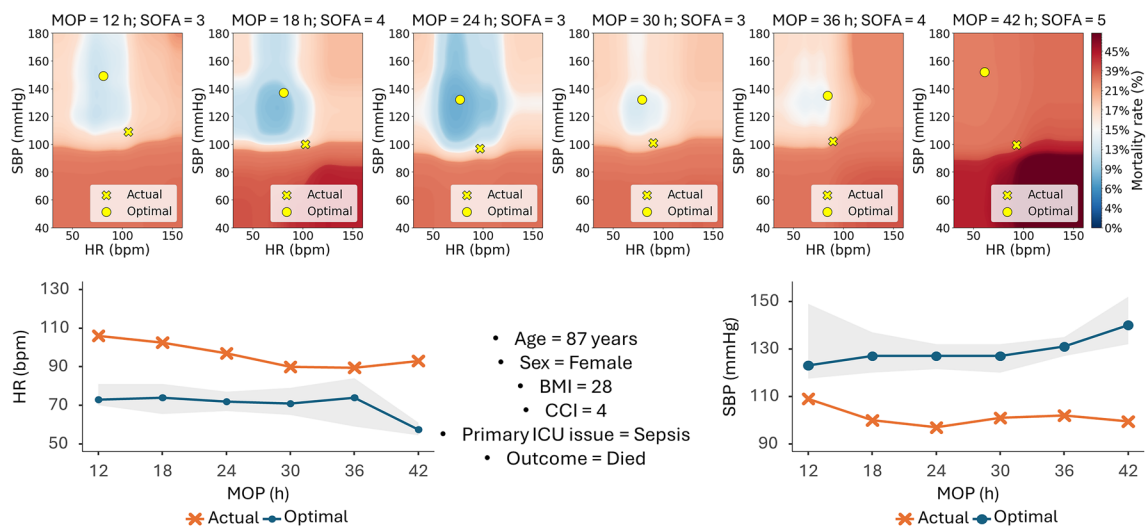
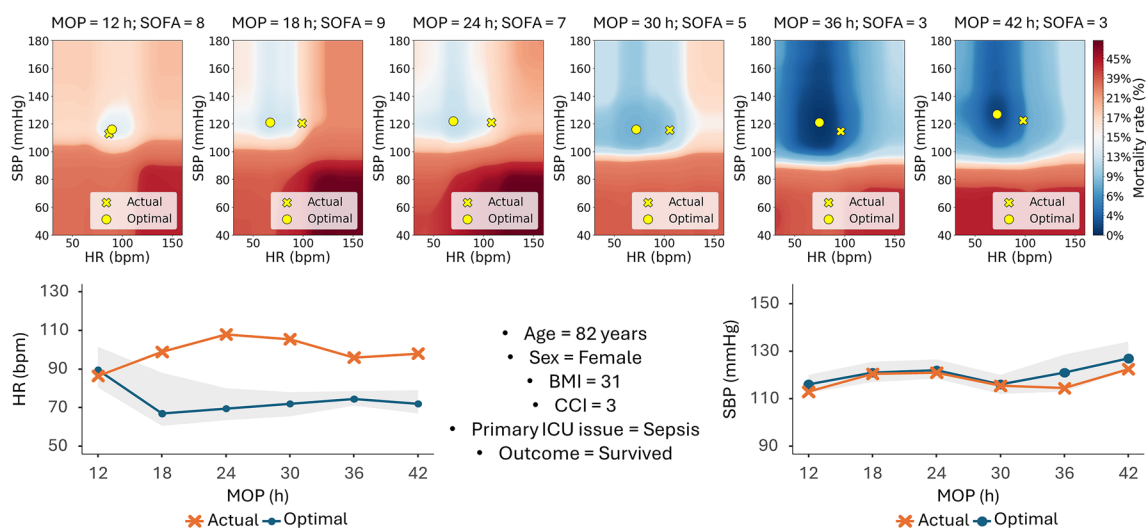
### DynaCEL targets have the potential to reduce mortality per deviation-stratified propensity-matched comparative effectiveness analysis

Some deviations from the hemodynamic targets occur as trade-offs for competing therapeutic goals, such as avoiding excessive vasopressor-related side effects or resulting from clinical conditions like sepsis, both of which are associated with mortality risks. To reduce confounding effects, we performed a comparative effectiveness analysis using propensity score matching for different subgroups with different deviations. Briefly, the temporal cohort representing patients who had been admitted to the ICU for 18 h (MOP = 18 hr,  $n = 63,310$  ICU stays) was labeled as "within targets" if HR and SBP were within 20% of the recommended values and "beyond targets" otherwise (Fig. 4a, Supplementary Tables 1 and 2). Propensity score matching was then used to construct matched subgroups, and odds ratios for 24-hour acute mortality were computed.

When "beyond targets" was defined as either HR or SBP deviating by more than 20% (Fig. 4b–d), the odds ratio for 24-hour mortality was 143.3 using DynaCEL's personalized dynamic targets, versus 2.9 for fixed population-based targets. If "beyond targets" required both HR and SBP to deviate by more than 20% (Fig. 4e–g), odds ratios increased to 587.0 for DynaCEL and 5.3 for the fixed targets. The corresponding 24-hour mortality rates under DynaCEL were 0.03% ("within targets") vs. 4.2% ("beyond targets") for the "HR or SBP" criterion, and 0.05% vs. 22.1% for the "HR and SBP" criterion. For the fixed targets, mortality was 0.9% vs. 2.5% ("HR or SBP"), and 1.0% vs. 5.3% ("HR and SBP"), respectively.

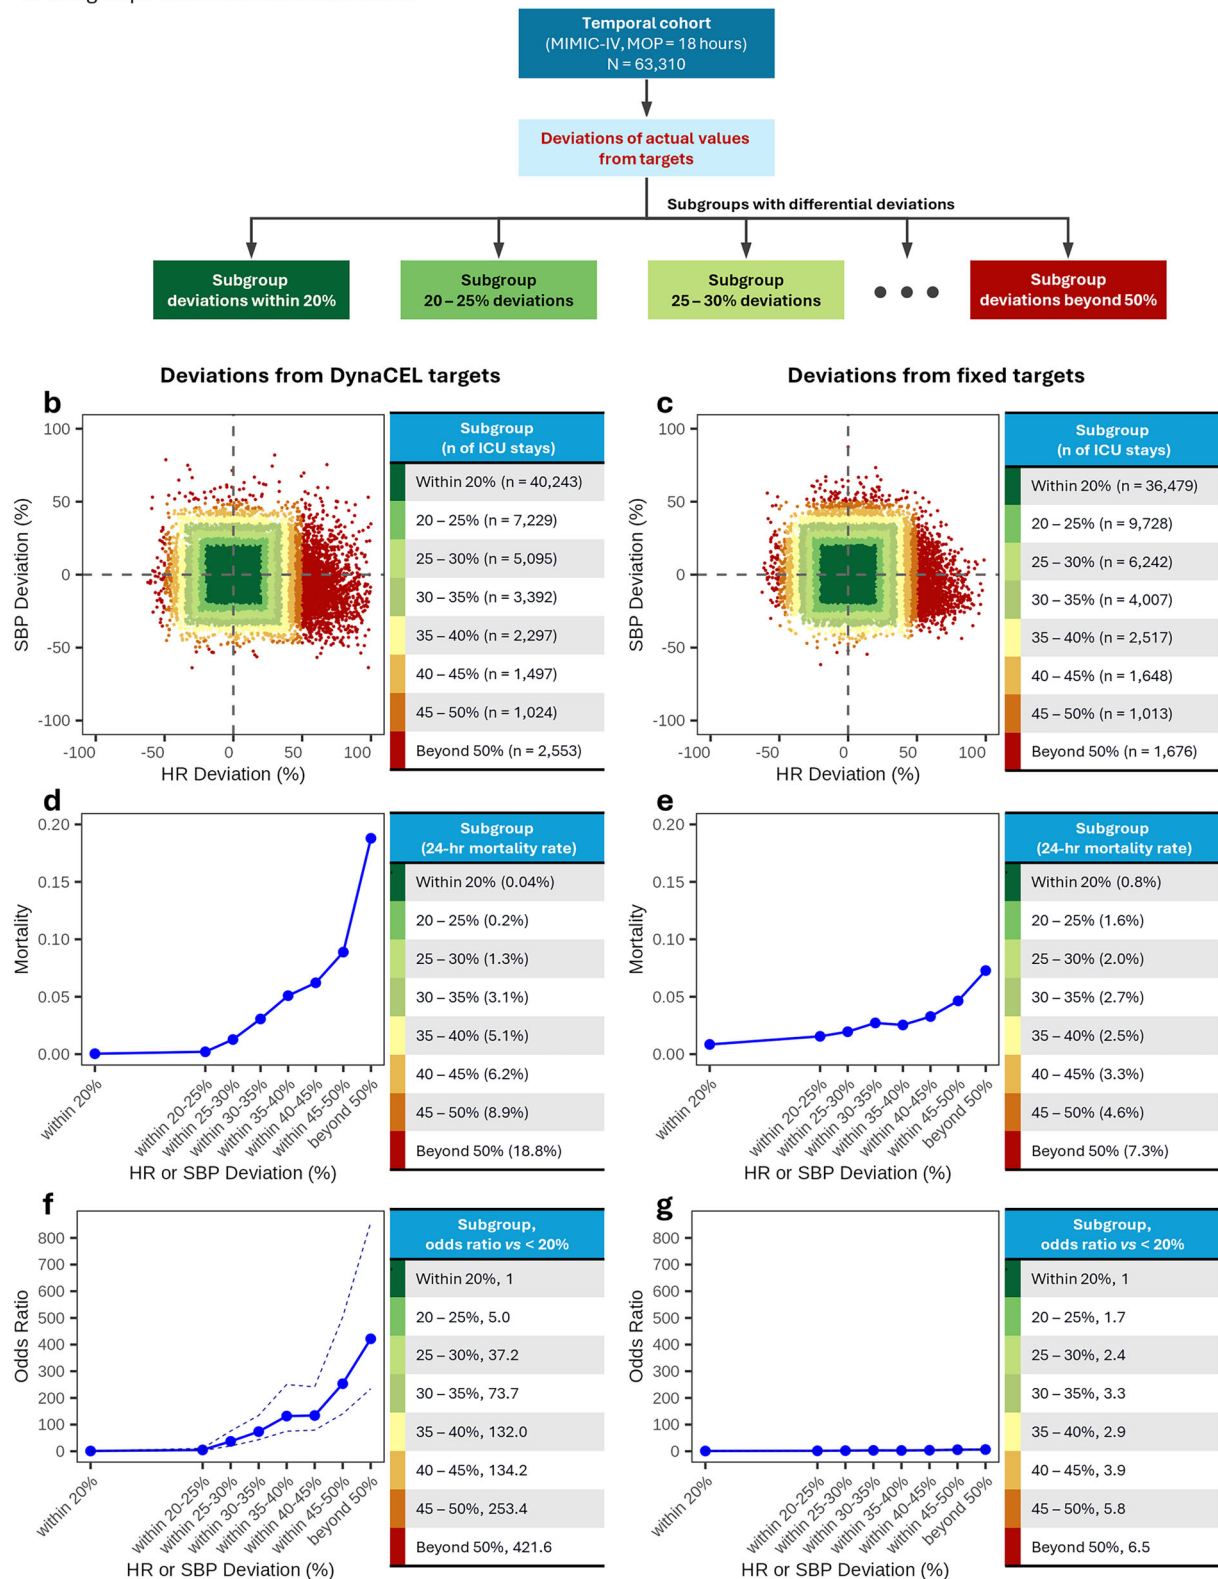
The comparative effectiveness results suggest that DynaCEL target-guided hemodynamic management is associated with lower observed mortality and may warrant further investigation in prospective interventional studies, based on the direct comparisons between subgroups close to and far away from these targets, after adjusting for observable confounding factors. Moreover, DynaCEL target-guided care appears to reduce mortality risks more significantly than fixed population-based targets. However, the final validation depends on future randomized controlled trials in which mortality outcomes are compared between groups whose hemodynamic care is guided by DynaCEL targets and those guided by usual targets, respectively.



