International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

Bridging Sensor Telemetry, Explainable Risk Models, and Clinical Workflows in ICU Decision Support

1st Mr. Atharv Khot
Department of Computer Engineering
D. Y. Patil Education Society (Deemed to be University)
School of Engineering and Management
Kolhapur, India

2nd Ms. Ishwari Kadam
Department of Computer Engineering
D. Y. Patil Education Society (Deemed to be University)
School of Engineering and Management
Kolhapur, India

3rd Ms. Rajlaxmi Ghatage
Department of Computer Engineering
D. Y. Patil Education Society (Deemed to be University)
School of Engineering and Management
Kolhapur, India

4th Ms. Priyanka Bote
Department of Computer Engineering
D. Y. Patil Education Society (Deemed to be University)
School of Engineering and Management
Kolhapur, India

5th Mr. Satish Madhukar Ranbhise
Department of Computer Engineering
D. Y. Patil Education Society (Deemed to be University)
School of Engineering and Management
Kolhapur, India

Abstract **- Patients in the Intensive Care Unit (ICU) can get worse very quickly, so keeping a close eye on their vitals and spotting risk early really matters. Many AI and machine learning models work well on past hospital records, yet most of them stay stuck in research instead of becoming part of real day-to-day ICU work. A lot of current tools only use old datasets, run as separate demos where staff have to type numbers in by hand, depend on too many variables, or give scores that are hard to trust or fit into busy hospital routines. This paper looks at research across three parts of ICU decision support: live vital-sign tracking from sensors, risk predictions that can be explained clearly, and how those alerts and insights fit into a doctor's or nurse's workflow. By bringing together studies on patient sensors, simpler feature sets, tree-based risk models, and explanation tools like SHAP, we point out what still gets in the way of real-world use. We also share a clear system layout that links vital-sign streams, explainable risk scores, and role-based alerts to support medical staff not replace doctors or make automated diagnoses.**

Index Terms **- Clinical Decision Support, Explainable AI, Intensive Care Unit, Mortality Risk, Patient Monitoring, Sensor Telemetry, SHAP, XGBoost.**

I. INTRODUCTION

Patients in the Intensive Care Unit (ICU) can get worse in minutes, and catching early warning signs quickly matters a great deal. In everyday ICU care, doctors and nurses watch vital signs like heart rate, respiratory rate, blood oxygen (SpO_2), and temperature closely to spot trouble early. Machine learning tools can already find useful risk patterns in hospital data, yet a few practical problems keep most of them out of real-world bedside routines

Many existing studies only test models on older, static records, rely on heavy lab inputs that aren't always ready at once, or sit in separate tools where staff have to enter numbers by hand. Even when a model predicts higher risk, a raw score without a clear reason is tough for a busy

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

clinician to trust in the middle of rounds. Too many noisy or poorly timed alerts only add to alarm fatigue.

This work looks at connecting three parts of ICU care: live vital-sign telemetry from sensors, risk predictions that explain what is driving them, and clear clinical workflows for doctors and nurses. It is meant to support medical staff not replace clinical judgment or make automated diagnoses. This paper reviews current research across these areas, points out where existing approaches fall short in practice, and outlines a clinician in the loop framework for ICU decision support.

A. Key Focus Areas of This Review

Existing ICU tools often treat sensor streams, risk scoring, and clinician alerts as separate pieces instead of one connected flow. This review focuses on:

1) *Live Vital-Sign Telemetry:* Tracking core patient vitals like heart rate, respiratory rate, SpO_2, and temperature continuously while handling missing or noisy readings honestly.
2) *Explainable Risk Modeling:* Pairing risk predictions (such as gradient-boosted models) with clear explanations (such as SHAP feature contributions) so staff can see *why* risk is rising.
3) *Clinical Workflow Fit:* Routing clear, timely decision support to the right clinical roles with proper alerts, acknowledgements, and logs rather than overwhelming staff.
4) *Safe Decision-Support Boundaries:* Keeping the system positioned as decision support for doctors and nurses never a replacement for human care or a certified medical device.

