Volume 10, Issue 6, June – 2025 International Journal of Innovative Science and Research Technology
ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25jun278

Real-Time Multivariate Vital Sign Forecasting in Intensive Care Units: A Comparative Study of Machine Learning Models with Emphasis on Time Series Mixer for Early Warning Systems

Manish Raj[1]; Goalla Kartheek[2]; Arikatla Sai Sumedha[3]; Chintala Guru Sunanda[4]; Bhagyashree G Malipatil[5]; Poojitha S[6]; Praveen Kumar Burra[7]; Bharani Kumar Depuru[8]

[1;2;3;4;5;6;7;8]Aispry

Publication Date: 2025/06/11

Abstract: Real-time monitoring and prediction of physiological parameters in critical care settings remains essential for preventing patient deterioration and enabling timely medical interventions. This research examines various computational approaches for short-term prediction of six critical vital signs: Heart Rate (HR), Systolic/Diastolic Blood Pressure (SBP/DBP), Respiratory Rate (RR), Oxygen Saturation (SpO2), and Temperature (Temp) using the VitalDB clinical database. Our investigation began with conventional time-series methods including Autoregressive Integrated Moving Average (ARIMA) models and progressed to sophisticated neural architectures such as Long Short-Term Memory (LSTM) networks. However, these approaches demonstrated limitations in modeling complex relationships between multiple physiological variables simultaneously. Subsequently, we implemented advanced hybrid architectures incorporating Bidirectional LSTM (BiLSTM) layers, Convolutional Neural Networks (CNN), and Graph Attention Network (GAT) mechanisms. Although this hybrid model achieved enhanced prediction accuracy, their computational complexity posed challenges for clinical deployment. Addressing practical implementation requirements, we evaluated Multi-Layer Perceptron (MLP)-based frameworks, specifically Patch Time Series Transformer (PatchTST) and Time Series Mixer (TSMixer) architectures. PatchTST effectively captures extended temporal dependencies but lacks comprehensive cross-variable interaction modeling. Conversely, TSMixer employs dual mixing mechanisms—temporal and feature-based—to simultaneously learn chronological patterns and inter-vital correlations. Utilizing 10-minute historical windows to forecast subsequent 3-minute intervals, TSMixer demonstrated superior performance across all evaluation metrics. The model achieved the lowest Root Mean Square Error (RMSE) values for all vital parameters while maintaining computational efficiency suitable for real-time applications. These findings establish TSMixer's potential as a practical solution for prospective integration into Intensive Care Unit (ICU) monitoring systems, offering both predictive accuracy and operational feasibility for clinical environments.

Keywords: *Intensive Care Unit, Vital Signs Forecasting, ARIMA, LSTM, BiLSTM, CNN, GAT, PatchTST, TSMixer, Transformer Models, Multivariate Forecasting, Real-Time Health Monitoring, Deep Learning, Clinical Decision Support*.

How to Cite: Manish Raj; Goalla Kartheek; Arikatla Sai Sumedha; Chintala Guru Sunanda; Bhagyashree G Malipatil; Poojitha S; Praveen Kumar Burra; Bharani Kumar Depuru (2025) Real-Time Multivariate Vital Sign Forecasting in Intensive Care Units: A Comparative Study of Machine Learning Models with Emphasis on Time Series Mixer for Early Warning Systems. *International Journal of Innovative Science and Research Technology*, 10(6), 134-140. https://doi.org/10.38124/ijisrt/25jun278

I. INTRODUCTION

Healthcare systems around the world have gone through major changes in recent years, mainly because of new advances in digital technology and Artificial Intelligence (AI) [3]. One of the biggest changes is how AI and machine learning are now being used more in hospitals and clinics. These new technologies are changing how healthcare works by making diagnoses faster, making treatments more personal for each patient, and helping predict medical problems before they happen [3]. AI is now being used in many areas like medical imaging, drug development, and checking patient risks, making it a very important part of modern medicine [3].