**a** Bedside application: real-time recommendation of personalized HR and BP targets**b** Series HR-BP-mortality maps over time in a patient who died at 68 hours**c** Series HR-BP-mortality maps over time in a patient who was discharged at 69 hours

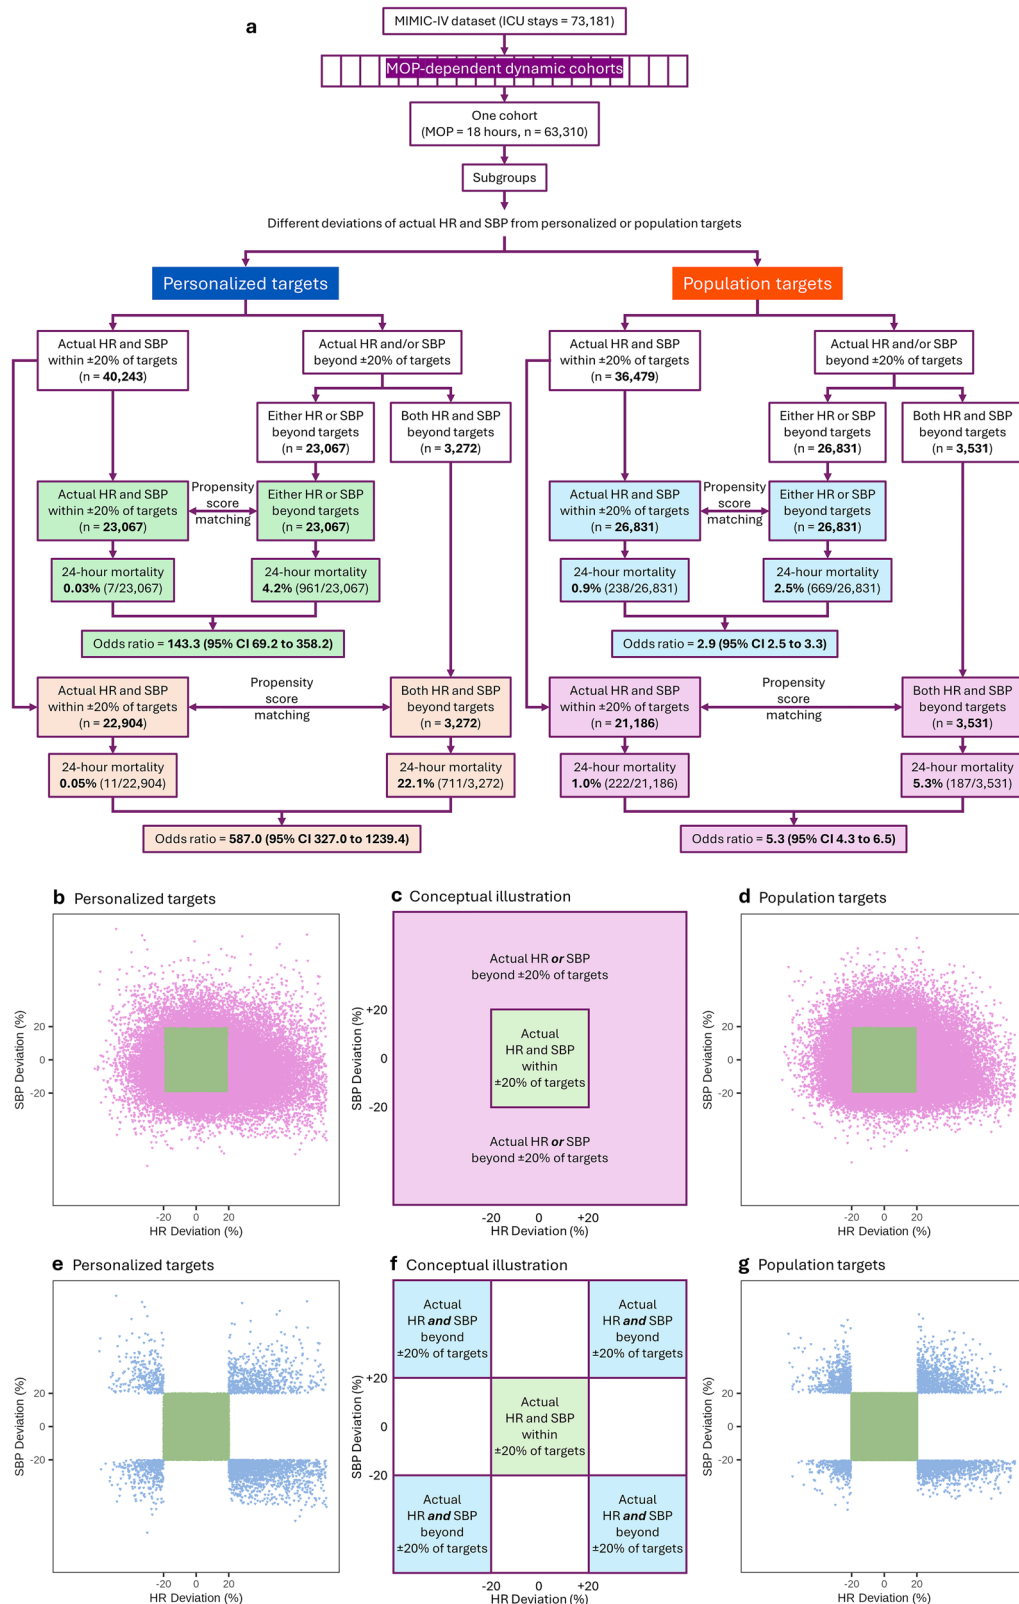
**Fig. 2 | Case studies illustrating DynaCEL's real-time visualization and clinical interpretability.** Panel **a** shows the bedside implementation of DynaCEL and its real-time visualization of recommended HR/SBP targets using HR-SBP-mortality contour maps. These maps were applied to two elderly female patients with sepsis to demonstrate the model's interpretability over time. In patient **b**, who died 68 h after ICU admission, mortality risk progressively increased alongside sustained deviations between actual and predicted SBP. In contrast, patient **c**, who was discharged after 69 h, showed a decline in mortality risk with closer alignment between actual

and predicted SBP values. These examples reveal distinct hemodynamic patterns, highlight discrepancies between actual values and model-predicted targets, and underscore the importance of personalized hemodynamic targets. (Artificial intelligence (ChatGPT image generator) was used to create the vignettes in **a**). DynaCEL Dynamic Cohort Ensemble Learning, HR heart rate, SBP systolic blood pressure, MOP moment of prediction, ICU intensive care unit, SOFA Sequential Organ Failure Assessment, BMI body mass index, CCI Charlson Comorbidity Index.

**a** Subgroups with differential deviations

**Fig. 3 | Dose-response analysis of deviation from HR/SBP targets and 24-hour mortality risk.** Using the 18-hour MOP cohort from the MIMIC-IV dataset, patients were divided into non-overlapping subgroups based on the degree of deviation between actual HR/SBP values and two target strategies: personalized targets predicted by DynaCEL and fixed population-based targets (HR = 80 bpm, SBP = 120 mmHg). **a** shows the stratification strategy based on deviation levels. **b** and **c** depict subgroup distributions by deviation from DynaCEL and fixed targets,

respectively. Panels **d** and **e** present the corresponding 24-hour mortality incidence, while **f** and **g** display odds ratios using the <20% deviation subgroup as the reference. This analysis quantifies the dose-response relationship between target deviation and mortality, highlighting a stronger and more discriminative risk gradient for personalized targets compared to population-based thresholds. DynaCEL Dynamic Cohort Ensemble Learning, HR heart rate, SBP systolic blood pressure, MOP moment of prediction, ICU intensive care unit.



**Fig. 4 | Propensity score-matched comparison of 24-hour mortality based on proximity to HR/SBP targets.** Propensity score-matched analyses were performed on the 18-hour MOP cohort from the MIMIC-IV dataset to assess mortality differences between subgroups whose actual HR and SBP were within or beyond 20% of target values. Two target sets were evaluated: personalized targets predicted by DynaCEL and fixed population-based targets (HR = 80 bpm, SBP = 120 mmHg). Panel **a** outlines the subgrouping strategy. Subgroups were stratified into those with