II. LITERATURE REVIEW

Most earlier work looks at one piece of ICU decision support at a time rather than connecting the full chain from bedside sensors to clinical action. On the monitoring side, continuous tracking of basic body signals such as heart rate, breathing rate, blood oxygen (SpO_2), and body temperature forms the foundation of ICU care [1]. Studies on sensor-based patient monitoring show that pairing lightweight sensors with microcontrollers makes live vital-sign tracking practical [2], though real-world streams still suffer from patient movement, loose wires, and missing data gaps that must be flagged rather than hidden [3].

On the modeling side, machine learning methods, especially tree-based algorithms like XGBoost and LightGBM, consistently outperform traditional ICU scoring tools such as SOFA and APACHE II on retrospective datasets like MIMIC-III and MIMIC-IV [4], [5]. Research on compact feature selection also shows that simpler models built on readily available vital signs can stay accurate without waiting for dozens of slow lab tests [6]. To make these predictions trustworthy, explainability methods such as SHAPley Additive explanations (SHAP) highlight which vital signs are pushing a patient's risk score up or down during clinical rounds [7].

Looking at clinical adoption, broader reviews of AI in critical care show that most models remain retrospective or run as standalone demos where staff must enter numbers manually [8]. Studies on model calibration further warn that probabilities trained in one hospital cohort can overpredict risk elsewhere unless adjusted locally [9], while unprioritized alarms lead to alarm fatigue among busy ICU staff [10]. Together, these findings suggest that live vital-sign telemetry, explainable risk scoring, and structured clinical workflows are each becoming practical on their own, yet their reliability and smooth coordination must be established before staff can depend on them in daily ICU routines.

Research Gap: Current systems tend to handle vital-sign telemetry, explainable risk prediction, and clinician alert workflows separately rather than around one connected care process. What is missing is a clinician in the loop framework that links live vital-sign streams with clear handling of missing or stale readings to compact, SHAP-explained risk models and auditable, role-based ICU workflows. Our proposed architecture addresses this gap by bringing continuous sensor telemetry, explainable risk scoring, and structured clinical alerting together as decision support for doctors and nurses.

III. SYSTEM ARCHITECTURE AND WORKFLOW

Our clinical decision-support system uses a structured workflow built around continuous bedside telemetry from ICU monitors (HR, RR, SpO2, and Temp). Clinicians, nurses, and supporting data stores work together around each patient's live stream rather than treating every reading in isolation. Instead of guessing missing numbers when sensors slip or raising noisy alarms on bad signals, the pipeline checks signal freshness and validity first, scores deterioration risk, and routes clear information and warnings to the care team.

A. Workflow Diagram

At a basic level, continuous vital-sign streams from ICU sensors are ingested and time-stamped, and the system checks whether the data is fresh and valid. When the signal is fresh and valid, it moves forward to extract vital and short-term trend features, uses XGBoost and SHAP to score deterioration risk and identify feature reasons, and checks whether risk exceeds the alert threshold. Readings above the threshold trigger a role-based alert for doctor or nurse review, while readings below threshold go to routine telemetry logging. When the signal is stale, noisy, or missing, the system flags a quality warning without guessing missing vitals and shows that warning in the decision view.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026



Fig. 1. ICU clinical decision-support workflow: from telemetry ingestion and quality checking through risk scoring, clinician review, and logging.

B. Implementation Architecture

The prototype divides its work across three tiers and supporting data stores. Tier 1 (Bedside Telemetry and Ingestion) collects HR, RR, SpO2, and Temp from vital sensors, timestamps and buffers the stream, and checks freshness and validity. Tier 2 (Feature and Risk Engine) extracts vital and short-term trend features, estimates deterioration risk with XGBoost, and generates signal-contribution reasons using SHAP. Tier 3 (Clinician Workflow and UI) sends role-based alerts to a clinician-in-the-loop decision view and logs responses through an acknowledgment handler, supported by a Telemetry Store for vital history, an Audit Log for review trails, and a Quality Warning store for bad or missing signal flags.

