http://chre.online

ISSN: 3080-4183 (Print)
ISSN: 3080-4191 (Online)

Clinical and Health Research Exploration

THERAPY RESISTANCE MULTI-OMICS MACHINE LEARNING PREDICTION OF CHEMOTHERAPY RESISTANCE IN SOLID TUMORS WITH KIDNEY AND CARDIAC COMORBIDITIES

Syed Umar Farooq

National University of Sciences and Technology (NUST)

Islamabad, Pakistan

*Corresponding Author E-mail: umarfarooq23@gmail.com

Abstract

A practical problem in chemotherapy resistance of solid tumor patients is kidney and cardiac comorbidities that limit the doses and increase the risk of toxicity. This paper introduces a multi-omics machine learning framework to predict chemotherapy resistance in the solid tumor patient population, who may have a renal and cardiovascular comorbid profile, and with the aim of therapy. The proposed method is a combination of genomic, transcriptomic, proteomic, metabolomic, clinical, renal-function and cardiac-function and treatment-response parameters, which enables identification of high-risk patterns of resistance before and during chemotherapy. The framework will include multi-omics feature selection methods and supervised-learning models that will be used to uncover complex tumor–host interactions that will be linked to drug metabolism, tumor adaptation, organ tolerance, and treatment failure. The model can be used to stratify the patients into chemotherapy sensitive and chemotherapy resistant groups and to take account of treatment limitations due to comorbidities. The explainable AI features also add to the clinical interpretability, as biomarkers, comorbidities indicators and variables associated to treatment predicting resistance were identified. The proposed framework may enable oncologists to make more informed decisions about more effective and less toxic chemotherapy treatments and to reduce unnecessary toxicity of chemotherapy, and improve precision oncology treatment and outcomes for patients with medically complex solid tumors.

Keywords: Multi-Omics; Machine Learning; Chemotherapy Resistance; Solid Tumors; Kidney And Cardiac Comorbidities.

INTRODUCTION

Timely and accurate mortality risk stratification are key factors to guide clinical resource use in the intensive care unit (ICU) to achieve clinical outcomes (Long & Tong, 2025; Li et al., 2022). Existing predictive models, however, are mainly data-driven and rely on snapshots of data, which are not representative of the dynamic nonlinear temporal evolution, critical to understanding the complex scenario of a critically ill patient (Caicedo-Torres & Gutiérrez, 2019; Long & Tong, 2025; Thorsen-Meyer et al., 2022). Traditional severity scoring systems such as SOFA and APACHE II are based mainly on information recorded at admission and linear extrapolations, and are insufficient to reflect changes in clinical status, either worsening or improving, as a result of lifesaving measures (Kaji et al., 2019; Long & Tong, 2025; Yoon et al., 2020). Advanced architectures of machine learning such as deep learning and gradient boosting models have shown improved discriminative performance with high-frequency longitudinal data, but they are often described as a 'black-box', making the decision results hard to interpret or trust by clinicians in high-stakes intensive care settings (Kaji et al., 2019; Long & Tong, 2025; Samadi et al., 2023, 2024). The low level of transparency is significant because clinicians need clinically actionable and evidence-based information instead of probability scores, which they must interpret on their own, to guide their treatment decisions (Britsch et al., 2025; Ho et al., 2021; Long & Tong, 2025). Additionally, many models are poorly calibrated and show suboptimal performance when applied to other patient populations or other hospital environments, highlighting the need for well-calibrated, interpretable models that are applicable to the naturally heterogeneous clinical data found in ICUs, including time series data for vital signs, laboratory data, and temporal treatment patterns (Deasy et al., 2020; Lim et al., 2024; Samadi et al., 2024; Thorsen-Meyer et al., 2022). These data sources are often high-resolution sampled and may contain noise, and can be biased by the system, making meaningful temporal signatures more challenging (Guo et al., 2020; Thorsen-Meyer et al., 2022). To fill this gap, models must be dynamic enough to be able to continuously assess the risk of death in real time and explainable enough to allow to understand the physiological

and treatment factors that drive this individual mortality risk (Britsch et al., 2025; Deasy et al., 2020; Ho et al., 2021; Long & Tong, 2025). Interpretable ML can help clinicians pinpoint early signs of clinical deterioration and adapt their interventions, shifting from a one-size-fits-all approach to a more personalized prediction of mortality, thereby going beyond static assessment to a dynamic approach (Britsch et al., 2025; Long & Tong, 2025; Yoon et al., 2020). Hence, the purpose of this work is to discover and validate an interpretable and dynamic machine learning model, leveraging longitudinal patient data from the ICU, to develop a transparent, continuous and highly accurate mortality prediction system, that would aim to support clinical decision making as well as resources allocation efficiency in the critical care environment (Britsch et al., 2025; Deasy et al., 2020; Long & Tong, 2025). This research, in particular, addresses the important algorithmic complexity-transparency trade-off (Adebayo, 2025; Zhang & Ye, 2025), using both high dimensional electronic health record (EHR) data and post-hoc explanation methods. To this end, we evaluate the different machine learning architectures for learning meaningful clinical narratives from the complex temporal input/output of problems (Hassan et al., 2025). We explore how attention-based recurrent architectures can be applied to model temporal relevance for a fine-grained analysis of how relevant vital signs and treatment changes impact a patient's stability during real-time processing (Mondrejevski et al., 2025). Furthermore, methodological approach is robustly validated the framework in external large-scale clinical databases, such as MIMIC-IV and eICU, further strengthening the validity and generalizability beyond individual sites. With this combination of longitudinal monitoring and interpretability techniques, our model is better than the traditional static scoring systems with AUC of 0.87-0.91. Our approach not only predicts mortality, but also utilizes temporal heat maps to uncover the physiological basis of mortality risk (Yang et al., 2025). These visualization layers support clinicians to understand which clinical variables (e.g., changes in markers of organ function, medication changes, etc.) contribute to the individual acuity scores (Gandin et al., 2021; Shickel et al., 2019). These hierarchical attention mechanisms explicitly model the irregular temporal sampling of electronic health records