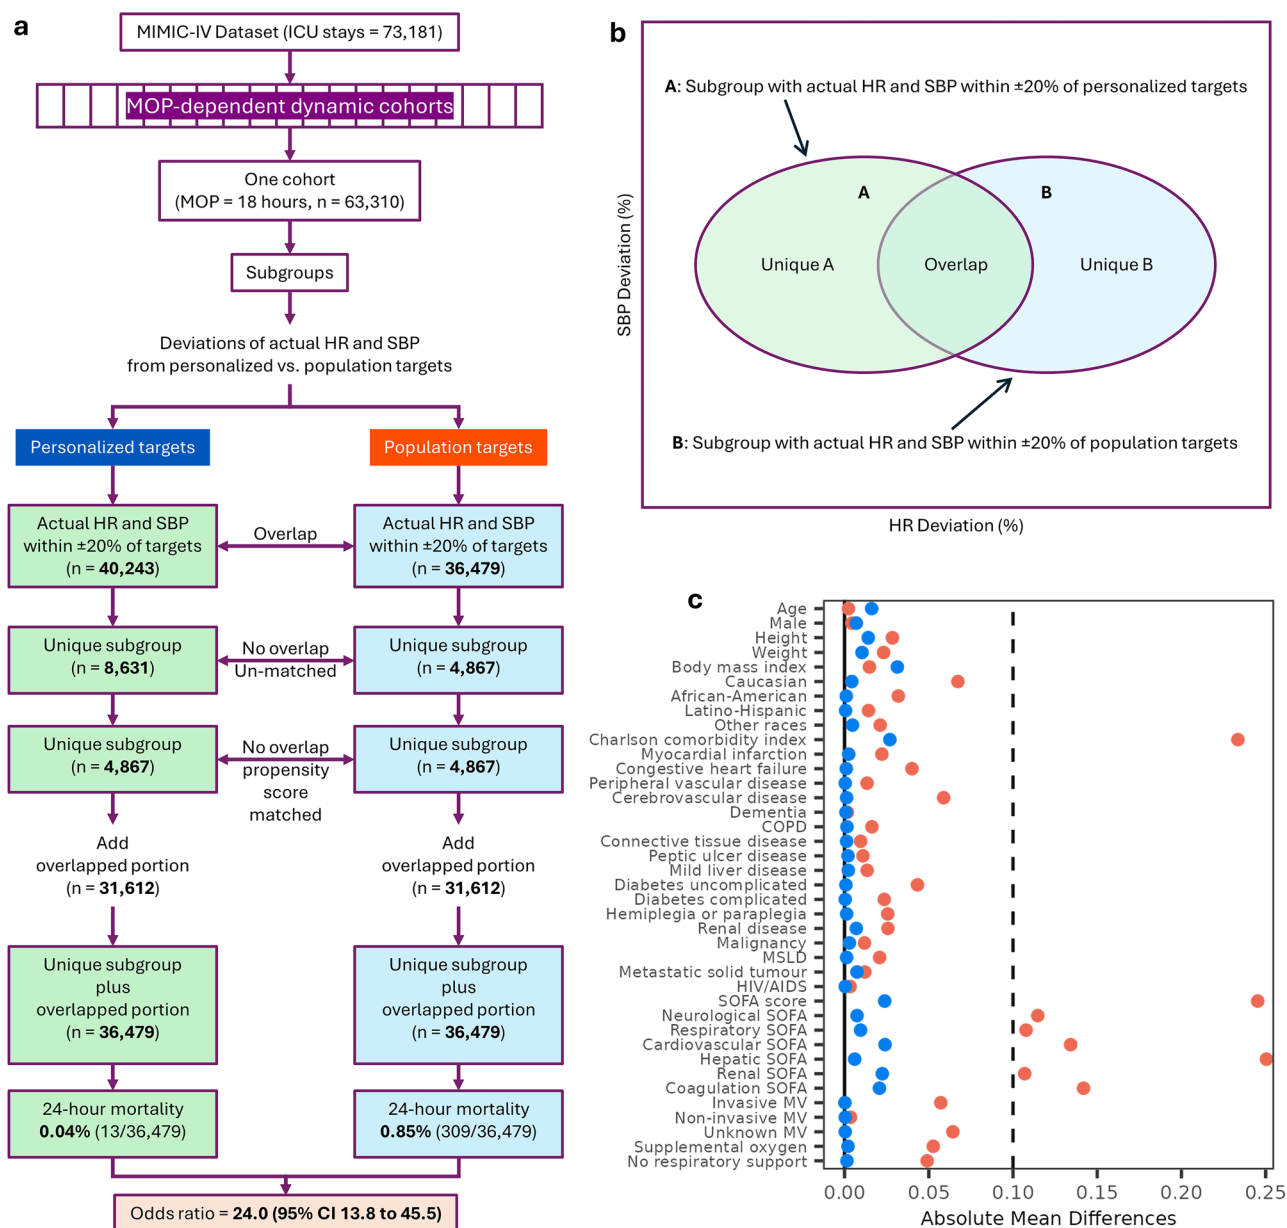
either HR or SBP beyond 20% of targets (**b–d**), and those with both HR and SBP beyond 20% (**e–g**). Panels show distributions of deviations from DynaCEL-predicted targets (**b, e**) and from population-based targets (**d, g**). DynaCEL Dynamic Cohort Ensemble Learning, HR heart rate, SBP systolic blood pressure, MOP moment of prediction, MIMIC-IV Medical Information Mart for Intensive Care, ICU intensive care unit.

## Direct comparison of DynaCEL and population-based targets in propensity-matched comparative effectiveness analysis

The effectiveness of DynaCEL and population-based targets was directly compared for the “within targets” subgroups using propensity score matching (Fig. 5). The rationale was that the above deviation-stratified, propensity-matched comparative effectiveness analyses were not head-to-head direct comparisons between DynaCEL and fixed targets (Fig. 4); therefore, differences in the study populations could introduce bias into the observed superior performance of DynaCEL targets. Another consideration is that HR and SBP values in some ICU cases within the “beyond targets” subgroup might be less modifiable. For example, some patients were less

responsive to vasopressor treatments due to their clinical conditions. To address these potential confounders, we generated matched “within targets” cohorts between DynaCEL and population-based targets (Fig. 5a), assuming that patients’ HR and SBP were modestly modifiable.

The two “within targets” subgroups partially overlapped (Fig. 5b). We performed propensity score matching on the non-overlapping portions (Fig. 5a and b, the Unique A and Unique B portions) and assumed self-matching for the overlapping segments. We observed a 24-hour mortality rate of 0.04% in the subgroup within 20% of DynaCEL personalized targets, compared to 0.85% in the subgroup within 20% of population targets. The odds ratio between DynaCEL and the population-based targets was 24.0,



**Fig. 5 | Propensity score-matched analysis of subgroups with actual HR/SBP within 20% of personalized vs. population targets.** This analysis, based on the 18-hour MOP cohort from the MIMIC-IV dataset, compares two subgroups: patients whose actual HR and SBP values were within 20% of DynaCEL-predicted personalized targets and those within 20% of fixed population-based targets (HR = 80 bpm, SBP = 120 mmHg). **a** Illustrates the subgrouping strategy and highlights the overlapping cases, which were excluded from propensity score matching but retained in the mortality analysis. **b** Shows the distribution of overlap between the two subgroups, and **c** presents absolute mean differences in clinical characteristics

before and after matching to assess covariate balance. Matching was performed using 1:1 nearest-neighbor matching between non-overlapping individuals, while overlapping cases were treated as self-matched. ICU intensive care unit, DynaCEL Dynamic Cohort Ensemble Learning, HR heart rate, SBP systolic blood pressure, MOP moment of prediction, MIMIC-IV Medical Information Mart for Intensive Care, COPD chronic obstructive pulmonary disease, MSLD moderate or severe liver disease, HIV/AIDS human immunodeficiency virus/acquired immunodeficiency syndrome, SOFA Sequential Organ Failure Assessment, MV mechanical ventilation.



indicating a 95% lower mortality associated with DynaCEL personalized targets (Fig. 5a and Supplementary Table 3).

The quality of the propensity score matching was evaluated using post-matching absolute mean differences of demographic and clinical features (blue dots, Fig. 5c). Before matching, the two “within targets” cohorts differed significantly in the Charlson comorbidity index and SOFA scores (red dots, Fig. 5c). After matching, all absolute mean differences were under 0.05, suggesting high matching quality.

This comparative effectiveness analysis suggests that among patients with modifiable HR and SBP, those within DynaCEL’s recommended ranges had a more than 20-fold lower observed 24-hour mortality compared with the currently widely used population-based fixed targets. Once again, the final validation depends on future randomized controlled trials, as causality cannot be inferred from these retrospective analyses.

### Applying DynaCEL at the bedside: two illustrative ICU cases

We conducted case studies to demonstrate the practical application of our DynaCEL model (Fig. 2a). Using HR-SBP-mortality contour maps, we visualized patient-specific relationships between HR, SBP, and mortality risk at different ICU time points (e.g., MOP = 12, 18, 24, 30, 36, and 42 h) (Fig. 2b and c). These maps, derived from patients with different demographics and baseline conditions (Supplementary Fig. 2) and varied critical illnesses (Supplementary Fig. 3), revealed unique patterns and highlighted discrepancies between actual values and predicted optimal targets.

Sequential maps for two elderly female patients with sepsis, who had different outcomes, illustrate the evolving nature of optimal HR/SBP targets and associated mortality risk (Fig. 2b and c and Supplementary Fig. 4). The patient who died showed an increasing mortality risk with SBP persistently below the predicted optimal range (actual 90–110 mmHg vs. target 120–140 mmHg) throughout her ICU stay (Fig. 2b). In contrast, the patient who was discharged from the ICU demonstrated a decreasing mortality risk, with SBP closely tracking the predicted targets (Fig. 2c). Notably, the deceased patient’s SBP was maintained around 100 mmHg, a level considered acceptable in current ICU sepsis care<sup>39</sup>. However, this value was below the DynaCEL-recommended optimal range, suggesting a potential need to reconsider therapeutic goals in such cases.

These two case studies demonstrate that the DynaCEL model offers real-time, intuitive visualizations of hemodynamic targets, safe and risk zones on the HR-SBP contour maps, deviations from recommended targets, and the trajectory of hemodynamic management for individual patients. The DynaCEL model is ready for bedside implementation and can support real-time clinical decision-making in ICU settings in future clinical trials.

### Feasibility assessment through ICU case review

To evaluate the clinical feasibility of aligning hemodynamic management with DynaCEL-predicted targets, we conducted a structured chart review of 20 representative ICU cases. These patients were randomly selected from the MIMIC-IV dataset and met the criteria of being alive and still in the ICU 18 h after admission (i.e., the 18-hour moment of prediction). To enhance the relevance of this analysis, we restricted the selection pool to cases in which either the model-predicted HR or SBP target deviated by at least 20% from the actual recorded values. This criterion enriched the sample for scenarios in which clinical practice diverged from model recommendations, thereby optimizing the utility of expert review for assessing the feasibility and safety of adopting DynaCEL targets.

For each case, we extracted demographic information, profiles of acute illness and chronic comorbidities, SOFA scores, and details of cardiovascular and respiratory support. Actual HR and SBP values were compared with the corresponding DynaCEL-predicted targets. Two independent ICU physicians, blinded to model development, reviewed each case and provided structured assessments of whether the recommended targets would have been clinically appropriate and safely achievable.

As summarized in Supplementary Table 18, the predicted targets were deemed consistent with individualized care strategies in the majority of cases. No safety concerns were identified based on the available clinical

context. In several instances, clinicians noted that the model-predicted targets may have offered helpful guidance for more personalized management. These findings provide preliminary support for the real-world applicability and safety of DynaCEL-recommended hemodynamic targets, warranting future prospective evaluation.

### Discussion

This study introduces a novel approach to personalized hemodynamic management in ICU patients. It leverages real-time, dynamic, data-driven HR and BP targets tailored to individual patient needs. Validation across the eICU, MIMIC-IV, and IUH datasets demonstrates the robustness and potential clinical applicability of this strategy. The DynaCEL models account for variability in patient demographics, comorbidities, and illness severity, delivering personalized hemodynamic targets that align with individual physiological states and are associated with improved outcomes in retrospective analyses.

Our study addresses the critical question of whether fine-tuning HR and BP management could be associated with improved outcomes without introducing additional harm. As ICU clinicians often prioritize treating underlying conditions—such as infection control in sepsis—rather than restoring “normative” vital sign values, we acknowledge that interventions solely targeting HR and BP must be interpreted with caution. While prior studies and our findings suggest associations between HR, BP, and mortality, achieving optimal hemodynamic targets should be considered complementary to, not a substitute for, definitive disease-specific therapies. The DynaCEL model does not imply that HR and BP management alone can overcome critical illness; rather, it reflects best practices embedded in large-scale ICU datasets to identify patient-specific hemodynamic ranges associated with lower mortality risk. Validation of model-guided targets through prospective randomized trials is necessary to determine whether aligning interventions to these targets improves patient outcomes.