The use of AI in healthcare started many years ago with simple rule-based computer systems and basic statistical

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25jun278

models. While these early methods helped doctors make decisions in a structured way, they were too rigid to handle the complex situations that happen in real medical settings. Today's healthcare world has changed a lot because of Electronic Health Records (EHRs), wearable devices that monitor patients, and systems that watch patients all the time [4]. This has created huge amounts of medical data. Modern AI and machine learning methods, especially deep learning, are now being used to study all this data and find patterns that older methods might miss [4].

This technology improvement is especially important in ICUs, where patients can get worse very quickly and without warning [4]. Even small changes in basic body functions like heart rate, blood pressure, or oxygen levels can be early signs of serious medical problems. Being able to spot and understand these small changes early is very important for preventing medical emergencies and saving patients' lives. However, most ICU monitoring systems today work by reacting to problems, using alarms that only go off when vital signs cross dangerous levels. These alarm systems often create false alerts, which leads to alarm fatigue among doctors and nurses and sometimes causes delays in responding to real medical emergencies.

Moving from reactive to proactive healthcare can be greatly improved by using AI-powered forecasting systems. These smart analytical models use past patient data to predict how vital signs might change in the future. This prediction ability could potentially give doctors and nurses important early warnings, allowing them to take action before patient conditions become critical and potentially preventing bad outcomes.

Also, integrating AI forecasting technologies into ICU work routines could provide many important benefits that go beyond just making predictions. These systems could help medical staff understand patient situations better by giving them complete real-time information about patient status trends. They could also reduce the mental workload on healthcare workers by automatically doing complex pattern recognition tasks that would otherwise need intensive manual work. Additionally, these technologies could support evidence-based clinical decisions by providing objective, data-driven insights that work together with clinical expertise. The implementation could also help improve how resources are used, making sure that medical staff and equipment are used most effectively across patient care situations. Through these many potential benefits, AI technologies could not only help existing human medical expertise but also help ensure that critical care delivery stays both highly responsive to immediate needs and ready for future requirements.

The development of an effective multivariate forecasting system specifically designed for ICU vital sign monitoring requires adherence to a systematic machine learning development pipeline. This research adopts a structured approach inspired by the Cross Industry Standard Process for Machine Learning with Quality Assurance CRISP-ML(Q) methodology, which represents a comprehensive, iterative framework specifically designed for developing high-quality AI solutions in critical application domains. This methodological approach seamlessly integrates essential components including thorough clinical problem analysis and understanding, comprehensive data processing and preparation, systematic model development and optimization, rigorous performance evaluation protocols, and extensive quality assurance procedures. The complete framework is readily available as an open-source resource through the 360DigiTMG website [Fig. 1].