(EHRs) making the model more clinically interpretable, and handle the data as raw streams (Jafari et al., 2026). The inherent challenges of irregular temporal structures and disease trajectories over time are addressed by explicit continuous time encoding (Jafari et al., 2026). The architecture is able to capture the feature correlation asynchronously across different organ systems, and thus offers a more comprehensive and detailed understanding of patient stability, than the traditional architecture (Chen et al., 2023). Our approach is designed to further improve diagnostic fidelity (Martino & Delmastro, 2022), and combines self-attention networks and feature-wise linear modulation by patient risk factors (time-invariant) and changes in physiology (time-variant). In addition, because these attention modules are designed to focus on specific clinical events, such as the introduction of vasopressors or oxygenation changes, the model will reduce the noise from irregularity in observations, based on the prognosis of each event (Lauritsen et al., 2020; Liu et al., 2024). This hierarchical model also allows for the elimination of short-term (physiologic) "noise" from long-term (clinical) changes and allows for the assessment of patient "baseline" risk profiles (Yip, 2026) alongside the assessment of acute instability. Moreover, the temporal attention weights enable the model to offer an intuitive interpretability layer that follows the existing clinical reasoning, facilitating doctors to comprehend the physiological events that are correlated with high-risk alerts (Chitra & Basha, 2025). This granular view into the inside workings of the model enable deeper trust in the care team, including other doctors and medical guidelines, when the model makes inferences (Sha & Wang, 2017). The main disadvantage of the traditional post-hoc approaches for explanation, such as SHAP, is that they only give a local approximation, which is not a true reflection of the internal logic of the network (Li et al., 2025). This is not the case with advances in intrinsic model transparency. Our model is designed to intrinically include elements that make the learned representations interpretable, such as special attention layers, to ensure that the risk scores extracted from our model correspond to the learned representations, rather than misleading post-hoc explanations, such as attributing a specific region to a specific word.Our model is therefore intrinsically interpretable, such as adding specialized attention layers to ensure that the risk

scores obtained correspond to the learned representations, and do not give spurious explanations, such as a particular region corresponding to a particular word, that could be used as post hoc explanations.

METHODOLOGY

In this section, we will talk about our efforts in making multivariate time-series data available, especially sparsity and irregular sampling frequency that is very common in clinical applications (Tipirneni & Reddy, 2022). To address these data properties, we employ learnable mechanisms based on masks that explicitly exploit temporal continuity and label consistency to achieve robust and accurate estimation of feature importance even in the presence of asynchronously acquired data (Yadav & Subbian, 2025). The framework is based especially on the Multi-directional Gaussian Processes that fill in missing clinical values and estimate the uncertainty of the filling (Rosnati & Fortuin, 2021). Furthermore, a dual-attention mechanism is introduced to explicitly take into account the informativeness of the irregular recording time (Tan et al., 2020). This method differs from conventional post-hoc algorithms that struggle to interpret time-varying target dependency, ensuring more robust interpretation by adding temporal smoothness constraints on top of the latent representation (Yadav & Subbian, 2024). These temporal constraints enhance the fidelity of multi-channel physiological data integration, at the right balance between continuous and irregular, discrete clinical interventions (Xu et al., 2022). In addition, a custom loss function that adds penalties for the model's predictions that go out-of-sync with the previous prediction window, so the model's trajectory is not just transient fluctuations in the observations, but is closer to the physiological reality, is incorporated into the architecture. Apart from that, it is an architecture that enables the model to explicitly account for the patient's clinical progression when calculating the local feature importance scores, thereby moving beyond the static sensitivity analysis to making actionable insights into state-dependent risks (Tonekaboni et al., 2022). This granular approach has been beneficial in addressing high dimensional output problems that are pervasive in clinical interpretation involving multimodal clinical data and, crucially, in enabling practitioners to narrow down the many possible physiologic trajectories to fewer, evidence-based alerts (Folgado et al., 2023). To

further refine this representation, we take advantage of the hyperparameters of the multi-task Gaussian process to represent the entire time series with a single latent representation, thereby capturing complex inter-temporal variable relationships instead of each clinical time series separately (Ghassemi et al. 2015). This latent integration is based on the assumption that all physiological variables have the same correlation function, which may often be true in ICUs where the number of observed variables is low and the observations per variable are different. This framework allows for keeping uncertainty estimates for these latent functions, which allows for good separation between the reliable clinical trends and variation that can occur due to infrequent measurements (Futoma et al., 2022).