C. Data Flow

The operational data flow proceeds as follows:

1. Bedside sensors stream heart rate, respiratory rate, oxygen saturation, and temperature into the ingestion layer.
2. Stream ingestion attaches timestamps and buffers incoming readings so telemetry stays ordered.
3. The quality checker tests each stream for freshness and validity.
4. Valid streams move into feature extraction, risk scoring, and explanation generation.
5. Degraded streams trigger a quality warning without guessing or imputing missing vitals.
6. Alerts, quality warnings, and clinician acknowledgments are recorded in the audit log.

The system routes valid signals and problem readings along two parallel paths that meet in the audit log. On the valid telemetry path, clean, time-stamped vitals go to the feature builder to extract vital values and short-term trends, where XGBoost and SHAP calculate the risk score and signal reasons and send them through the alert router to the clinician decision view. On the quality warning path, missing, stale, noisy, or dropped signals raise a quality warning flag without guessing missing numbers, displaying a warning in the decision view for a bedside clinician check. Both paths log all alerts, warnings, and acknowledgments in the shared Audit Log for full traceability.

IV.IMPLEMENTATION METHODOLOGY

The implementation follows a fixed sequence of steps, chosen to keep the delay between telemetry and alert short and the prototype reliable.

A. Methodology Framework

The workflow moves through three main stages into a final validation summary. Stage 1 prepares and windows ICU telemetry streams (HR, RR, SpO2, and temperature). Stage 2 splits data by patient, models deterioration risk with XGBoost, and explains feature contributions with SHAP. Stage 3 evaluates predictive accuracy, system latency, and alert and quality-flag handling before compiling the final validation summary.

B. Execution Steps

1. *Telemetry Ingestion and Timestamping:* Bedside monitors stream HR, RR, SpO2, and Temp continuously, and the ingestion layer adds timestamps and buffers them in order.
2. *Quality Filtering and Synchronization:* Incoming readings are checked for freshness, noise, and dropouts; good data is time-aligned, while bad or missing signals trigger a quality warning instead of guessing missing vitals.
3. *Sliding Window Feature Extraction:* Clean vital streams are grouped into sliding windows to pull out current vital values and short-term trends showing recent patient changes.
4. *Patient-Level Split, Modeling, and SHAP Attribution:* Data is split by patient for fair testing, XGBoost estimates deterioration risk, and SHAP highlights which vital signs pushed that score up.
5. *Alert Routing and Clinician Review:* High-risk scores and SHAP reasons are sent to the clinician decision view for doctor or nurse review, while signal-quality warnings prompt a bedside sensor check.
6. *Audit Logging and Prototype Evaluation:* All alerts, quality warnings, and clinician acknowledgments are saved in the audit log, and the prototype is evaluated on accuracy (AUROC, AUPRC, sensitivity, specificity), latency, and alert/quality-flag handling.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