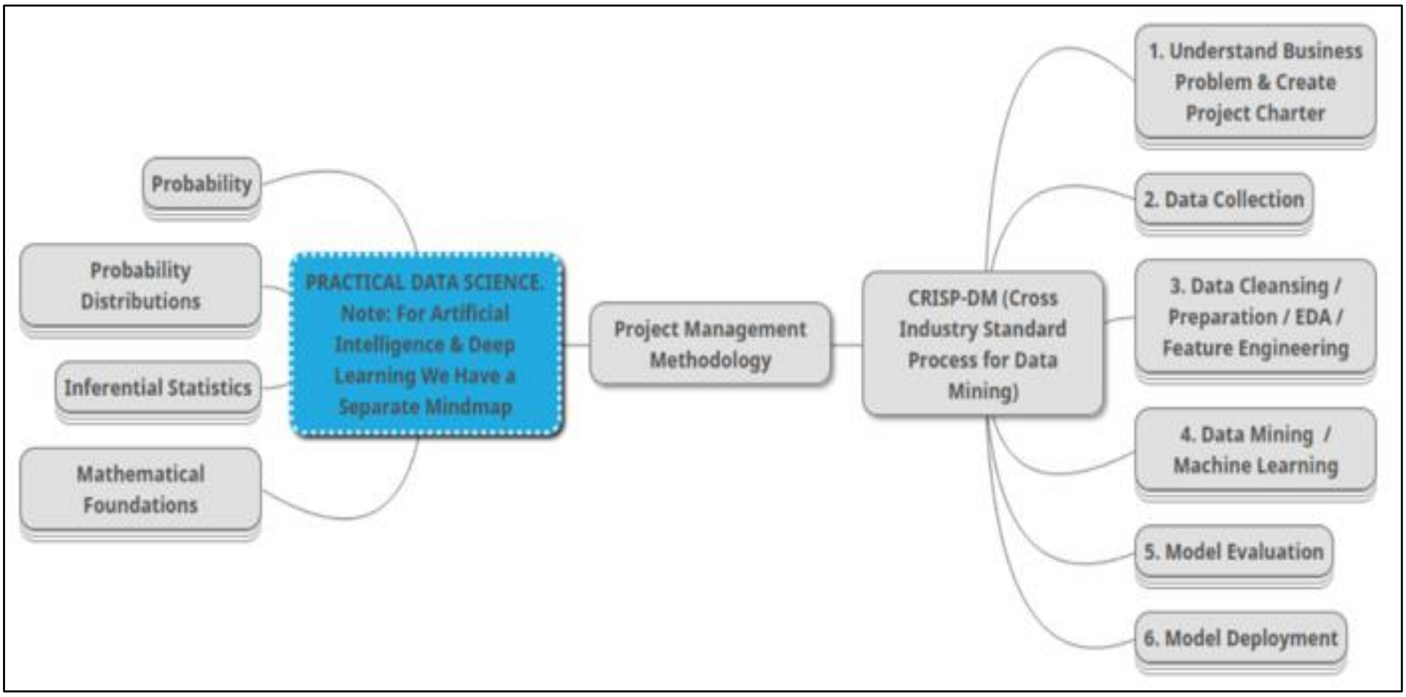

Fig 1 CRISP-ML(Q) - Approach for Quality Assurance across all Six Phases.
(Source: Mind Map - 360DigiTMG)

The fundamental objective of this research initiative centers on enhancing short-term forecasting accuracy for critical physiological parameters that are essential for patient monitoring and care. These parameters encompass HR measurements, both SBP and DBP readings, RR monitoring, SpO_2 levels, and Tem measurements. The improved forecasting capabilities are designed to potentially support real-time ICU monitoring systems and facilitate early detection of patient deterioration patterns. The study implements a comprehensive evaluation framework that systematically assesses multiple forecasting model architectures, including traditional statistical approaches such as ARIMA, sophisticated recurrent neural network architectures including LSTM, and cutting-edge modern MLP-based architectures such as PatchTST and TSMixer models.

Through systematic comparison of these diverse modeling approaches across multiple evaluation criteria including predictive accuracy, model interpretability, and computational efficiency requirements, the research successfully identifies the optimal solution for potential future implementation within critical care environments. Among all evaluated architectures, the TSMixer model demonstrates the most promising performance characteristics, effectively capturing both complex temporal dynamics inherent in physiological data and intricate inter-vital sign dependencies, while simultaneously maintaining low inference latency that would be essential for real-time clinical applications.

The comprehensive workflow methodology, as depicted in [Fig 2], follows a carefully structured pipeline that begins with thorough Clinical and Data Understanding phases to precisely define the research problem and identify relevant physiological parameters for analysis. This foundational phase is succeeded by systematic Data Collection procedures utilizing open-source ICU datasets, followed by comprehensive Data Preparation processes that encompass data cleaning protocols, normalization procedures, and sequence generation methodologies. The Model Building phase involves the systematic implementation and optimization of diverse forecasting models including ARIMA, LSTM, hybrid model (CNN+BiLSTM+GAT) [1] [5], PatchTST, and TSMixer architectures through careful hyperparameter tuning processes. Model Evaluation procedures are conducted using key performance metrics, particularly RMSE measurements, to provide comprehensive assessment of model performance capabilities. The research concludes with the identification and validation of the optimal modeling approach for potential future ICU applications, determined through careful analysis of both predictive accuracy and computational inference requirements.