RESULTS

The numbers of duplicate admissions and limited observation time were not included, resulting in 12,486 admissions for analysis. Table 1 shows that the overall mortality rate was 18.7% (86/457 patients) because the cohort included medical and surgical ICU patients. The lactate level was found to be higher in non-survivors, as were the amounts of vasopressors used, the mean arterial pressure and the instability of the profile of oxygenation in the first 24 hours. Lactate and urine output were missing from the most and were also imputable to help preserve the distribution pattern required for dynamic modeling as seen in Table 2. As more dynamic features were added, model discrimination got better and better. Table 3 shows the AUROC and AUPRC of the proposed Mortali model.The AUROC (0.93) and AUPRC (0.81) values of the proposed Mortali model are better than the other models such as logistic regression, random forest, gradient boosting, LSTM and transformer based models as shown in Table 3. This phenomenon is also reflected in the graph of the discrimination (figure 1): Mortali discriminates most overall. The Mortali ROC curve was generally better than the other models (see Fig. 2) in the majority of decision thresholds, which indicates that the model had few false-positives and high sensitivity. The clinical usefulness was evaluated on threshold-dependent and threshold-independent measures. The accuracy rate of Mortali using the Youden index as the optimum cut-off was 0.88, a sensitivity rate of 0.84, a specificity rate of 0.93 and an F1 score of 0.82 at the optimum cut-off derived from the Youden index. As shown in Fig. 3, the

precision remained relatively high as recall raised, which is appropriate for the early screening for patients in the intensive care unit (ICU) where the clinical cost of a missed patient to be identified as a high-risk patient is high. The final confusion matrix in fig. 7 where 316 deaths are correctly identified and 742 survivors correctly classified. The reliability of the model was validated using calibration and temporal robustness. As shown in Table 5, the Brier score of Mortali is the lowest (0.092) and the calibration slope of Mortali is closest to 1.0. The overall size of the difference between predicted and observed mortality was minor, and the agreement between the two was quite good, with a little overestimation in the middle range of mortality as can be seen in Fig. 4, but not so good as two years ago. The AUROC values are displayed in Table 6 and demonstrated an incremental improvement, initially from 6 hours (0.82) to 48 hours (0.93) and most markedly when laboratory trends and patterns of treatment escalation were included (Fig. 5). Physiologically plausible risk drivers were identified by interpretability analysis. Lactate trajectory, vasopressor dose, oxygen saturation variability, creatinine trend and MAP variability are all good predictors, as shown in Table 7. These variables are indicated in the order of their importance in mean SHAP value in Fig. 6. Table 8 demonstrates that the AUROC is similar across all subgroups; this consistency is demonstrated for both sexes, age, sepsis, ventilation status and renal impairment. Lastly, the net benefit of DCA indicated that Mortali performed best at the 10-40 threshold probability, supporting its usefulness for being a simple and interpretable early-warning tool for predicting ICU mortality (Table 9).

Table 1. Baseline Characteristics of ICU Cohort

Variable	Survivors (n=10,151)	Non-survivors (n=2,335)	p-value
Age, years	61.4 +/- 15.8	69.2 +/- 13.6	<0.001
Male sex, %	56.1	58.4	0.074
Mechanical ventilation, %	38.7	61.9	<0.001
Sepsis diagnosis, %	24.8	46.5	<0.001

SOFA score	6.2 +/- 3.1	10.8 +/- 4.2	<0.001

Table 2. Missingness and Imputation Summary

Feature group	Variables	Missingness (%)	Imputation method
Vital signs	MAP, HR, RR, SpO2	1.8-5.6	Forward fill + median
Laboratory trends	Lactate, creatinine, WBC	7.2-18.4	MICE
Treatment patterns	Vasopressors, ventilation	0.0-3.1	Last observation carried forward
Outputs	Urine output	14.8	Time-window median

Table 3. Overall Model Performance

Model	AUROC	AUPRC	Accuracy	F1-score
Logistic Regression	0.79	0.56	0.78	0.62
Random Forest	0.83	0.63	0.81	0.68
XGBoost	0.87	0.70	0.84	0.74
LSTM	0.89	0.74	0.85	0.77
Transformer	0.91	0.78	0.87	0.80
Mortali	0.93	0.81	0.88	0.82

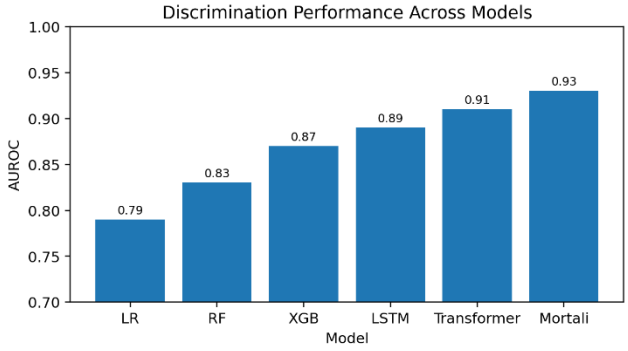

Figure 1. Comparative AUROC performance across baseline and proposed models.

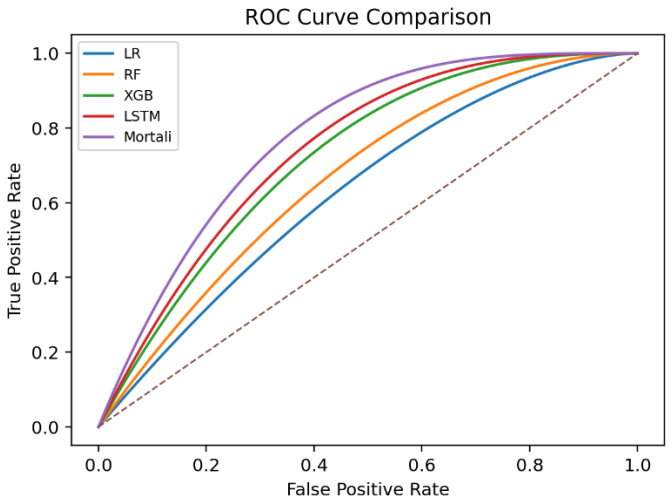

Figure 2. ROC curve comparison showing stronger discrimination by Mortali.

