AI-Driven ICU Pharmacophysiologic Intelligence & Safety Platform

Real-time PK/PD concordance engine · Knowledge graph reasoning · LLM-synthesised recommendations · HITL safety gating

THE PROBLEM / Current Challenge-

Current Intensive Care Unit (ICU) monitoring systems separately track vital parameters—ECG waveforms, blood pressure, oxygen levels, and medication histories—without reasoning across them together.

This data remains in isolated, fragmented data silos.

Key Challenges:

Clinician Overload: Physicians and nurses operating under intense time pressure must manually synthesize this disconnected data, increasing cognitive strain.

Hidden Risks: Critical, patient-specific pharmacological or physiological interactions are frequently missed until they manifest as emergencies, compromising patient safety.

THE SOLUTION-

We are developing a comprehensive, AI-Driven ICU Safety and Intelligence Hub. This platform bridges the fragmented data gaps by natively integrating multiple data streams into a personalized, continuous physiological model.

This is achieved through four specialized engines:

Core Intelligent Processors:

1. Signal Intelligence: Continuous ECG waveform classification to detect subtle diagnostic shifts.
2. PK/PD Engine: Patient-specific Pharmacokinetic (PK) and Pharmacodynamic (PD) modeling to forecast personalized drug effects.
3. Knowledge Graph: A deterministic reasoning engine mapping mechanisms and interactions.
4. Verification Gate: A mandatory safety layer validating all recommendations.

Clinical Problem

📋 Illustrative ICU Scenario

Patient on norepinephrine + metoprolol + amiodarone . What does the monitor see?

58 MAP mmHg ↓ | **42** HR bpm ↓ | **92** SpO_2 % ↓ | ⚠ Bradycardia

Existing clinical monitoring lacks the capability to discriminate look-alike states: expected pharmacologic effects, dangerous drug–drug interactions, and underlying disease progression.

System Architecture -

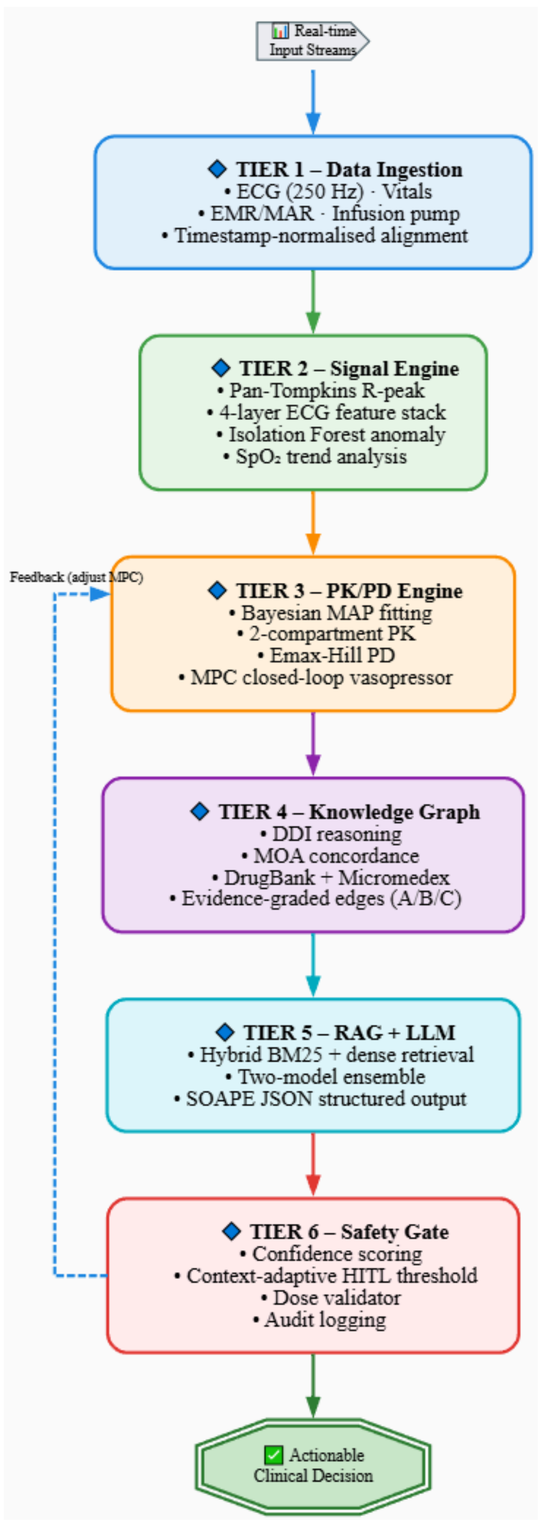

Concordance Engine-

1. ECG Signal Processing

- **Pan-Tompkins → real-time R-peak detection (250 Hz)**
- **1D CNN (4-layer: Conv1D + BatchNorm + Pooling + Dense) → classifies beat morphology / arrhythmia**
- **Outputs: HR, rhythm stability, PVC burden → used as covariates in discordance detection**

2. Predicted vs. Observed MAP – Discordance Alert

- **Predicted MAP: Emax-Hill model**
- **Observed MAP: Arterial line (1 Hz)**
- **Discordance: >3.5MAPobs-MAPpred >3.5 → HITL**