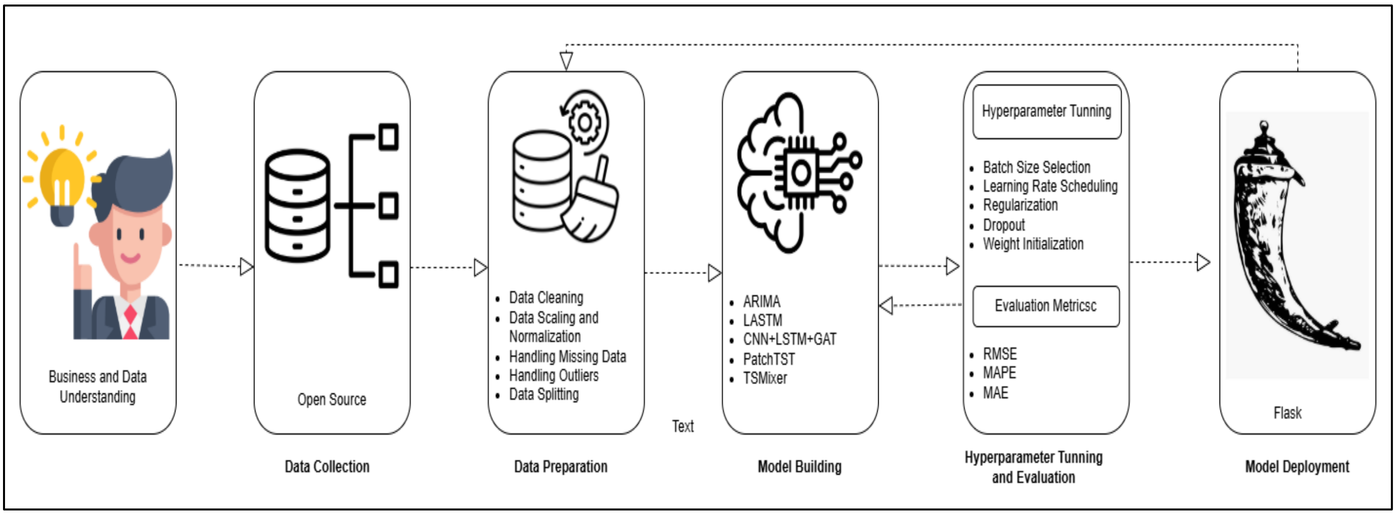

Fig 2 Project Architecture Diagram – A Structured Workflow from Clinical Understanding to Model Evaluation for Real-Time ICU Vital Signs Forecasting.
(Source: Machine Learning Workflow – 360DigiTMG)

II. METHODS AND TECHNIQUES

➢ *Data Collection:*

This study leverages the VitalDB dataset an open-source, high-resolution database curated by Seoul National University Hospital and specifically designed to support research in critical care and time-series analysis. VitalDB comprises detailed physiological waveform and numeric data collected from ICU patients, often during surgeries or intensive monitoring scenarios. The dataset is publicly available and has been widely adopted in studies involving ICU patient monitoring, early warning systems, and machine learning-based health analytics.

For the purpose of this research, we concentrate on six vital physiological parameters that are routinely tracked in ICU settings and serve as crucial indicators of a patient's clinical condition: HR, SBP, DBP, RR, SpO_2, Temp.

➢ *Data Preprocessing:*

The dataset extracted from VitalDB consists of high-resolution time-series recordings sampled at a fixed interval of two seconds. However, upon initial inspection, several challenges were noted that necessitated thorough preprocessing to prepare the data for modeling. The raw data contained numerous empty rows, where all physiological measurements were missing likely due to signal dropouts or

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25jun278

temporary sensor failures. These rows were systematically removed to ensure the continuity and completeness of the time series.