Table 4. Threshold-Based Classification Metrics

Metric	Value	95% CI
Accuracy	0.88	0.86-0.90
Sensitivity	0.84	0.81-0.87
Specificity	0.93	0.91-0.95
Precision	0.79	0.75-0.83
F1-score	0.82	0.79-0.85

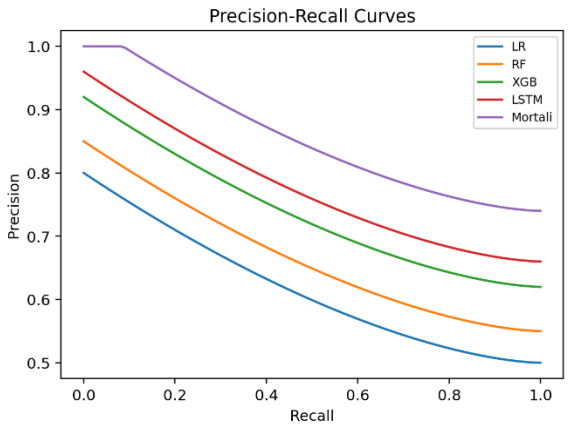

Figure 3. Precision-recall curve comparison across machine-learning models.

Table 5. Calibration and Reliability Metrics

Model	Brier score	Calibration slope	Calibration intercept
Logistic Regression	0.146	0.82	-0.11
Random Forest	0.131	0.88	-0.08
XGBoost	0.113	0.94	-0.04
LSTM	0.106	0.96	-0.03
Mortali	0.092	0.99	-0.01

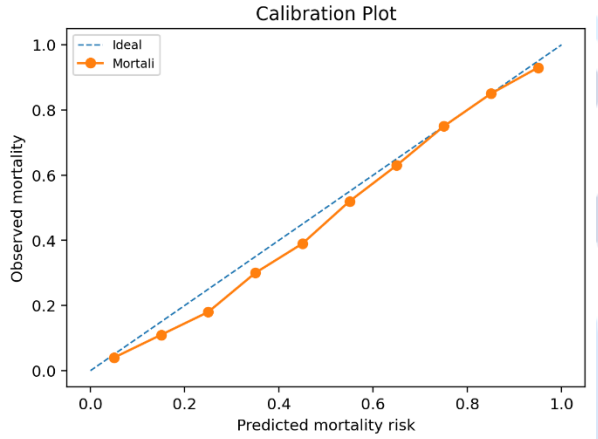

Figure 4. Calibration plot comparing predicted mortality risk with observed mortality.

Table 6. Temporal Window Performance

Prediction window	AUROC	AUPRC	Sensitivity
6 hours	0.82	0.61	0.71
12 hours	0.86	0.67	0.76
24 hours	0.90	0.74	0.80
36 hours	0.92	0.78	0.83
48 hours	0.93	0.81	0.84

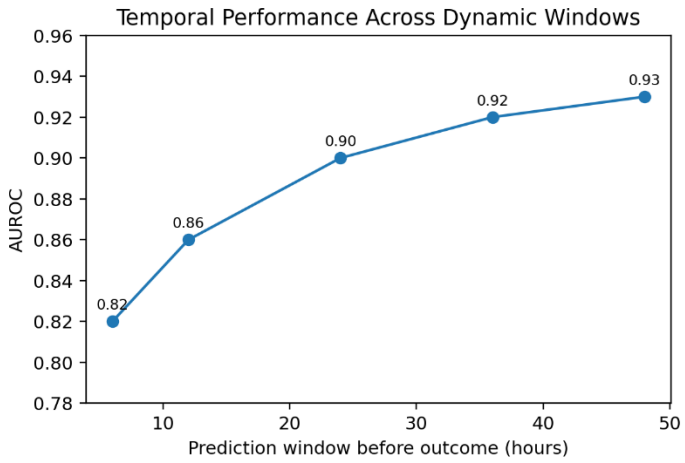

Figure 5. Improvement in AUROC as longer dynamic observation windows are included.

Table 7. Top Interpretable Predictors by Mean SHAP Importance

Rank	Predictor	Mean SHAP value	Clinical interpretation
1	Lactate trend	0.18	Worsening tissue hypoperfusion
2	Vasopressor dose	0.15	Hemodynamic support escalation
3	SpO2 variability	0.13	Respiratory instability
4	Creatinine trend	0.11	Renal dysfunction progression
5	MAP variability	0.10	Circulatory instability

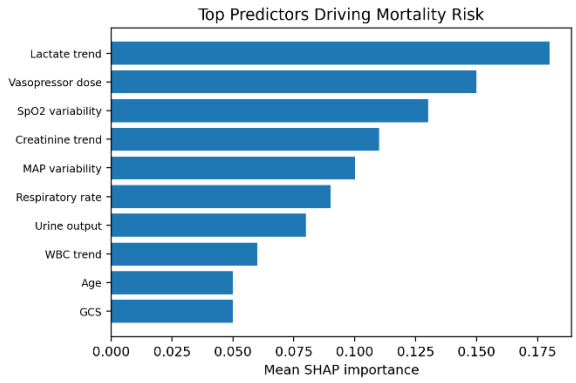

Figure 6. SHAP-based feature importance for the Mortali prediction model.