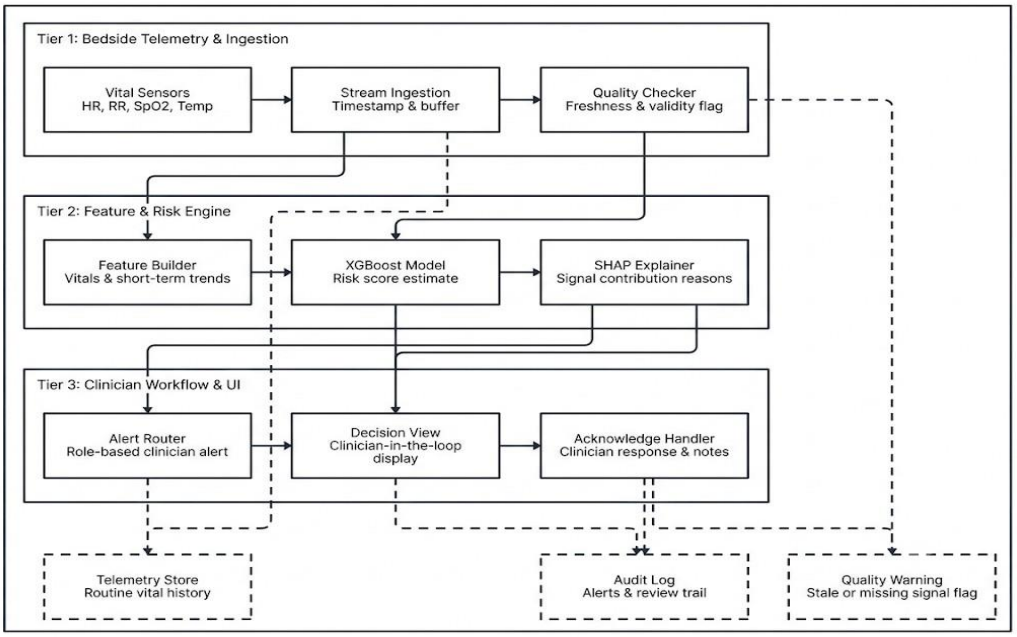


Fig. 2. Implementation architecture. Solid frames are system tiers; dashed boxes are supporting data stores.

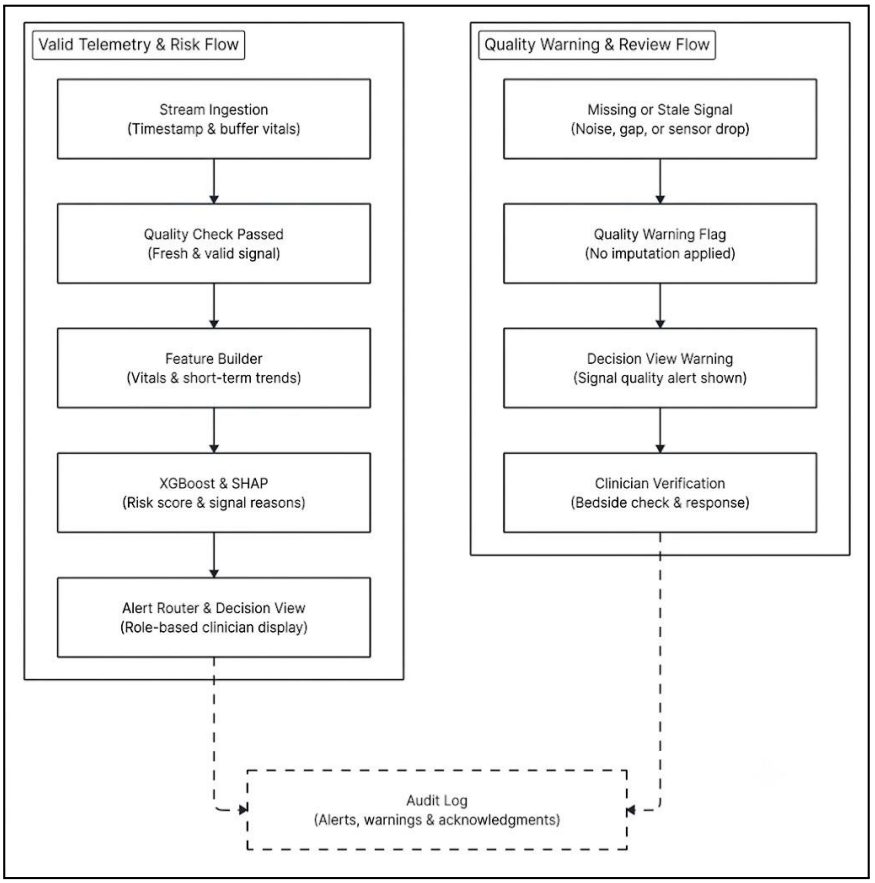


Fig. 3. Implementation data flow. Parallel paths handle valid telemetry risk scoring and signal-quality warnings, converging into the shared audit log.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 15 Issue 11 , November - 2026