3. Possible Causes (Deterministic + ML)

- **DDI (Micromedex severity 3-4)**
- **Disease progression (Isolation Forest on vitals + HR trend from 1D CNN)**
- **Pump failure (infusion log mismatch)**
- **Receptor resistance (PK parameter drift: EC_{50} ↑ >2σ)**

4. Drug Scope – Phase 1

- **Vasopressors: Norepinephrine, Epinephrine, Vasopressin**
- **β-blockers / Antiarrhythmics: Amiodarone, Metoprolol, Esmolol**

5. MPC Controller

- **Horizon: 10 min prediction, 4 min control (30 s sample)**
- **Performance: 40% ↓ MAP variability vs PI (Rinehart 2012)**
- **Fallback: PI if model confidence < 0.80**

Confidence Scoring-

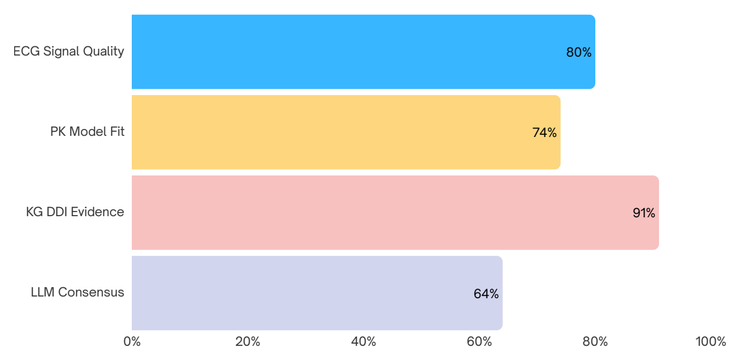

PK/PD Intelligence Engine-

Data Processes:

- **Infusion history + MAP + patient parameters (EMR) → real-time physiological digital twin**

Method:

- **2-compartment PK + Bayesian MAP fitting → Emax-Hill PD → MPC predicts MAP & adjusts infusion every few seconds**

Results output:

- **Personalised drug concentration (Cp), predicted MAP response, closed-loop dose recommendations, deviation alerts**

Project impact:

- **Core automated decision-making for precision vasopressor therapy — reduces manual titration, logs all actions for safety, and scales to other ICU drugs**

Knowledge Graph Layer-

Drug–Drug Interaction (DDI) Edges

- **Source: Micromedex (severity 3–4)**
- **Cardinality: ~200–400 clinically significant pairs**
- **Evidence grading: A (high) / B (moderate) / C (low)**
- **Example: Norepinephrine + Amiodarone → QT prolongation (Level A)**

Mechanism of Action (MOA) Edges

- **Source: DrugBank API (real-time query)**
- **Semantic encoding: Agonism / antagonism at receptor level**
- **Example: α_1-adrenergic agonism (norepinephrine) + β_1-antagonism (metoprolol) → opposing MAP/HR effects**
- **Application: Concordance engine uses MOA for directional response inference**

Contraindication & ADR Edges

- **Sources: SIDER 4.1 + FAERS (auto-populated)**
- **Formulary coverage: 80–120 drugs**
- **Key property: Deterministic rule base – no LLM hallucination**
- **Use case: Safe recommendation filtering, absolute contraindication blocking**

Safety Architecture-

✅ **Dose Validator**
Deterministic range check. Every LLM recommendation is verified against formulary limits before display.

🔒 **Watchdog Process**
Independent OS process monitoring: heartbeat, log timestamps, HTTP health. Triggers audible alarm + SMS on system failure.

🔒 **FMEA Required**
IEC 62304 Class C (patient safety). FMEA for every data pathway. RPN > 100 → mandatory redesign. Compliant with EU MDR 2017/745.

🔒 **Differential Privacy**
Federated learning gradient updates protected by Gaussian mechanism ($\varepsilon = 1.0$, $\delta = 10^{-9}$). Fully GDPR-compliant.

RAG + LLM Recommendation Engine-

🔍 **Hybrid Retrieval — BM25 + Dense (HNSW)**

Reciprocal Rank Fusion combines exact medical term matching (BM25) with semantic similarity (BioBERT embeddings). 5–15% Improvement on clinical NDCG@10 vs dense-only retrieval.

SNOMED-CT synonym expansion handles: norepinephrine ≡ noradrenaline ≡ NorEpi ≡ Levophed.

🧠 **Two-Model LLM Ensemble**

Claude Sonnet 4 (cloud) + BioMedLM (local) run in parallel.
Disagreement on drug name, dose, or conclusion → LOW_CONSENSUS → triggers forced HITL (override even if confidence is high).
Expected outcome: 30% reduction in false-negative validator bypass.

- **Dr. Sohel Kalim Tamboli,** Pharm. D, M.Sc.
Digital Health Lab, UKD / HHU, Düsseldorf
- **Mihamad Abdulkader,** M.Sc.
Digital Health Lab, UKD / HHU, Düsseldorf
- **Dr. Falko Schmid ,** PhD
Digital Health Lab, UKD / HHU, Düsseldorf
- **PD Dr. med. Hug Aubin,** MHBA
Dep. of Cardiovascular Surgery , UKD / HHU, Düsseldorf