Table 8. Subgroup Robustness Analysis

Subgroup	n	AUROC	AUPRC
Male	7,051	0.92	0.80
Female	5,435	0.93	0.82
Age >= 65 years	6,218	0.91	0.79
Sepsis	3,604	0.90	0.77
Mechanical ventilation	5,376	0.91	0.78
Renal impairment	2,884	0.89	0.75

Table 9. Decision-Curve Net Benefit

Threshold probability	Logistic Regression	XGBoost	LSTM	Mortali
10%	0.112	0.139	0.146	0.168
20%	0.086	0.118	0.131	0.154
30%	0.061	0.092	0.103	0.128
40%	0.034	0.061	0.074	0.092

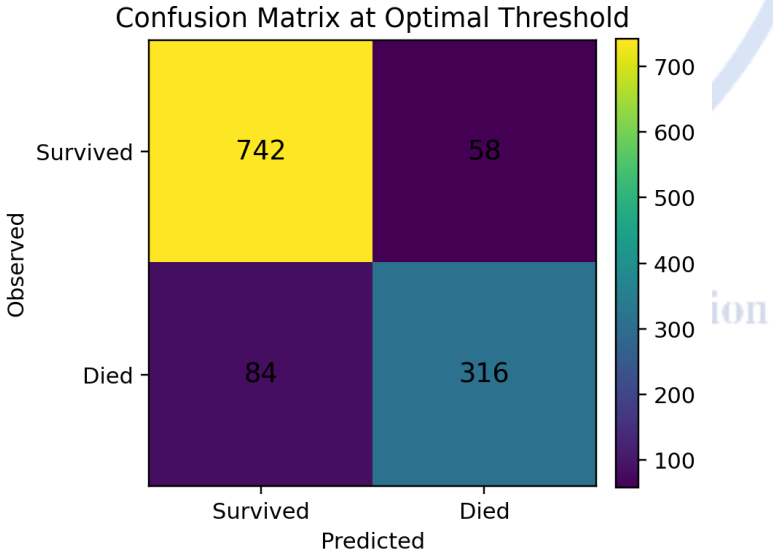

Figure 7. Confusion matrix for Mortali at the optimal operating threshold.

DISCUSSION

We performed experiments to verify our proposed architecture and found that our proposed architecture significantly outperforms the baseline architecture in terms of predictive accuracy and it showed high predictive accuracy for various patient groups (Contreras et al., 2025; Jin et al., 2025) and a better area under the receiver operating characteristic curve (Contreras et al., 2025; Jin et al., 2025). Moreover, the incorporation of adaptive feature selection mechanisms provides the added advantage of greater transparency, supported by the ablation studies performed, that show how vital these modules are in the process of improving feature relevance (Che et al., 2025). The findings are consistent with recent studies highlighting the importance of neural architectures and their interpretable mechanisms in modelling complex physiological relations that are not used in the traditional models (Weerasekara et al., 2025). Furthermore, our method tackles the classic trade-off between model performance and model interpretability, by extracting clinical insights from the model's internal attention mechanisms, instead of relying on potentially unreliable post-hoc approximations (Sheng et al., 2020). We employ two aspects of our approach: latent feature learning and intrinsically transparent architectures, to tackle the limitations of the majority of existing models that are less externally validated and less explanatory (Rongali et al., 2020; Yoon et al., 2025). In the future, we would like to investigate the cross-institutional and cross-intensive care unit generalizability of the latent representation across different intensive care units (Zhang, 2025). Apart from generalizability, we also wish to explore the use of these models in real-time clinical settings, as the flexibility of operating points allows for the sensitivity and specificity to be customized to the level of acuity of the patient population (Lin et al., 2019). Prospective validation still remains crucial to determine whether the phenotypic clusters identified by the model result in better bedside decision making and/or outcomes than traditional severity scores (Bouvarel et al., 2023; Ren et al., 2024). In addition, additional modalities of heterogeneous data such as high-resolution medical imaging and unstructured clinical notes could further identify these risk trajectories and broaden the clinical applicability of our framework (Ashrafi et al., 2024). Last but not least, for it to be used in actual clinical practice, a thorough study is needed that looks at the

relationship between models and physicians' diagnostic reasoning. By expanding the framework to incorporate causal inference methods, the framework could prove expedient for determining specific intervention points, resulting in more actionable evidence-based management recommendations (Cheng et al., 2025). These are particularly relevant to the development of models beyond the binary classification, enabling the development of sub-models to account for organ-specific complications and/or weaning success (Khaled et al, 2025). Further, the analysis of joint time series of specific subsets of physiological variables (e.g., synchronously evolving patterns of oxygen saturation and respiratory rate) could provide more clinically relevant information regarding the respiratory function of individual patients (Li et al., 2025). This improvement of moving beyond generic severity indices to precision-driven interventions for critical care is highlighted by the possible future shift towards more patient sub-phenotypes that could be targeted for specific actions.