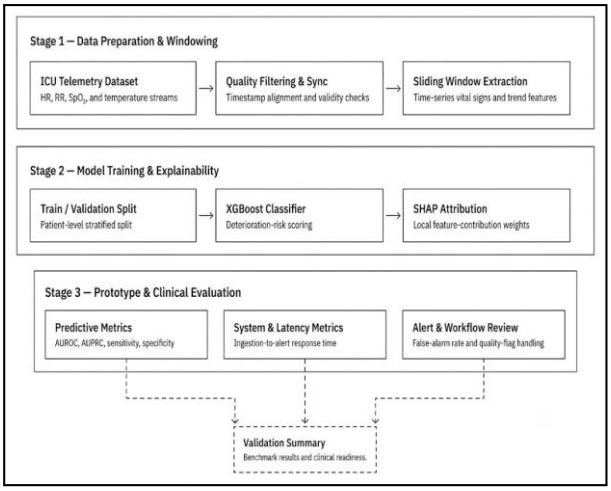


Fig. 4. Implementation methodology: three-stage end-to-end operational lifecycle.

TABLE I

PROPOSED EVALUATION CRITERIA FOR THE ICU DECISION-SUPPORT PROTOTYPE

Criterion	What is measured
Response Speed	Time from vital capture to risk score or quality warning in the UI.
Signal & Data Accuracy	Timestamp alignment of HR, RR, SpO2, and Temp, plus quality flags without imputation.
Deterioration Risk & Explainability	XGBoost risk accuracy across patient splits and SHAP feature reasons.
Alert & Audit Traceability	Routing of alerts, warnings, and acknowledgments into the Audit Log.
Usability	Ease of use and fit within daily ICU routines.

V. EXPECTED OUTCOMES AND DISCUSSION

We expect this prototype to deliver a working ICU decision-support tool that collects continuous vital signs (HR, RR, SpO2, and Temp) from bedside sensors, checks signal freshness and validity, and computes deterioration risk scores with XGBoost alongside SHAP feature explanations. Rather than guessing missing numbers through imputation when a sensor drops or becomes noisy, the system raises clear quality warnings so clinicians know when to check the bedside or sensor directly. It also routes role-based alerts and quality flags to the clinician view, captures clinician acknowledgments and notes, and records every event in a shared Audit Log. Overall, we expect this setup to keep alerts easier to understand, cut down on noisy or misleading cues, and support smoother team handoffs without adding extra clutter during busy ICU shifts.

A. Evaluation Criteria

Table I sets out the criteria we plan to use to evaluate the prototype.

VI. CONCLUSION

This prototype takes a clear, practical approach to helping ICU teams spot patient deterioration earlier without adding extra noise around the bed. Instead of treating every incoming reading as equal or guessing missing numbers when a sensor drops, it checks signal freshness first, flags doubtful data openly, and uses XGBoost with SHAP to show both the risk score and the specific vital signs or short-term trends behind it. Routing alerts by role, capturing clinician responses, and keeping everything in a shared Audit Log helps care teams stay aligned during shift handoffs and daily rounds.

This system is meant to support clinicians not replace their judgment or existing hospital monitoring workflows. By keeping alerts focused, explanations plain, and signal quality transparent, the goal is to make bedside decisions faster and more dependable. Our next step is to run the evaluation outlined above and report real measurements across all five criteria.

REFERENCES

1. N. Kudrot, Y. Si, J. Sanjaya, *et al.*, “First-24-hour machine learning for 30-day mortality prediction in ICU trauma patients: development in MIMIC-III and cross-database evaluation in MIMIC-IV,” *medRxiv*, Jul. 30, 2026.
2. Y. Pranata, R. Mirwanti, and Y. A. Bahtiar, “Translational potential and explainability of artificial intelligence-based clinical decision support for adults in intensive care: A scoping review,” *J. Multidiscip. Healthc.*, 2026.
3. M. Zarantonello *et al.*, “Artificial intelligence in intensive care units: a scoping review addressing the translational gap to clinical practice,” 2026.
4. T. Chen and C. Guestrin, “XGBoost: A scalable tree boosting system,” in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2016, pp. 785–794.
5. S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in *Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 4765–4774.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)