The DynaCEL strategy harnesses large-scale ICU data and temporal ensemble learning to enable precise, dynamic hemodynamic management. Its adaptability across different time points, mortality windows, and patient subpopulations ensures both flexibility and robustness. DynaCEL supports seamless adjustments and is compatible with various machine learning models, positioning it as a versatile tool for precision medicine. Compatible with advanced AI and machine learning pipelines, DynaCEL offers a scalable framework for enhancing personalized care in ICU settings.

Hemodynamic management in critically ill patients remains complex due to dynamic physiological changes and the interplay of multiple factors. Although associations between HR, BP, and mortality have been established, existing approaches lack a real-time, personalized framework that adapts as clinical conditions evolve. Our findings highlight that optimal HR and BP targets vary with individual characteristics, clinical context, and ICU timing, underscoring the need for tailored approaches beyond static, population-based targets.

The DynaCEL models address this need by integrating HR and BP, recognizing their interdependence and the challenges of managing them both independently and jointly in clinical practice. BP interventions, such as vasopressors, often influence HR due to their interconnected regulation and shared responses to treatment. However, in clinical settings, elevated HR is less frequently a direct target of intervention than BP unless it reaches extremes that cause hemodynamic instability or myocardial ischemia. In many cases, HR serves as a marker of underlying physiological stress rather than a primary therapeutic target. Nonetheless, by incorporating both HR and BP, DynaCEL provides a comprehensive assessment of a patient’s hemodynamic state, potentially enhancing decision-making through a more nuanced and holistic perspective. At this stage of discovery, these data-driven targets can support clinical decision-making, but clinical judgment remains essential, particularly given the absence of trial-based evidence.

Systematic and rigorous validations of the DynaCEL targets were conducted to assess model performance across multiple dimensions. Propensity score-matched analyses demonstrated that patients with HR and BP within personalized target ranges were associated with lower mortality

compared to those within population-based targets. These findings reflect associations derived from observed clinical patterns, not causal effects; whether aligning interventions to these targets improves outcomes remains unknown and must be tested in prospective, randomized trials. Although our comparison was not based on a classical treatment assignment, we applied propensity score matching to adjust for baseline risk when comparing outcomes across patient groups stratified by physiologic states (e.g., deviations from predicted HR and SBP targets). This approach reflects the broader application of propensity score matching in observational prognosis-related modeling rather than treatment evaluation, which has gained traction in predictive analytics where no clear intervention is applied<sup>40</sup>.

While our study focused on identifying associations between HR/SBP and short-term mortality, we acknowledge that these parameters may act as either causal factors or bystanders influenced by underlying conditions or therapeutic interventions, such as vasopressors or beta-blockers. Distinguishing between causal and non-causal pathways remains a central challenge in critical care analytics. As highlighted by recent work on subphenotyping and causality in ICU patients<sup>41,42</sup>, observational models must be interpreted within the limitations of retrospective data and unmeasured confounding. Our findings should be understood as identifying predictive targets associated with lower observed mortality risk, rather than establishing definitive causal pathways. Future studies incorporating causal inference methods and prospective interventional designs will be crucial in determining whether achieving these personalized targets improves patient outcomes through causal mechanisms.

Practical implementation of DynaCEL will require integration with electronic health records and real-time monitoring systems, as well as clinician-facing visualization tools such as the HR-BP-mortality contour maps developed in this study. These maps translate model outputs into intuitive, actionable insights, while robust error-detection mechanisms ensure reliability in high-pressure ICU environments. Ongoing refinement through real-world data and clinical feedback is essential to maintain accuracy and clinical relevance across diverse settings.

Future research should explore the practical integration of this model into ICU workflows and assess its influence on patient outcomes. Prospective randomized trials are essential to determine whether maintaining personalized HR and BP targets leads to improved survival. Although this study focused on HR and SBP, many clinical protocols prioritize mean arterial pressure. Investigating personalized mean arterial pressure targets represents a key direction for broader clinical adoption. In addition, expanding model inputs to include laboratory data, imaging, and other clinical variables could enhance predictive performance. Evaluating the framework's utility in resource-limited settings will also be crucial in addressing global disparities in critical care delivery.

In addition, we acknowledge that ICU populations are heterogeneous, and the hemodynamic sensitivity and mortality risk associated with deviations from HR and BP targets likely vary across clinical subgroups. For instance, patients with septic shock may have a narrower tolerance to hypotension than those with chronic obstructive pulmonary disease or heart failure. While our current model incorporates demographic and clinical features and performs robustly across multiple subpopulations (Fig. 6d and e), future work should investigate whether subgroup-specific targets or models are necessary to optimize outcomes further. This may involve stratifying patients based on physiologic or diagnostic clusters, such as hyperinflammatory versus immunosuppressed phenotypes<sup>43</sup>, or leveraging causal inference techniques and stratified modeling to clarify whether observed differences reflect underlying biology or treatment effect heterogeneity. Such approaches are especially relevant in light of growing evidence that causal structures differ across ICU subpopulations and require tailored modeling strategies, as discussed by Zhang et al. using marginal structural models for longitudinal data<sup>42</sup>. This concept builds on the broader recognition of syndromic heterogeneity in critical illness, such as in sepsis<sup>44</sup>. Addressing this heterogeneity is essential to enhance the precision, fairness, and clinical utility of the DynaCEL framework.

This study highlights the potential of personalized hemodynamic management using machine learning, but several considerations must be acknowledged. First, while we used large and diverse ICU datasets, these primarily represent healthcare systems in the United States. Variations in clinical practice, patient demographics, and data availability in lower-resource environments may affect model performance and generalizability. Future studies should prioritize external validation across diverse healthcare settings, particularly in low- and middle-income countries.

Second, the reliability of our model depends on the quality and completeness of input data. Although the demographic and clinical features used are widely available in most ICU settings, and standardized imputation techniques were applied to address missing data, these methods cannot fully account for unmeasured confounders or capture the complexity of nuanced clinical situations. Additionally, as a retrospective study, data quality is inherently constrained by documentation practices in electronic health records, introducing potential biases. Prospective validation in controlled environments is necessary to ensure robustness and clinical applicability.

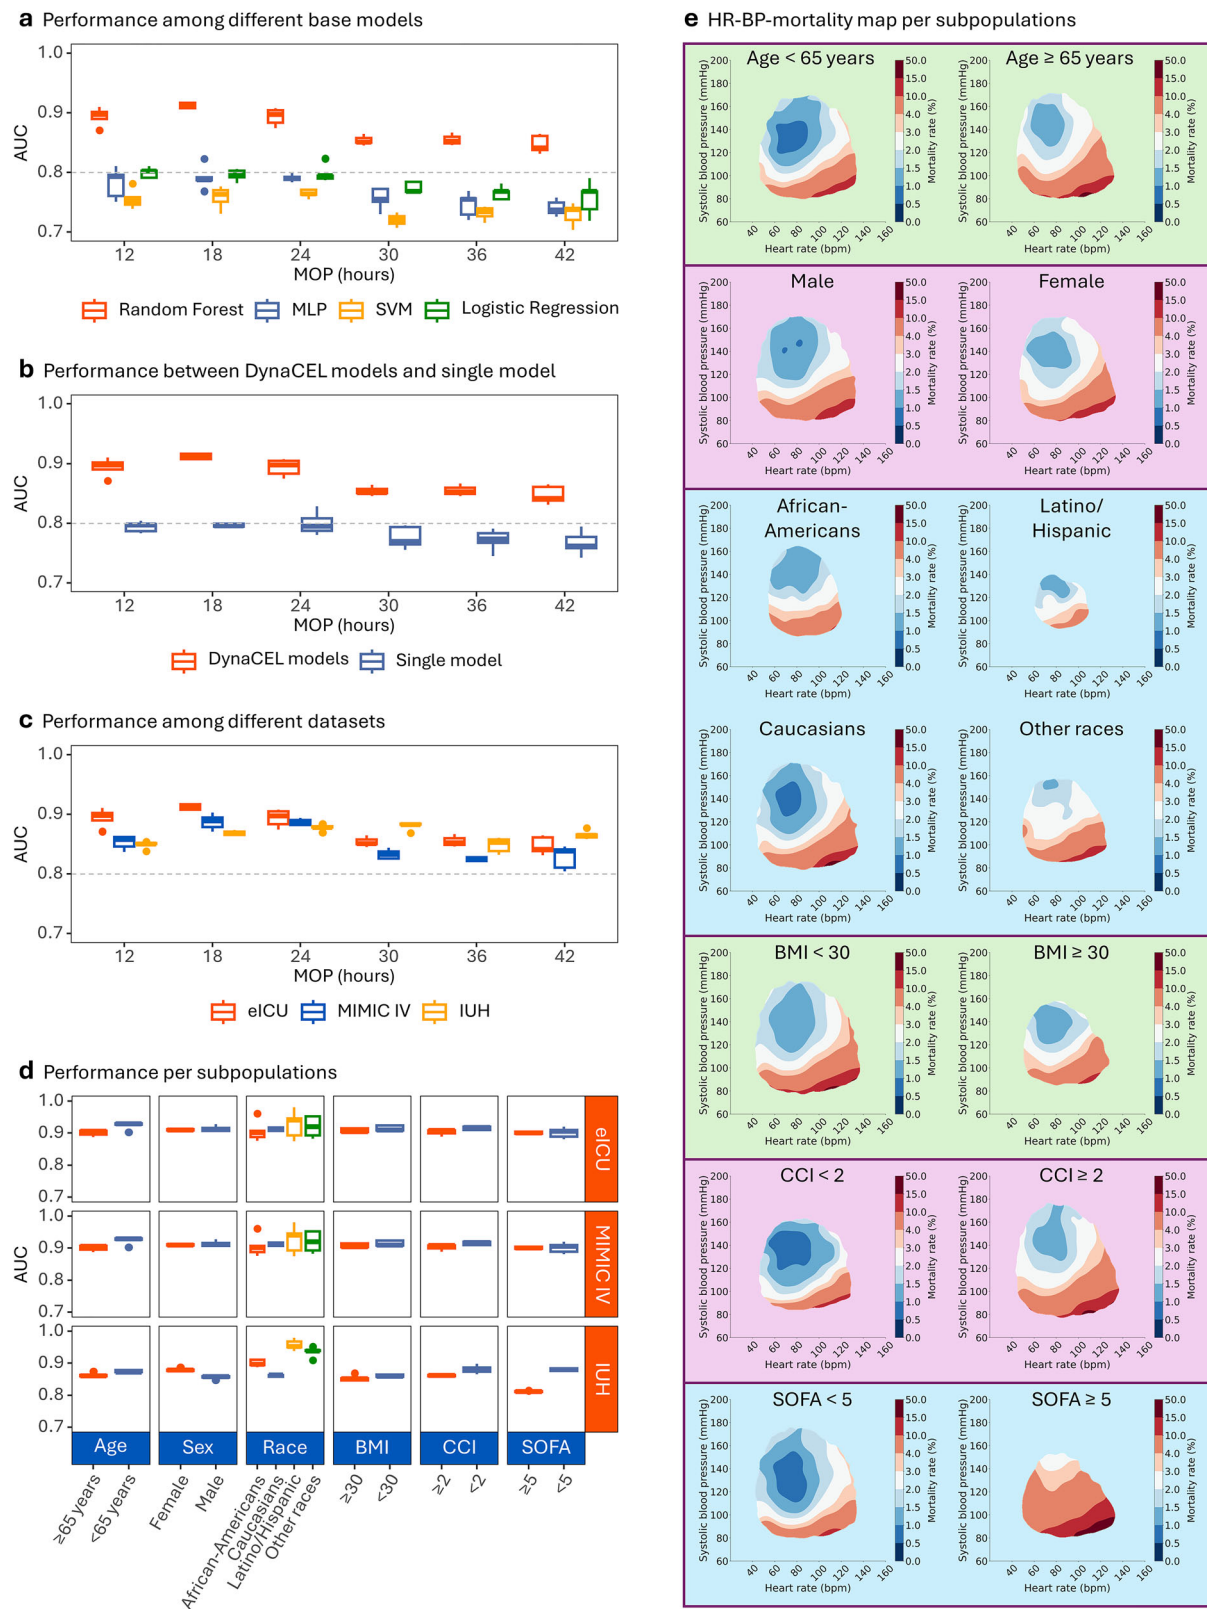
Third, although we implemented rigorous cross-validation and hyperparameter tuning to mitigate overfitting, and model performance was not sensitive to cohort characteristics or data structure, the model has already been externally validated in multiple hospitals through the IUH and MIMIC datasets, representing diverse regional features of ICU patients and clinical practices. Nevertheless, because machine learning models may exhibit reduced accuracy when applied to entirely new populations with distinct clinical patterns, local hospitals are encouraged to perform their own validation using institutional data prior to clinical implementation. Alternatively, hospitals may use our open-source algorithm to train a new DynaCEL model based on their own patient data, ensuring optimal performance and contextual relevance.