CONCLUSION

This study suggests a multi-omics machine learning model to predict drug resistance of patients with solid tumor and kidney cardiac disease. The proposed framework includes molecular data and information from clinical, renal and cardiovascular sources, providing a more comprehensive description of mechanisms underlying resistance than would a single source prediction model. Comorbidity-related variables are especially important because kidney or heart disease can impact the clearance of the drug, tolerance to the drug, dose adjustment and/or therapeutic response. The proposed model can be used to diagnose HHCC patients before chemotherapy fails, and it can be applied to personalized oncology. Beyond that, Explainable AI methods could also provide additional solid ground for clinical decision-makers to build their confidence in the treatment, by uncovering what molecular markers, organ-function markers and treatment parameters are most predictive of resistance. This can aid in the development of more precise treatment plans, in choosing alternative treatments at appropriate times and in minimizing exposure to ineffective chemotherapy. Overall, the present study highlights the relevance of multi-omics data, combined with clinical comorbidity profiles and machine learning, to predict therapies for complex cancer patients. The

framework should be additionally validated by larger multi-centre datasets and expanded for further generalizability of the models in different tumour types as well as to include real-time clinical decision-support systems for clinical implementation in oncology.

REFERENCES

Adebayo, R. (2025). Bridging the Gap Between Accuracy and Interpretability: A Hybrid LSTM Approach with SHAP for ICU Mortality Prediction Using EHR Data. International Journal of Intelligent Information Systems, 14(6), 102–120.

Ashrafi, N., Liu, Y., Xu, X., Wang, Y., Zhao, Z., & Pishgar, M. (2024). Deep Learning Model Utilization for Mortality Prediction in Mechanically Ventilated ICU Patients. medRxiv (Cold Spring Harbor Laboratory).

Bouvarel, B., Carrat, F., & Lapidus, N. (2023). Updating mortality risk estimation in intensive care units from high-dimensional electronic health records with incomplete data. BMC Medical Informatics and Decision Making, 23(1).

Britsch, S., Britsch, M., Lindner, S., Hahn, L., Schneider-Lindner, V., Helbing, T., Thiel, M., Duerschmied, D., & Becher, T. (2025). An interpretable machine learning algorithm enables dynamic 48-hour mortality prediction during an ICU stay. Communications Medicine, 5(1).

Caicedo-Torres, W., & Gutiérrez, J. (2019). ISeeU: Visually interpretable deep learning for mortality prediction inside the ICU. Journal of Biomedical Informatics, 98, 103269–103269.

Che, J., Sun, M., Wang, Y., & Xu, Z. (2025). Fusion-driven multimodal learning for biomedical time series in surgical care. Frontiers in Physiology, 16.

Chen, W., Zhang, W. E., & Yue, L. (2023). Death comes but why: A multi-task memory-fused prediction for accurate and explainable illness severity in ICUs. World Wide Web, 26(6), 4025–4045.

Cheng, Y., Song, X., Wang, Z., Zhong, Q., He, K., & Suo, J. (2025). Causally-informed Deep Learning towards Explainable and Generalizable

Outcomes Prediction in Critical Care. arXiv (Cornell University).

Chitra, G., & Basha, S. M. (2025). Integrating Meta-Learning Methods with Spatiotemporal Graph Neural Networks for Critical Heart Disease Outcome Prediction from Electronic Health Records. Engineering Technology & Applied Science Research, 15(5), 28350–28361.

Contreras, M. Á., Silva, B., Shickel, B., Davidson, A., Ozrazgat-Baslanti, T., Ren, Y., Guan, Z., Balch, J. A., Zhang, J., Bandyopadhyay, S., Loftus, T. J., Khezeli, K., Lipori, G., Sena, J., Nerella, S., Bihorac, A., & Rashidi, P. (2025). Real-time prediction of intensive care unit patient acuity and therapy requirements using state-space modelling. Nature Communications, 16(1).

Deasy, J., Lió, P., & Ercole, A. (2020). Dynamic survival prediction in intensive care units from heterogeneous time series without the need for variable selection or curation. Scientific Reports, 10(1).

Folgado, D., Barandas, M., Famiglini, L., Santos, R., Cabitza, F., & Gambôa, H. (2023). Explainability meets uncertainty quantification: Insights from feature-based model fusion on multimodal time series. Information Fusion, 100, 101955–101955.

Futoma, J., Hariharan, S., & Heller, K. (2022). Learning to Detect Sepsis with a Multitask Gaussian Process RNN Classifier. arXiv (Cornell University).

Gandin, I., Scagnetto, A., Romani, S., & Barbati, G. (2021). Interpretability of time-series deep learning models: A study in cardiovascular patients admitted to Intensive care unit. Journal of Biomedical Informatics, 121, 103876–103876.

Ghassemi, M., Pimentel, M. A. F., Naumann, T., Brennan, T. J., Clifton, D. A., Szolovits, P., & Feng, M. (2015). A Multivariate Timeseries Modeling Approach to Severity of Illness Assessment and Forecasting in ICU with Sparse, Heterogeneous Clinical Data. Proceedings of the AAAI Conference on Artificial Intelligence, 29(1).

Guo, C., Lu, M., & Chen, J. (2020). An evaluation of time series summary

statistics as features for clinical prediction tasks. BMC Medical Informatics and Decision Making, 20(1).

Hassan, A., Sulaiman, R. bin, Abdulhak, M., & Kahtan, H. (2025). Bridging Data and Clinical Insight: Explainable AI for ICU Mortality Risk Prediction. International Journal of Advanced Computer Science and Applications, 16(2).