Additionally, the original time sequence was irregular and inconsistent due to prior trimming or data loss. To resolve this, the time column was reset to form a uniform sequence representing consistent time progression, facilitating temporal alignment across multiple vital sign channels. Beyond structural inconsistencies, the dataset included several anomalous readings, such as negative values or numerically implausible entries that fall outside clinically acceptable ranges. These were treated as outliers and were replaced with missing values (NaNs) to avoid skewing the model.

To handle missing values introduced during this cleaning process, a forward-fill (also known as "last observation carried forward") imputation strategy was employed. This method assumes that vital signs do not fluctuate erratically from one second to the next and is appropriate in ICU time-series contexts where trends generally evolve gradually. Finally, to increase the temporal resolution and allow finer-grained forecasting, the data was resampled and interpolated from its original 2-second interval to a 1-second interval. These preprocessing steps collectively ensured that the dataset was cleaned, regularized, and suitable for multivariate forecasting using deep learning models.

➢ *Model Training:*

After cleaning the dataset, the modeling journey began with traditional statistical and deep learning baselines and progressively evolved toward advanced neural architectures tailored for the complex nature of multivariate ICU time-series data.

The initial approach utilized the ARIMA model due to its established effectiveness in univariate time-series forecasting. Despite its simplicity, ARIMA revealed several critical shortcomings when applied to ICU data. Its assumptions of stationarity and linearity are incompatible with the highly dynamic and nonlinear characteristics of physiological signals. ICU data often contain abrupt changes due to clinical interventions or sudden deteriorations, which ARIMA is ill-equipped to handle. Moreover, ARIMA models lack the capability to process multiple time-dependent variables simultaneously, which is essential when working with interdependent vitals like heart rate, blood pressure, and oxygen saturation. The need for manual differencing and parameter tuning further limited its utility in a scalable ICU forecasting system.

Recognizing these limitations, the next phase involved transitioning to LSTM networks. LSTM, known for its ability to capture long-term dependencies in sequential data, offered a more flexible alternative to ARIMA. In early tests, LSTM models showed improved performance in forecasting vital sign trends over time. However, when trained individually on each vital sign, LSTM treated each time series in isolation, failing to utilize the rich inter-variable relationships inherent in physiological data. In practice, this meant the model could identify how a single signal changed over time but was blind to causal or correlated shifts between vitals for instance, how a rising heart rate might correspond with a drop in blood pressure. Furthermore, LSTM models often underperformed in handling the high variability and noise levels typical of ICU data, particularly under conditions involving rapid physiological changes.

To bridge this gap, a hybrid model combining BiLSTM, CNN, and Graph Attention Networks (GAT) was developed [1] [5]. This architecture aimed to comprehensively model temporal patterns and inter-variable interactions. The BiLSTM component was responsible for learning long-term trends from both past and future contexts, enabling better detection of gradual physiological changes [1]. In parallel, the Convolutional Neural Network (CNN) focused on short-term temporal windows to detect quick shifts, such as sudden desaturations or heart rate spikes [1]. The inclusion of GAT allowed the model to dynamically learn how different vital signs influence one another using attention scores, rather than relying on predefined physiological rules. This made the model more adaptive and data-driven [5].

While the hybrid model produced promising results with low RMSE values, several practical drawbacks emerged. It often struggled to maintain trend consistency across test samples and demonstrated performance variability across different patient profiles. The GAT module, while beneficial for learning cross-variable interactions, added considerable computational overhead, which became a bottleneck for real-time clinical deployment. Moreover, the requirement to construct input graphs on-the-fly for every prediction introduced additional complexity and latency, making the solution less feasible in high-throughput ICU environments.

In response to these challenges, the modeling strategy pivoted toward more efficient, scalable architectures based on MLPs, with a focus on two models: PatchTST and TSMixer [7] [8]. The PatchTST model was first selected for its excellent track record in long-sequence forecasting tasks [7]. It operates by segmenting the input sequence into fixed-size temporal patches and applying linear projections to extract temporal features [7]. However, despite its simplicity and computational efficiency, PatchTST treated each variable independently, lacking any built-in mechanism to model inter-variable dependencies. This limitation was particularly critical in an ICU setting, where understanding how vitals influence one another is key to accurate prediction and early warning.