Furthermore, our analysis focused on HR and SBP, as these parameters are commonly monitored in ICU settings and readily available in real-time data streams. We recognize that other hemodynamic parameters, such as central venous pressure, cardiac output, and indicators of volume status, are clinically relevant and may enhance model performance. However, these variables were excluded due to their limited availability and inconsistent documentation across the datasets, with most patients lacking reliable measurements. Our goal was to develop a broadly applicable model using parameters that are routinely and consistently captured in ICU patients. Nonetheless, the DynaCEL framework is modular and extensible, allowing users to incorporate additional hemodynamic features when available to enhance predictions and adapt the model to local practice.

For institutions where other hemodynamic variables, such as mean arterial pressure, are prioritized, our framework offers flexibility. While we used SBP in this study, we acknowledge that mean arterial pressure is more commonly used in clinical practice to guide vasopressor therapy and assess perfusion. However, in large retrospective ICU datasets like MIMIC-IV and eICU, SBP is often recorded more frequently and consistently, whereas mean arterial pressure is sometimes derived from systolic and diastolic values and not always directly charted<sup>45</sup>. To maximize data completeness and applicability across diverse settings, we selected a single, directly measured blood pressure parameter for modeling. Nonetheless, our open-source algorithm allows users to substitute SBP with mean arterial pressure or other hemodynamic variables, supporting integration into institution-specific protocols while preserving the flexibility and robustness of the modeling framework.

Lastly, while our findings indicate an association between personalized HR and BP targets and lower mortality, this study does not establish causality. Whether adherence to model-inferred targets improves outcomes remains unknown and must be rigorously tested through prospective randomized controlled trials. Additionally, feasibility studies are needed to evaluate how this model can be effectively integrated into ICU workflows, ensuring that its recommendations are actionable and supportive rather than burdensome to clinical practice.

Despite these considerations, this study represents an important step toward data-driven, personalized hemodynamic management. Addressing



these challenges through prospective validation, broader population studies and real-world implementation research will be critical for translating this approach into clinical practice.

Taken together, this study presents a data-driven framework for personalized hemodynamic management, leveraging machine learning to deliver real-time, adaptive HR and BP targets tailored to individual patients.

By integrating patient-specific characteristics, this model offers a potential alternative to static, population-based targets, with the capacity to refine clinical decision-making. While findings suggest an association between personalized targets and lower mortality risk, additional validation remains essential. Prospective studies should assess feasibility, clinical adoption, and integration into ICU workflows, while randomized controlled trials are



**Fig. 6 | Summary of DynaCEL performance across models, datasets, and patient subpopulations.** This figure summarizes DynaCEL's predictive performance across several modeling configurations and validation settings. **a** compares four base learners—random forest, MLP, SVM, and logistic regression—using a 12-hour predictor window and 24-hour outcome window, applied to the eICU test set cohort at MOP = 18 hr. **b** compares DynaCEL's temporally stratified ensemble to a single pooled model trained across all MOPs with MOP included as a covariate; both used random forest and the same prediction window configuration as in panel a, evaluated on the same cohort. In **c**, external validation of DynaCEL is shown using the same random forest model applied to 18-hour MOP cohorts from the eICU test set, MIMIC-IV, and IUH datasets. **d** presents model performance across subpopulations

defined by age, sex, race/ethnicity, BMI, comorbidity burden, and SOFA score using 18-hour MOP cohorts from the same three datasets. **e** shows HR–SBP–mortality contour maps from the MIMIC-IV cohort at MOP = 18 hr, illustrating mortality risk variation across HR–SBP combinations by subpopulation. DynaCEL Dynamic Cohort Ensemble Learning, AUC area under the receiver operating characteristic curve, MOP moment of prediction, PW predictor window, OW outcome window, MLP multilayer perceptron, SVM support vector machine, eICU eICU Collaborative Research Database, MIMIC-IV Medical Information Mart for Intensive Care, IUH Indiana University Health, BMI body mass index, CCI Charlson Comorbidity Index, SOFA Sequential Organ Failure Assessment, HR heart rate, SBP systolic blood pressure.

needed to determine their impact on patient outcomes. Expanding model parameters, including mean arterial pressure, may further enhance its applicability in routine practice. With continued validation, this approach could advance precision in hemodynamic management, enabling more individualized and responsive ICU care.

## Methods

This study employs a machine learning-based approach to personalize HR and BP management targets for critically ill patients. Distinct from prior research, it advances beyond model development to demonstrate practical applications in personalized medicine and validate effectiveness in guiding individualized hemodynamic management. The study was approved by the Indiana University Institutional Review Board on 2 February 2024 (#22056), with patient consent waived due to its retrospective nature.

## Datasets

Three large ICU datasets were used: the eICU Collaborative Research Database (v2.0)<sup>16</sup>, the MIMIC-IV dataset (v2.2)<sup>17</sup>, and the IUH ICU dataset (v1.0)<sup>15</sup>. The eICU dataset includes 177,187 admissions from over 200 hospitals in the United States (2014–2015), offering a broad and diverse sample. The MIMIC-IV dataset (2008–2019) includes 73,181 ICU stays at Beth Israel Deaconess Medical Center (Boston, Massachusetts, the United States). The IUH dataset contains 4179 ICU admissions from 11 hospitals in Indiana, the United States (2022–2024), offering relevance to post-COVID-19 clinical practice. Supplementary Table 4 provides additional dataset details.

## Participants and ICU stays

Eligible participants were ICU patients aged  $\geq 18$  years with available HR and BP data, documented sex, and known survival status. As patients could have multiple ICU stays, this study used ICU stays—rather than unique individuals—as the unit of analysis, reflecting clinical practice where care is based on current presentation regardless of prior ICU admissions.

## Predictors

Predictors included a comprehensive range of demographic, baseline, comorbidity, and dynamic variables (Supplementary Table 5). Demographic and baseline data comprised age, sex, height, weight, body mass index, and race (categorized as Caucasian, African American, Latino/Hispanic, or other). Comorbidities were assessed using the Charlson Comorbidity Index, which assigns weighted scores based on 17 predefined medical conditions<sup>46,47</sup>. Dynamic variables included SOFA scores<sup>48</sup>, respiratory support modes, and HR and SBP measurements. For modeling, the SOFA score closest to the MOP, the most intensive respiratory support mode, and the median HR and SBP values within the prediction window were used.

## Outcomes

The outcome was acute mortality within a predefined short-term outcome window following each MOP (which indicates the time after ICU admission), representing real-time clinical decision-making. Short-term mortality was chosen because acute HR and BP changes are more strongly associated with short-term outcomes, and improvements in short-term survival are

likely to influence long-term outcomes. ICU mortality was used, defined as death occurring during the ICU stay. Mortality status was classified as deceased or survived based on discharge disposition and time of death, as recorded in ICU-specific tables within each dataset.

## Data processing

ICU stays missing critical data (i.e., age, HR/BP values, sex, and survival status) were excluded (Supplementary Fig. 5). Dynamic variables were computed hourly. Missing values were imputed as follows: height, weight, and BMI were imputed using the median; SOFA scores were forward-filled using the nearest previous value when necessary (Supplementary Table 6 and Supplementary Fig. 6). When respiratory support data were unavailable, the absence of support was assumed. Hourly missing rates of heart rate and systolic blood pressure were low (Supplementary Table 7). Since ICU stays without either HR or SBP (i.e., complete missingness) during the prediction window were rare (Supplementary Table 8), these stays were excluded from model training and validation. Therefore, for HR and SBP, missing hourly values within the prediction window were not imputed, as partial missingness did not affect the derivation of median values.

## Class imbalance

To address the imbalance between survival and mortality outcomes, the Synthetic Minority Over-sampling TEchnique (SMOTE) was applied to generate synthetic samples for the minority class, resulting in a balanced dataset for model training.

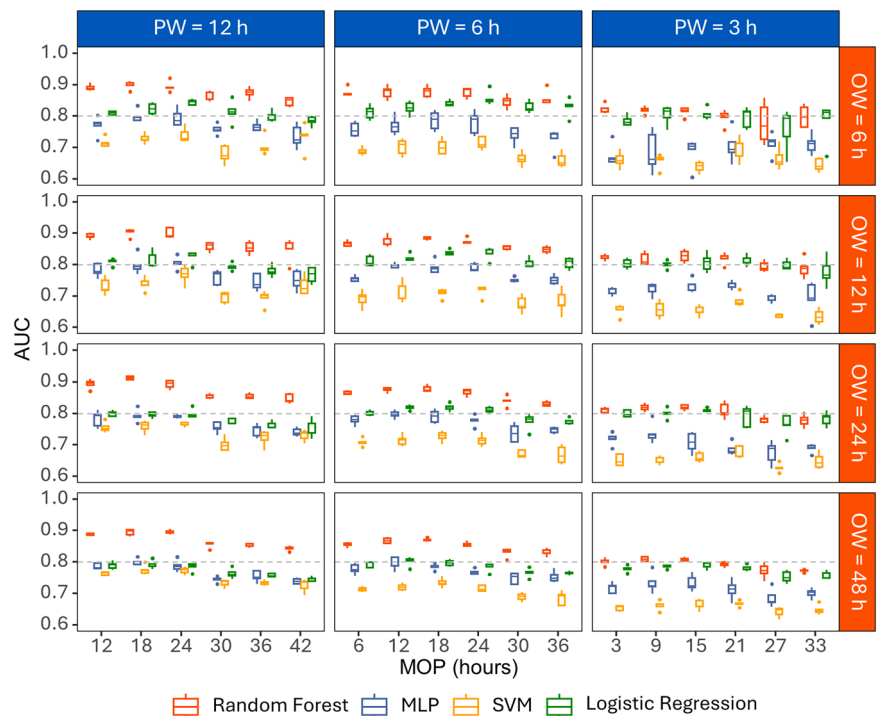
## Model development

The model was designed to identify personalized hemodynamic targets associated with the lowest mortality risk for individual ICU patients (Fig. 1). Model development followed a two-phase strategy. First, prediction models were constructed to establish associations between input variables (including HR and BP) and acute mortality (model development phase). These models were then used to infer individualized HR and BP targets associated with the lowest mortality risk (model application phase). In clinical practice, care providers could refer to these personalized targets to support real-time decision-making.