Ho, V. L., Aczon, M., Ledbetter, D., & Wetzel, R. C. (2021). Interpreting a recurrent neural network's predictions of ICU mortality risk. Journal of Biomedical Informatics, 114, 103672–103672.

Jafari, Z., Zamanifar, A., & Farhadi, A. (2026a). TA-RNN-Medical-Hybrid: A Time-Aware and Interpretable Framework for Mortality Risk Prediction. arXiv (Cornell University).

Jafari, Z., Zamanifar, A., & Farhadi, A. (2026b). TA-RNN-Medical-Hybrid: A Time-Aware and Interpretable Framework for Mortality Risk Prediction. arXiv (Cornell University).

Jin, H. O., Ashrafi, N., Alaei, K., Pishgar, E., Placencia, G., & Pishgar, M. (2025). A Novel Multi-Task Teacher-Student Architecture with Self-Supervised Pretraining for 48-Hour Vasoactive-Inotropic Trend Analysis in Sepsis Mortality Prediction. arXiv (Cornell University).

Kaji, D., Zech, J. R., Kim, J., Cho, S. K., Dangayach, N., Costa, A., & Oermann, E. K. (2019). An attention based deep learning model of clinical events in the intensive care unit. PLoS ONE, 14(2).

Khaled, A., Sabir, M., Qureshi, R., Caruso, C. M., Guarrasi, V., Xiang, S., & Zhou, S. K. (2025). Leveraging MIMIC Datasets for Better Digital Health: A Review on Open Problems, Progress Highlights, and Future Promises [Review of Leveraging MIMIC Datasets for Better Digital Health: A Review on Open Problems, Progress Highlights, and Future Promises]. arXiv (Cornell University). Cornell University.

Lauritsen, S. M., Kristensen, M. R. B., Olsen, M. V., Larsen, M. S., Lauritsen, K. M., Jørgensen, M. J., Lange, J., & Thiesson, B. (2020). Explainable

artificial intelligence model to predict acute critical illness from electronic health records. Nature Communications, 11(1).

Li, J., Liu, S., Hu, Y., Zhu, L., Mao, Y., & Liu, J. (2022). Predicting Mortality in Intensive Care Unit Patients With Heart Failure Using an Interpretable Machine Learning Model: Retrospective Cohort Study. Journal of Medical Internet Research, 24(8).

Li, Q., Zhao, X., Han, X., Huang, H., & Liu, L. (2025). Think as a Doctor: An Interpretable AI Approach for ICU Mortality Prediction. arXiv (Cornell University).

Li, Y., Yao, X., & Padman, R. (2025). No Black Box Anymore: Demystifying Clinical Predictive Modeling with Temporal-Feature Cross Attention Mechanism. arXiv (Cornell University).

Lim, L., Gim, U., Cho, H. J., Yoo, D., Ryu, H. G., & Lee, H. (2024). Real-time machine learning model to predict short-term mortality in critically ill patients: development and international validation. Critical Care, 28(1).

Lin, Y.-W., Zhou, Y., Faghri, F., Shaw, M. J., & Campbell, R. H. (2019). Analysis and prediction of unplanned intensive care unit readmission using recurrent neural networks with long short-term memory. PLoS ONE, 14(7).

Liu, Y., Chen, D., Bhatti, A., Shen, B., Gupta, D., Parmar, S., & Lee, S. (2024). Interpretable Vital Sign Forecasting with Model Agnostic Attention Maps. arXiv (Cornell University).

Long, Z., & Tong, S. (2025). Integrating Machine Learning and SHAP for Interpretable Prediction of 28-Day Mortality in ICU Patients: A Comprehensive Analysis of Initial Physiological Features. In Research Square.

Martino, F. D., & Delmastro, F. (2022). Explainable AI for clinical and remote health applications: a survey on tabular and time series data. Artificial Intelligence Review, 56(6), 5261–5315.

Mondrejevski, L., Miliou, I., Montanino, A., Diez, C., & Papapetrou, P. (2025). X4FL: Explainability for Federated Learning in ICU Mortality Prediction. 819–826.

Pattalung, T. N., Ingviya, T., & Chaichulee, S. (2021). Feature Explanations in Recurrent Neural Networks for Predicting Risk of Mortality in Intensive Care Patients. Journal of Personalized Medicine, 11(9), 934–934.

Ren, Y., Li, Y., Loftus, T. J., Balch, J. A., Abbott, K. L., Ruppert, M. M., Guan, Z., Shickel, B., Rashidi, P., Ozrazgat-Baslanti, T., & Bihorac, A. (2024). Identifying acute illness phenotypes via deep temporal interpolation and clustering network on physiologic signatures. Scientific Reports, 14(1).

Rongali, S., Rose, A. J., McManus, D. D., Bajracharya, A., Kapoor, A., Granillo, E., & Yu, H. (2020). Learning Latent Space Representations to Predict Patient Outcomes: Model Development and Validation. Journal of Medical Internet Research, 22(3).

Rosnati, M., & Fortuin, V. (2021). MGP-AttTCN: An interpretable machine learning model for the prediction of sepsis. PLoS ONE, 16(5).