To address this, the TSMixer model was introduced—a novel, MLP-only architecture inspired by the MLP-Mixer framework from computer vision [9]. TSMixer is explicitly designed to handle multivariate time-series data by jointly learning temporal dynamics and feature-level relationships [9]. Its architecture comprises alternating layers of two key operations: Time Mixing and Feature Mixing. In the Time Mixing block, MLPs are applied across the time dimension for each individual vital sign, enabling the model to capture both short-term fluctuations and long-range dependencies. In the Feature Mixing block, MLPs operate across the feature

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25jun278

dimension (i.e., across different vital signs) at each time step, allowing the model to learn complex interdependencies between variables such as heart rate, SpO_2, and blood pressure.

Each MLP block in TSMixer is followed by residual connections and layer normalization, which contribute to stable gradient flow and improved generalization. Unlike transformer-based models, TSMixer avoids attention mechanisms, making it significantly lighter and faster to train, while still maintaining the ability to model sophisticated relationships in the data. This architecture proved especially effective in capturing the nuanced interplay of ICU vitals and forecasting their future trajectories with high accuracy and efficiency, making it a strong candidate for real-time clinical applications.

➢ *Model Evaluation:*

Table 1 RMSE Comparison of Models Across Vital Signs

Model	HR	SpO_2	SBP	DBP	RR	Temp
Hybrid Model (CNN + BiLSTM + GAT)	4.49	1.49	3.92	2.08	6.41	0.32
PatchTST	6.5948	3.9663	5.6873	3.8954	1.6291	0.1393
TSMixer	2.4593	0.6810	2.3139	1.5834	0.7511	0.0579

Evaluating predictive models represents a crucial step in their development and deployment, as it validates their dependability and precision when applied to real-world applications, especially in high-stakes applications such as ICU vital sign forecasting. In this study, we evaluated the forecasting accuracy of three different models—Hybrid (CNN + BiLSTM + GAT) [1] [5] [6], PatchTST, and TSMixer—across six physiological parameters: HR, SpO_2, SBP, DBP, RR, and Temp. RMSE was chosen as the primary evaluation metric due to its sensitivity to large deviations, which is vital in clinical monitoring scenarios.

The Hybrid model, combining CNN, BiLSTM, and Graph Attention Networks, demonstrated reasonable accuracy across most vitals, particularly in modeling short- and long-term temporal dependencies and learning inter-variable relationships [1] [5] [6]. The model achieved an RMSE of 4.49 for Heart Rate, 1.49 for SpO_2, 6.41 for Respiratory Rate, 3.92 for Systolic BP, 2.08 for Diastolic BP, and 0.32 for Tem. While these results indicated the model's effectiveness, especially in SpO_2 and Temperature prediction, its performance on Respiratory Rate and Heart Rate showed room for improvement, suggesting difficulties in trend-following for rapidly fluctuating signals.

PatchTST model was tested next, leveraging its capability for long-sequence forecasting through temporal patch extraction [7]. However, this model treated each physiological signal independently and did not incorporate inter-signal dependencies, which impacted its overall performance. The PatchTST model reported higher RMSEs in most categories compared to the hybrid model: 6.5948 for Heart Rate, 3.9663 for SpO_2, 5.6873 for Systolic BP, 3.8954 for Diastolic BP, 1.6291 for Respiratory Rate, and 0.1393 for Temperature. Although it performed relatively well in Temperature and Respiratory Rate forecasting, the model struggled significantly with cardiovascular parameters due to its univariate treatment of input features.