To enable practical application, ICU admission was standardized as time zero, with each subsequent hour defined as an MOP. At each MOP, the model generated recommendations for optimal HR and BP values for that specific point in time. For this study, we selected MOPs at 12, 18, 24, 30, 36, and 42 h for validation purposes based on the median ICU stay length; however, the DynaCEL framework can be applied flexibly to any time point as needed by clinical users. Model inputs included a combination of fixed variables (e.g., age) and dynamic variables (e.g., HR), with only pre-MOP dynamic data within the prediction window used to forecast outcomes in the outcome window. This structure allowed clinicians to base decisions on recent data, simulating real-time guidance in ICU settings. Prediction window lengths (e.g., 3, 6, or 12 h) were treated as hyperparameters and optimized accordingly. Various base models, including random forest, multilayer perceptron, support vector machine, and logistic regression, were evaluated, and the best-performing model was selected.



**Fig. 7 | Model performance across base models and prediction configurations.** This figure evaluates the effect of base model selection and prediction configuration on DynaCEL performance. The base model was treated as a hyperparameter, and AUC values were calculated across combinations of PW and OW using cohorts defined by MOPs in the eICU test set. Performance is shown for random forest, MLP, SVM, and logistic regression. Random forest consistently achieved the highest AUC across most configurations and was selected as the base model for subsequent analyses. AUC area under the receiver operating characteristic curve, MOP moment of prediction, PW predictor window, OW outcome window, MLP multilayer perceptron, SVM support vector machine, eICU eICU Collaborative Research Database.



ICU populations experience dynamic changes in cohort size due to patient discharges and deaths, which vary across MOPs (Supplementary Fig. 1). Additionally, patient conditions evolve over time. To address these changes, MOP-specific models were trained on corresponding MOP-specific cohorts using the DynaCEL framework. Further explanation of this approach is provided in the Introduction under ‘DynaCEL: Rationale and Framework.’

Each MOP-specific cohort included ICU stays where patients were alive and still hospitalized at the time of the MOP. Characteristics of the 18-hour MOP cohort from the eICU, MIMIC-IV, and IUH datasets are provided in Supplementary Table 9. To improve model sensitivity, we also evaluated expanded cohorts that included ICU stays ending in death (positive controls) or discharge (negative controls) during the prediction window preceding each MOP, and compared them with unexpanded cohorts that excluded these cases. The final model was selected by comparing the performance between models trained on expanded versus unexpanded cohorts. Models trained on expanded cohorts, which integrated ongoing ICU stays with those concluding during the prediction window, consistently outperformed those trained on unexpanded cohorts (Supplementary Table 10). Accordingly, all models in this study were trained on expanded cohorts.

### Model training, testing, and validation

For the validation within the eICU dataset, a 5-fold cross-validation design was used. The model was trained on four folds and tested on the remaining fold, with this process repeated five times so that each fold served as the test set once (Supplementary Fig. 7). The resulting models were validated on the MIMIC-IV and IUH datasets to assess their generalizability across populations and time periods. Specifically, five independent models were trained on different cross-validation splits of the eICU dataset, and each model was independently applied to the external MIMIC-IV and IUH datasets, resulting in five AUCs for each external validation to assess robustness. Model performance was evaluated using the area under the receiver operating characteristic curve (AUC), with median and interquartile range (IQR) reported across the five models. Results were visualized using boxplots.

### Performance across different prediction windows, outcome windows, and base models

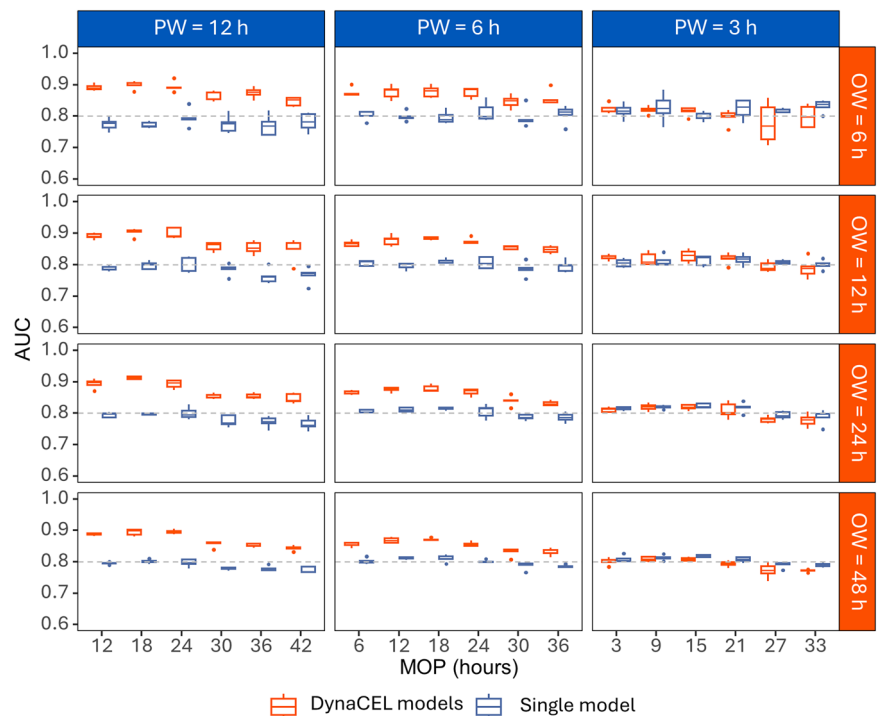
We evaluated three prediction windows and four outcome window lengths across six MOPs using the eICU test set, MIMIC-IV dataset, and IUH dataset (Figs. 7–9). DynaCEL models with a 12-hour prediction window and a 24-hour outcome window consistently outperformed models using other durations. Model performance across different base algorithms was assessed using the eICU test set under various predictions and four outcome window combinations and MOPs (Fig. 7 and Supplementary Table 11). Random forest consistently outperformed other models, achieving an AUC of 0.912, compared to 0.788 for multilayer perceptron, 0.763 for support vector machine, and 0.796 for logistic regression in the 18-hour MOP cohort using a 12-hour prediction window and a 24-hour outcome window (Fig. 7). As a result, all DynaCEL models used in this study were based on the random forest algorithm with a 12-hour prediction window and a 24-hour outcome window.

We specifically compared the performance of the XGBoost base model with the random forest model in reconstructing mortality risk (Supplementary Table 12) and in reliably inferring hemodynamic management targets (Supplementary Fig. 8) for the MIMIC-IV MOP 18-hour cohort, using a 12-hour prediction window and a 24-hour outcome window. The results suggested no noticeable difference in reconstructing mortality risk; however, the XGBoost base model demonstrated artifacts in target inference.

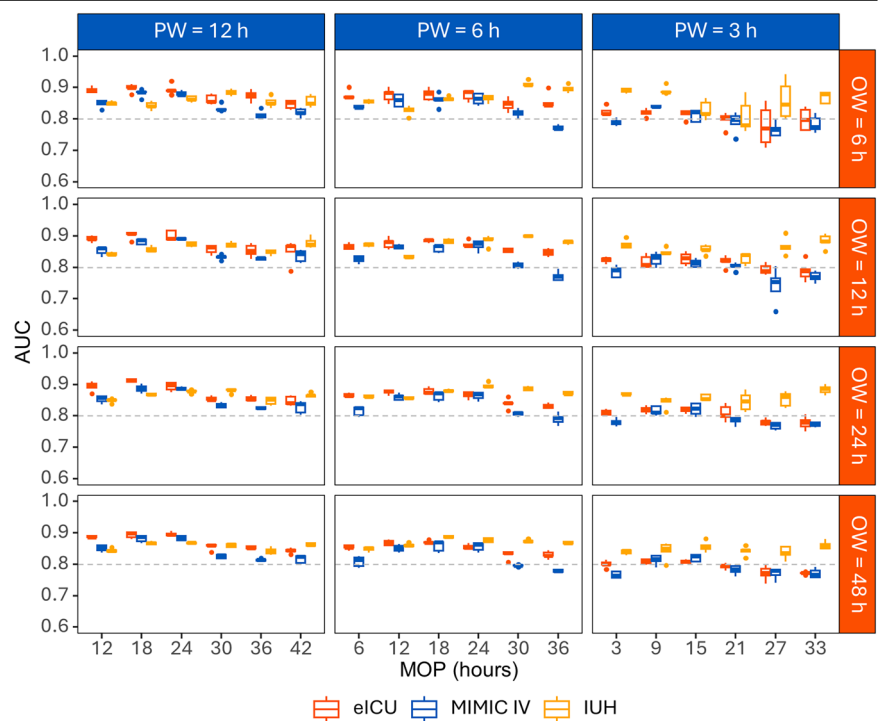
### Comparison with a single model

The DynaCEL strategy used multiple dynamic models, each trained on a specific MOP cohort. To evaluate its effectiveness, we compared this approach to a single model trained on the pooled data (originally used to train individual DynaCEL models), with MOP included as an input feature. DynaCEL outperformed the single model when the prediction window was 6 or 12 h (but not 3 h) across various outcome windows and MOPs (Fig. 8). For example, in the 18-hour MOP cohort from the eICU test set, DynaCEL achieved an AUC of 0.912 versus 0.795 for the single model (Fig. 6b and Supplementary Table 13). Both models used a random forest algorithm with a 12-hour prediction window and a 24-hour outcome window.

**Fig. 8 | Comparison of DynaCEL with a single model trained on pooled data.** This figure compares the performance of DynaCEL models trained separately for each MOP cohort with a single model trained on pooled data across all MOPs, where MOP was included as an input feature. Both models used random forest and were evaluated across various PW and OW configurations. AUCs were derived from the eICU test set using MOP-specific cohorts to reflect dynamic population changes and time-dependent shifts in clinical characteristics. DynaCEL Dynamic Cohort Ensemble Learning, AUC area under the receiver operating characteristic curve, MOP moment of prediction, PW predictor window, OW outcome window, eICU eICU Collaborative Research Database.