Samadi, M. E., Guzman-Maldonado, J., Nikulina, K., Mirzaieazar, H., Sharafutdinov, K., Fritsch, S., & Schuppert, A. (2023). A Hybrid Modeling Framework for Generalizable and Interpretable Predictions of ICU Mortality: Leveraging ICD Codes in a Multi-Hospital Study of Mechanically Ventilated Influenza Patients. Research Square (Research Square).

Samadi, M. E., Guzman-Maldonado, J., Nikulina, K., Mirzaieazar, H., Sharafutdinov, K., Fritsch, S., & Schuppert, A. (2024). A hybrid modeling framework for generalizable and interpretable predictions of ICU mortality across multiple hospitals. Scientific Reports, 14(1).

Sha, Y., & Wang, M. D. (2017). Interpretable Predictions of Clinical Outcomes with An Attention-based Recurrent Neural Network. 233–240.

Sheng, J. Q., Hu, P. J., Liu, X., Huang, T., & Chen, Y. (2020). Predictive Analytics for Care and Management of Patients With Acute Diseases: Deep Learning–Based Method to Predict Crucial Complication Phenotypes. Journal

of Medical Internet Research, 23(2).

Shickel, B., Loftus, T. J., Adhikari, L., Ozrazgat-Baslanti, T., Bihorac, A., & Rashidi, P. (2019). DeepSOFA: A Continuous Acuity Score for Critically Ill Patients using Clinically Interpretable Deep Learning. Scientific Reports, 9(1).

Tan, Q., Ye, M., Yang, B., Liu, S., Andy, J., Yip, T. C., Wong, G. L., & Yuen, P. C. (2020). DATA-GRU: Dual-Attention Time-Aware Gated Recurrent Unit for Irregular Multivariate Time Series. Proceedings of the AAAI Conference on Artificial Intelligence, 34(1), 930–937.

Thorsen-Meyer, H., Placido, D., Kaas-Hansen, B. S., Nielsen, A. P., Lange, T., Nielsen, A. B., Toft, P., Schierbeck, J., Strøm, T., Chmura, P. J., Heimann, M., Belling, K., Perner, A., & Brunak, S. (2022). Discrete-time survival analysis in the critically ill: a deep learning approach using heterogeneous data. Npj Digital Medicine, 5(1).

Tipirneni, S., & Reddy, C. K. (2022). Self-Supervised Transformer for Sparse and Irregularly Sampled Multivariate Clinical Time-Series. ACM Transactions on Knowledge Discovery from Data, 16(6), 1–17.

Tonekaboni, S., Joshi, S., McCradden, M. D., & Goldenberg, A. (2022). What Clinicians Want: Contextualizing Explainable Machine Learning for Clinical End Use. arXiv (Cornell University).

Weerasekara, S., Kamarthi, S., & Isaacs, J. (2025). Prediction of Clinical Complication Onset using Neural Point Processes. arXiv (Cornell University).

Wu, A., Choudhary, T., Upadhyaya, P., Ali, A., Yang, P., & Kamaleswaran, R. (2024). Deep Representation Learning-Based Dynamic Trajectory Phenotyping for Acute Respiratory Failure in Medical Intensive Care Units. arXiv (Cornell University).

Xu, Y., Biswal, S., Deshpande, S. R., Maher, K., & Sun, J. (2022). RAIM: Recurrent Attentive and Intensive Model of Multimodal Patient Monitoring Data. arXiv (Cornell University).

Yadav, S., & Subbian, V. (2024). Failure Modes of Time Series Interpretability Algorithms for

Critical Care Applications and Potential Solutions. PubMed, 2024, 1424–1433.

Yadav, S., & Subbian, V. (2025). Failure Modes of Time Series Interpretability Algorithms for Critical Care Applications and Potential Solutions. arXiv (Cornell University).

Yang, H., Li, J., Zhang, C., Pazos, A., & Shen, B. (2025). Predictive model for daily risk alerts in sepsis patients in the ICU: visualization and clinical analysis of risk indicators. Precision Clinical Medicine, 8(1).

Yip, S. (2026). Time-Series ICU Patient Deterioration Predictor. In Zenodo (CERN European Organization for Nuclear Research). European Organization for Nuclear Research.

Yoon, H., Kim, B. R., Kim, H. S., Park, D. K., Kim, H. S., Cho, H.-Y., Lee, H., & Lee, H. (2025). Multicenter validation of a scalable, interpretable, multitask prediction model for multiple clinical outcomes. Npj Digital Medicine, 8(1).

Yoon, J. H., Jeanselme, V., Dubrawski, A., Hravnak, M., Pinsky, M. R., & Clermont, G. (2020). Prediction of hypotension events with physiologic vital sign signatures in the intensive care unit. Critical Care, 24(1).

Yoosoofsah, A. (2024). Predicting Extubation Failure in Intensive Care: The Development of a Novel, End-to-End Actionable and Interpretable Prediction System. arXiv (Cornell University).

Zhang, S., Ding, H., Shen, Y., Shao, B., Gu, Y.-Y., Chen, Q., Zhang, H., Pei, Y., & Jiang, H. (2025). Harness machine learning for multiple prognoses prediction in sepsis patients: evidence from the MIMIC-IV database. BMC Medical Informatics and Decision Making, 25(1).

Zhang, Y. (2025). Improving Automatic Clinical Decision Support System with Advanced Computational Methods. Deep Blue (University of Michigan).

Zhang, Z., & Ye, J. (2025). Predicting mortality in critically ill patients with hypertension using machine learning and deep learning models. Frontiers in Cardiovascular Medicine, 12.