In contrast, the TSMixer model, inspired by MLP-Mixer architecture, demonstrated superior forecasting performance across all vital signs by jointly modeling temporal patterns and feature-wise dependencies [8] [9]. This model achieved the lowest RMSEs among all three models tested: 2.4593 for Heart Rate, 0.6810 for SpO_2, 2.3139 for Systolic BP, 1.5834 for Diastolic BP, 0.7511 for Respiratory Rate, and an exceptionally low 0.0579 for Temperature. The TSMixer's architecture, which incorporates time-mixing and feature-mixing blocks, allowed it to capture both intra-variable temporal dynamics and inter-variable physiological relationships more effectively than the other models.

These results clearly indicate that the TSMixer model outperformed both the Hybrid and PatchTST models in terms of accuracy and consistency. Its ability to forecast multivariate clinical signals with lower prediction error makes it a highly promising approach for real-time ICU monitoring systems. Table I presents the RMSE-based performance comparison of three forecasting models - Hybrid (CNN + BiLSTM + GAT), PatchTST, and TSMixer—across six critical ICU vital signs, highlighting each model's predictive accuracy and robustness in multivariate physiological signal forecasting.

III. RESULTS AND DISCUSSION

This study evaluated the forecasting performance of three models—Hybrid (CNN + BiLSTM + GAT), PatchTST, and TSMixer—across six critical ICU vital signs: HR, SpO_2, SBP, DBP, RR, and Temp. The RMSE was used as the primary evaluation metric. Among the three models, TSMixer demonstrated superior performance, achieving the lowest RMSE values across all variables: 2.46 (HR), 0.68 (SpO_2), 2.31 (SBP), 1.58 (DBP), 0.75 (RR), and 0.06 (Temp). The Hybrid model exhibited moderate performance, particularly performing well on SpO_2 (1.49) and Temp (0.32), but underperformed on dynamic signals such as HR and RR. PatchTST yielded the highest RMSEs, notably for HR (6.59) and SpO_2 (3.97), indicating limited effectiveness in this multivariate ICU setting.

The experimental results clearly highlight the efficacy of TSMixer in forecasting multivariate ICU vital signs. TSMixer's architecture leverages both temporal and feature mixing mechanisms, enabling it to effectively learn intricate interdependencies among multiple vital parameters. This structural advantage is particularly relevant in an ICU context where patient vitals are not only interdependent but also highly dynamic.

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25jun278

While the Hybrid model integrated CNN for spatial feature extraction, BiLSTM for temporal sequence learning, and GAT for graph-based relationship modeling, its increased architectural complexity did not translate into superior forecasting accuracy. Conversely, PatchTST, being primarily tailored for univariate long-term forecasting, failed to capture the multidimensional relationships essential for accurate ICU predictions.

From a clinical standpoint, TSMixer's consistent and accurate performance makes it a promising candidate for real-time deployment in ICU monitoring systems. Accurate forecasting of vital signs allows for anticipatory decision-making, where clinicians can identify deteriorating trends and intervene proactively. Such predictive capabilities are crucial in critical care, where early detection and timely intervention can significantly influence patient outcomes

IV. CONCLUSION

This study demonstrated that among the three evaluated models, TSMixer emerged as the most robust and efficient for multivariate ICU vital sign forecasting. Its superior accuracy, minimal computational overhead, and suitability for real-time applications underscore its potential for integration into clinical monitoring systems. The comparative analysis also revealed the limitations of traditional hybrid models and univariate time-series methods in capturing complex physiological dynamics in critical care settings.

The findings from this research lay a strong foundation for the development of intelligent ICU systems. As the field advances, incorporating more sophisticated forecasting models will be essential for supporting real-time clinical decisions and enhancing patient outcomes in intensive care environments.

FUTURE SCOPE

While the current study focused on accurately forecasting ICU vital signs, future research will explore predictive modeling of life-threatening clinical conditions based on the forecasted trends. These conditions such as respiratory failure, cardiac arrest, septic shock, and early-stage sepsis. By integrating multivariate time-series forecasting with disease classification models and clinical rule-based inference, the goal is to create a comprehensive early warning system capable of identifying patient deterioration before it becomes critical.