**Fig. 9 | DynaCEL performance across datasets and prediction configurations.** AUCs were evaluated for DynaCEL models using random forest across multiple PW and OW combinations to assess generalizability across datasets and time points. Models were trained on MOP-specific cohorts from the eICU training set and externally validated on the eICU test set, as well as the MIMIC-IV and IUH datasets. Consistent performance across varied PW and OW configurations demonstrates robustness to both temporal and institutional variations. DynaCEL Dynamic Cohort Ensemble Learning, AUC area under the receiver operating characteristic curve, PW predictor window, OW outcome window, MOP moment of prediction, eICU eICU Collaborative Research Database, MIMIC-IV Medical Information Mart for Intensive Care, IUH Indiana University Health.



### Performance across different datasets

Using a random forest base model with a 12-hour prediction window and a 24-hour outcome window, DynaCEL models achieved AUCs of 0.841–0.912 in the eICU test set, 0.825–0.889 in MIMIC IV, and 0.850–0.883 in the IUH dataset across MOPs at 12, 18, 24, 30, 36, and 42 h (Fig. 9 and Supplementary Table 13). At the 18-hour MOP, the AUC was 0.912 for eICU, 0.889 for MIMIC-IV, and 0.868 for IUH (Fig. 6c and Supplementary Table 13). Notably, IUH models trained on later MOPs

yielded higher AUCs than those trained on earlier MOPs in the eICU and MIMIC-IV datasets.

### Sample size sensitivity analysis

A sensitivity analysis assessed whether sample sizes were sufficient for training the DynaCEL models. Models were trained on varying proportions of the eICU training data (e.g., 1%, 5%, 10%, 50%) and evaluated using consistent test and validation datasets. For each proportion, five models

were trained with randomly selected data from different combinations of four eICU training folds. Performance was assessed on the held-out eICU test sets and validated on the MIMIC-IV and IUH datasets to evaluate robustness across sample sizes. Results showed that performance saturation varied by dataset (Supplementary Fig. 9). In the eICU and MIMIC-IV datasets, model performance (AUC) increased with larger training sample sizes and saturated at approximately 80% of the training data. In the IUH dataset, saturation was observed at around 30%.

### Determination of personalized HR/SBP targets

The primary objective was to identify HR and SBP targets associated with the lowest mortality risk at specific MOPs for personalized hemodynamic management in critically ill patients. Optimal targets were inferred via a grid search of clinically plausible HR/SBP combinations, with DynaCEL models estimating associated mortality risk. Because each MOP includes only one observed HR/SBP pair, theoretical alternatives were evaluated to determine the combination with the lowest predicted mortality (Fig. 1). Personalized HR/SBP targets were defined as the median values from the 100 combinations yielding the lowest predicted mortality probabilities.

We used the 18-hour MOP cohort from the MIMIC-IV dataset ( $n = 63,310$ ) to illustrate personalized target estimation. The median actual HR in the 12-hour prediction window was 83 bpm (IQR, 72–95), compared with a predicted median target of 79 bpm (IQR, 74–83). For SBP, the actual median was 115 mmHg (IQR, 105–128) versus a predicted target of 121 mmHg (IQR, 115–127). The median HR deviation (actual minus predicted) was 3 bpm (IQR, −4 to 14), and the median SBP deviation was −4 mmHg (IQR, −14 to 4).

### Visualization of personalized HR/SBP targets

To improve clinical usability, we visualized the relationship between HR, SBP, and mortality risk at individual MOPs using color contour maps with Gaussian smoothing (Fig. 1). These maps depict the mortality risk landscape across HR/SBP combinations, marking both the predicted optimal HR/SBP targets and the patient's actual HR/SBP levels. Discrepancies between predicted and actual values may indicate a potential need for intervention at the treating physician's discretion.

### Validation of personalized HR/SBP targets

Validating the proposed personalized HR and SBP targets is essential. While randomized controlled trials comparing personalized, target-guided management with usual care represent the gold standard, the following analyses provide preliminary insights into the validity of this approach.

First, we assessed the impact of deviations between actual HR and SBP values and model-predicted personalized targets. A  $\pm 20\%$  relative deviation threshold was selected based on several considerations. Optimal HR and SBP values vary substantially across patients and clinical trajectories in the ICU. A relative threshold enables flexible modeling of inter-individual variability and evolving physiologic states, whereas absolute cutoffs may be overly rigid. To support this rationale, we visualized the distribution of HR and SBP values in the MIMIC-IV dataset at 18 h post-admission, highlighting the wide variability in these parameters across ICU patients (Supplementary Fig. 10). We recognize that the  $\pm 20\%$  cutoff is somewhat arbitrary, and future work could test alternative thresholds (e.g.,  $\pm 25\%$ ) or directionally weighted thresholds based on clinical context.

We acknowledge that defining “within targets” as within  $\pm 20\%$  of the model-recommended or fixed values may appear arbitrary. For instance, a heart rate of 60 bpm is widely accepted as the lower limit of normal, yet it falls outside a  $\pm 20\%$  range around a fixed benchmark of 80 bpm. However, a relative deviation threshold offers several advantages in the ICU context. First, it allows for personalization and flexibility, accommodating the broad physiologic diversity and shifting hemodynamic demands seen in critically ill patients. Second, it enables consistent comparisons across a wide range of baseline values. Nevertheless, this threshold should be interpreted as a pragmatic starting point rather than a definitive clinical boundary. Future research could explore sensitivity to threshold selection, incorporate

directional weighting for hypotension versus hypertension, and evaluate clinical relevance through prospective testing.

The MOP-specific cohort was stratified based on whether the actual HR and SBP values fell within or beyond the  $\pm 20\%$  threshold. Propensity scores were constructed using demographic variables, comorbidities, the worst SOFA subscores observed during the prediction window (i.e., the 12-hour period preceding the MOP), and respiratory support status at MOP (excluding HR and SBP). This ensured comparability in illness severity across groups before comparing subsequent outcomes. We applied nearest neighbor matching with a caliper width of 0.1 and used a variable matching ratio to maximize the number of ICU stays retained after matching. Covariate balance was assessed using a standardized difference threshold of 10%, and odds ratios were calculated to compare mortality between groups.

To further assess the validity of our stratification and ensure robustness of the propensity score-matched comparison, we performed a supplementary analysis using the MIMIC-IV dataset at MOP = 18 h. Patients were categorized into two mutually exclusive subgroups: (1) both HR and SBP within  $\pm 20\%$  of the DynaCEL-recommended targets, and (2) either HR or SBP beyond  $\pm 20\%$ . Propensity score matching was conducted as previously described. We then examined the use of vasoactive drugs and antibiotics—variables not included in the matching model but clinically relevant to hemodynamic status. These variables were expressed as categorical indicators of administration during the 12-hour prediction window. The results, shown in Supplementary Table 14, demonstrated comparable medication usage across subgroups, with absolute standardized differences  $< 0.111$  for all variables. This analysis suggests minimal confounding from unadjusted hemodynamic medications and supports the validity of our subgroup comparisons.

Second, a dose-response analysis stratified the MOP-specific cohort into subgroups with increasing deviations from targets (e.g.,  $\pm 20$ – $25\%$ ,  $\pm 25$ – $30\%$ ,  $\pm 30$ – $35\%$ ). This tested the hypothesis that greater deviations correlate with higher mortality risk. Crude mortality and odds ratios were reported for each subgroup, with the  $\pm 20\%$  subgroup serving as a reference.

Lastly, we compared patients within  $\pm 20\%$  of the personalized targets to those within  $\pm 20\%$  of population-based fixed targets (e.g., HR = 80 bpm, SBP = 120 mmHg). If the personalized targets are valid, we hypothesized lower mortality among patients closer to them. This comparison also used propensity score-matched analyses as previously described.

### Subpopulation analysis and model fairness

We conducted subpopulation analyses to assess model performance across demographics (age, sex, race), baseline features (body mass index), and clinical characteristics (comorbidity burden, illness severity) to evaluate fairness and generalizability. HR–BP–mortality contour maps with Gaussian smoothing visualized risk distributions across HR/BP combinations within MOP-specific cohorts. These maps highlighted mortality risk differences across subpopulations and served as an additional validation layer, with accurate predictions expected to reflect known patterns—for example, higher mortality among older or more severely ill patients.

Subpopulation characteristics are summarized in Supplementary Table 15. Model performance was assessed using the 18-hour MOP cohort from the eICU test set and MIMIC-IV and IUH datasets (Fig. 6d, Supplementary Table 16). The model performed comparably across vulnerable subgroups, including patients aged  $\geq 65$  years, females, African Americans, and Latino/Hispanic individuals. Performance was consistent across BMI categories ( $< 30$  vs.  $\geq 30$ ), comorbidity burden (Charlson Comorbidity Index  $\geq 2$  vs.  $< 2$ ), and illness severity (SOFA score  $\geq 5$  vs.  $< 5$ ). However, in the IUH dataset, performance was lower among patients with SOFA  $\geq 5$  compared to those with SOFA  $< 5$ .

The relationship between HR, SBP, and mortality was also visualized across subgroups (Fig. 6e). As expected, older patients, those with more comorbidities, and patients with higher illness severity in the 18-hour MOP MIMIC-IV cohort had greater mortality risk. Among racial/ethnic groups, Caucasians showed lower overall mortality compared to others.

### Feature importance analysis

We used Gini impurity reduction (scikit-learn version 1.6.1) and Shapley additive explanations (SHAP) to evaluate the contribution of each feature (Supplementary Table 17). Heart rate and systolic blood pressure ranked among the top 10 of 44 features and are highlighted in Supplementary Table 17 and Supplementary Fig. 11.

### Sensitivity analysis using an encoder-decoder model

To assess whether other temporal features influence the inference of hemodynamic targets, we conducted a sensitivity analysis using an encoder-decoder model with a masking approach to reconstruct the input features. Specifically, we masked selected features and used the encoder-decoder model to predict their values, thereby estimating their dependence on HR and SBP. The reconstructed input features were then fed into the pretrained DynaCEL model to infer hemodynamic management targets, which were compared to those derived from the original, unmasked inputs.

The encoder-decoder model comprised a gated recurrent unit (GRU)-based encoder to extract latent representations from sequential clinical data, followed by a GRU decoder that reconstructed the original time series. This architecture captures temporal dependencies and compresses the multivariate clinical state into a hidden latent space. The model processed hourly time series data from 14 clinical variables—including HR, SBP, SOFA-related variables, and respiratory support parameters. The latent space had 64 dimensions, and the model was trained to minimize the mean squared error between the original and reconstructed sequences, thereby learning the joint distribution of temporal clinical features.

We used a dynamic cohort defined by MOP = 18 h, a 12-hour prediction window, and a 24-hour outcome window from the MIMIC-IV dataset. Results from this analysis are shown in Supplementary Fig. 12. The inferred HR and SBP targets remained consistent when reconstructed inputs were used in place of original inputs. These findings suggest that the DynaCEL model is not sensitive to dependencies between HR/SBP and other temporal features, supporting the robustness of its recommendations in the presence of potential feedback relationships.