In subsequent developments, additional patient data sources—such as laboratory results, and medication history will be incorporated to enhance the system's contextual understanding. This will enable the generation of personalized early warning scores and prescriptive recommendations, empowering ICU staff with timely, accurate, and actionable insights. The long-term objective is to build a real-time Clinical Decision Support System (CDSS) that not only monitors and forecasts patient vitals but also assists in preventing life-threatening complications through proactive intervention strategies.

ACKNOWLEDGMENTS

We gratefully acknowledge that, with permission from 360DigiTMG, we have utilized the CRISP-ML(Q) methodology (ak.1) and the Machine Learning Workflow (ak.2), both of which are openly accessible as part of the open-source resources provided on the official 360DigiTMG website.

REFERENCES

[1]. Jilin Zhang, Lishi Ye, Yongzeng Lai. 2023. Stock Price Prediction Using CNN-BiLSTM-Attention Model. https://www.mdpi.com/2227-7390/11/9/1985.

[2]. Oliver Chojnowski, Dario Luipers, Caterina Neef, Anja Richert. 2023. Forecasting Vital Signs in Human–Robot Collaboration Using Sequence-to-Sequence Models with Bidirectional LSTM: A Comparative Analysis of Uni- and Multi-Variate Approaches. https://www.mdpi.com/2673-4591/58/1/103

[3]. Michael R. Pinsky, Armando Bedoya, Azra Bihorac, Leo Celi, Matthew Churpek, Nicoleta J. Economou-Zavlanos, Paul Elbers, Suchi Saria, Vincent Liu, Patrick G.Lyons, Benjamin Shickel, Patrick Toral, David Tscholl, Gilles Clermont. 2024. Use of artificial intelligence in critical care: opportunities and obstacles. https://ccforum.biomedcentral.com/articles/10.1186/s13054-024-04860-z

[4]. Aswin Thacharodi, Prabhakar Singh, Ramu Meenatchi, Z H Tawfeeq Ahmed, Rejith R S Kumar, Neha V, Sanjana Kavish, Mohsin Maqbool, Saqib Hassan. 2024. Revolutionizing healthcare and medicine: The impact of modern technologies for a healthier future—A comprehensive review. https://pmc.ncbi.nlm.nih.gov/articles/PMC11520245/

[5]. Ugochukwu Orji, Cicek Guven, Dan Stowell. 2025. Enhanced Load Forecasting with GAT-LSTM: Leveraging Grid and Temporal Features. https://arxiv.org/abs/2502.08376

[6]. Tianlong Wu, Feng Chen, Yun Wan. 2018. Graph Attention LSTM Network: A New Model for Traffic Flow Forecasting. https://ieeexplore.ieee.org/document/8612556

[7]. Hussein Ahmad Ahmad, Seyyed Kasra Mortazavi, Mohamed El Bahnasawi, Fadi Al Machot, Witesyavwirwa Vianney Kambale, Kyandoghere Kyamakya. 2024. Enhanced Time Series Forecasting: Integrating PatchTST with BERT Layers. https://ieeexplore.ieee.org/document/10771354

[8]. Si-An Chen, Chun-Liang Li, Nate Yoder, Sercan O. Arik, Tomas Pfister. 2023. TSMixer: An All-MLP Architecture for Time Series Forecasting. https://arxiv.org/abs/2303.06053

[9]. Vijay Ekambaram, Arindam Jati, Nam Nguyen, Phanwadee Sinthong, Jayant Kalagnanam. 2023. TSMixer: Lightweight MLP-Mixer Model for Multivariate Time Series Forecasting. https://arxiv.org/abs/2306.09364

ISSN No:-2456-2165 https://doi.org/10.38124/ijisrt/25jun278

[10]. Haowei Ni, Shuchen Meng, Xieming Geng, Panfeng Li, Zhuoying Li, Xupeng Chen, Xiaotong Wang, Shiyao Zhang. 2024. Time Series Modeling for Heart Rate Prediction: From ARIMA to Transformers. https://arxiv.org/abs/2406.12199