### Patient involvement and ethical considerations

This retrospective study did not involve patients in formulating the research question, defining outcomes, or developing recruitment or implementation strategies. Patients were also not involved in interpreting or reporting results. However, findings will be shared with patients and families via scientific presentations and publications to promote future engagement. All analyses followed ethical standards, including data anonymization and privacy safeguards, as approved by the institutional review board.

### Data availability

The Indiana University Health dataset is available upon request via email to Lingzhong Meng (menglz@iu.edu) and is pending institutional approval. The eICU and MIMIC-IV datasets are accessible via their corresponding sources.

### Code availability

The codes used in the current study to develop the algorithm are provided in GitHub (<https://github.com/Su-informatics-lab/DynaCEL>).

Received: 15 April 2025; Accepted: 3 July 2025;

Published online: 24 July 2025

### References

- Vincent, J. L., Parquier, J. N., Preiser, J. C., Brimioulle, S. & Kahn, R. J. Terminal events in the intensive care unit: review of 258 fatal cases in one year. *Crit. Care Med.* **17**, 530–533 (1989).
- Omar, M. A. K., Aram, F. O. & Banafa, N. S. Causes of mortality among critically ill patients admitted in intensive care unit. *Bahrain Med. Bull.* **37**, 178–180 (2015).
- Morelli, A. et al. Effect of heart rate control with Esmolol on hemodynamic and clinical outcomes in patients with septic shock: a randomized clinical trial. *JAMA* **310**, 1683–1691 (2013).
- Mareshwari, K. et al. The relationship between ICU hypotension and in-hospital mortality and morbidity in septic patients. *Intensive Care Med.* **44**, 857–867 (2018).
- Meng, L. Heterogeneous impact of hypotension on organ perfusion and outcomes: a narrative review. *Br. J. Anaesth.* **127**, 845–861 (2021).
- McGuigan, P. J. et al. The effect of blood pressure on mortality following out-of-hospital cardiac arrest: a retrospective cohort study of the United Kingdom Intensive Care National Audit and Research Centre database. *Crit. Care* **27**, 4 (2023).
- Nguyen, D., Kritek, P. A., Greco, S. A. & Prutkin, J. M. Bradycardia at the onset of pulseless electrical activity arrests in hospitalized patients is associated with improved survival to discharge. *Heliyon* **6**, e03491 (2020).
- Yong, J., Hibbert, P., Runciman, W. B. & Coventry, B. J. Bradycardia as an early warning sign for cardiac arrest during routine laparoscopic surgery. *Int. J. Qual. Health Care* **27**, 473–478 (2015).
- Asfar, P. et al. High versus low blood-pressure target in patients with septic shock. *N. Engl. J. Med.* **370**, 1583–1593 (2014).
- Font, M. D., Thyagarajan, B. & Khanna, A. K. Sepsis and Septic Shock - Basics of diagnosis, pathophysiology and clinical decision making. *Med. Clin. North Am.* **104**, 573–585 (2020).
- Gavelli, F., Castello, L. M. & Avanzi, G. C. Management of sepsis and septic shock in the emergency department. *Intern. Emerg. Med.* **16**, 1649–1661 (2021).
- Han, J. et al. Care guided by tissue oxygenation and haemodynamic monitoring in off-pump coronary artery bypass grafting (Bottomline-CS): assessor blind, single centre, randomised controlled trial. *BMJ* **388**, e082104 (2025).
- De Backer, D. et al. A plea for personalization of the hemodynamic management of septic shock. *Crit. Care* **26**, 372 (2022).
- Sharma, S. et al. ICU mortality in patients with Coronavirus disease 2019 infection: highlighting healthcare disparities in rural Appalachia. *Crit. Care Explor* **3**, e547 (2021).
- Khan, S. H. et al. Development of a population-level prediction model for intensive care unit (ICU) survivorship and mortality in older adults: A population-based cohort study. *Health Sci. Rep.* **6**, e1634 (2023).
- Pollard, T. J. et al. The eICU Collaborative Research Database, a freely available multi-center database for critical care research. *Sci. Data* **5**, 180178 (2018).
- Johnson, A. E. W. et al. MIMIC-IV, a freely accessible electronic health record dataset. *Sci. Data* **10**, 1 (2023).
- Pinsky, M. R. et al. Effective hemodynamic monitoring. *Crit. Care* **26**, 294 (2022).
- Thorsen-Meyer, H. C. et al. Dynamic and explainable machine learning prediction of mortality in patients in the intensive care unit: a retrospective study of high-frequency data in electronic patient records. *Lancet Digit. Health* **2**, e179–e191 (2020).
- Chan, B. et al. Generalizable deep temporal models for predicting episodes of sudden hypotension in critically ill patients: a personalized approach. *Sci. Rep.* **10**, 11480 (2020).
- Thorsen-Meyer, H. C. et al. Discrete-time survival analysis in the critically ill: a deep learning approach using heterogeneous data. *NPJ Digit. Med.* **5**, 142 (2022).
- Alves, T., Laender, A., Veloso, A. & Ziviani, N. Dynamic Prediction of ICU Mortality Risk Using Domain Adaptation. in *2018 IEEE International Conference on Big Data (Big Data)* 1328–1336 (2018).



23. Bednarski, B. P. et al. Temporal convolutional networks and data rebalancing for clinical length of stay and mortality prediction. *Sci. Rep.* **12**, 21247 (2022).
24. Deasy, J., Lio, P. & Ercole, A. Dynamic survival prediction in intensive care units from heterogeneous time series without the need for variable selection or curation. *Sci. Rep.* **10**, 22129 (2020).
25. Wardi, G. et al. Predicting progression to septic shock in the emergency department using an externally generalizable machine-learning algorithm. *Ann. Emerg. Med.* **77**, 395–406 (2021).
26. Nemati, S. et al. An interpretable machine learning model for accurate prediction of sepsis in the ICU. *Crit. Care Med.* **46**, 547–553 (2018).
27. Fan, H. et al. Development and validation of a dynamic delirium prediction rule in patients admitted to the Intensive Care Units (DYNAMIC-ICU): A prospective cohort study. *Int J. Nurs. Stud.* **93**, 64–73 (2019).
28. Kurtz, D. M. et al. Dynamic risk profiling using serial tumor biomarkers for personalized outcome prediction. *Cell* **178**, 699–713.e619 (2019).
29. Gu, Y. et al. Predicting medication adherence using ensemble learning and deep learning models with large scale healthcare data. *Sci. Rep.* **11**, 18961 (2021).
30. Zhao, Y. et al. Ensemble learning predicts multiple sclerosis disease course in the SUMMIT study. *NPJ Digit. Med.* **3**, 135 (2020).
31. Tanveer, M. et al. Ensemble deep learning for Alzheimer's disease characterization and estimation. *Nat. Ment. Health* **2**, 655–667 (2024).
32. Biderman, D. et al. Lightning Pose: improved animal pose estimation via semi-supervised learning, Bayesian ensembling, and cloud-native open-source tools. *bioRxiv* (2024).
33. Peng, H., Wang, H., Kong, W., Li, J. & Goh, W. W. B. Optimizing differential expression analysis for proteomics data via high-performing rules and ensemble inference. *Nat. Commun.* **15**, 3922 (2024).
34. Hu, W. et al. An interpretable ensemble learning model facilitates early risk stratification of ischemic stroke in intensive care unit: Development and external validation of ICU-ISPM. *Comput. Biol. Med.* **166**, 107577 (2023).
35. Ren, N., Zhao, X. & Zhang, X. Mortality prediction in ICU using a stacked ensemble model. *Comput. Math. Methods Med.* **2022**, 3938492 (2022).
36. Toptas, M. et al. Factors affecting the length of stay in the intensive care unit: our clinical experience. *Biomed. Res. Int.* **2018**, 9438046 (2018).
37. Resche-Rigon, M., Azoulay, E. & Chevret, S. Evaluating mortality in intensive care units: contribution of competing risks analyses. *Crit. Care* **10**, R5 (2006).
38. Williamson, J. M., Satten, G. A., Hanson, J. A., Weinstock, H. & Datta, S. Analysis of dynamic cohort data. *Am. J. Epidemiol.* **154**, 366–372 (2001).
39. Sevransky, J. E. et al. Hemodynamic goals in randomized clinical trials in patients with sepsis: a systematic review of the literature. *Crit. Care* **11**, R67 (2007).
40. Fu, E. L. et al. Merits and caveats of propensity scores to adjust for confounding. *Nephrol. Dial. Transplant.* **34**, 1629–1635 (2018).
41. Yang, J. et al. Identification of clinical subphenotypes of sepsis after laparoscopic surgery. *Laparosc., Endosc. Robot. Surg.* **7**, 16–26 (2024).
42. Zhang, Z. et al. Causal inference with marginal structural modeling for longitudinal data in laparoscopic surgery: A technical note. *Laparosc., Endosc. Robot. Surg.* **5**, 146–152 (2022).
43. Seymour, C. W. et al. Derivation, validation, and potential treatment implications of novel clinical phenotypes for sepsis. *JAMA* **321**, 2003–2017 (2019).
44. Leligdowicz, A. & Matthay, M. A. Heterogeneity in sepsis: new biological evidence with clinical applications. *Crit. Care* **23**, 80 (2019).
45. Sauer, C. M. et al. Systematic review and comparison of publicly available ICU Data Sets—A decision guide for clinicians and data scientists. *Crit. Care Med.* **50**, e581–e588 (2022).
46. Quan, H. et al. Updating and validating the Charlson comorbidity index and score for risk adjustment in hospital discharge abstracts using data from 6 countries. *Am. J. Epidemiol.* **173**, 676–682 (2011).
47. Charlson, M. E., Carrozzino, D., Guidi, J. & Patierno, C. Charlson Comorbidity Index: A critical review of clinimetric properties. *Psychother. Psychosom.* **91**, 8–35 (2022).
48. Vincent, J. L. et al. Use of the SOFA score to assess the incidence of organ dysfunction/failure in intensive care units: results of a multicenter, prospective study. Working group on “sepsis-related problems” of the European Society of Intensive Care Medicine. *Crit. Care Med.* **26**, 1793–1800 (1998).

## Acknowledgements

We thank Rick V. Tuason from Clinical Research Systems, Enterprise Analytics, Indiana University Health, located in Indianapolis, Indiana, USA, for his help in preparing the Indiana University Health dataset. J.S., X.L., and Z.L. are financially supported by the Indiana University Precision Health Initiative. We also thank institutional and/or departmental sources for their support.

## Author contributions

L.M. and J.S. conceived and supervised the project. L.M., X.L., J.S., and X.Z. designed the model and computational framework. L.M. designed the study and contributed to the initial drafting of the paper. X.L., Z.L., J.L., Y.S., A.D.P., G.L., and A.B. collected and prepared data. X.L. and J.C. contributed to model development and validation. J.L., Y.S., and Z.L. contributed to the graphic illustration. J.L. and Z.L. (2nd) performed the statistical analysis. All authors provided critical comments, reviewed the paper, discussed the results, and approved the final version.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41746-025-01863-0>.

**Correspondence** and requests for materials should be addressed to Jing Su.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025