SepsisGuard: An Uncertainty-Aware Ensemble Early Warning System for Reliable and Scalable Detection of Sepsis and Acute Clinical Deterioration

Lin, Jophy (School: Manalapan High School)

Sepsis, responsible for 11,000,000+ deaths globally and 270,000+ annually in the U.S., is very time-sensitive. Mortality increases by 7.6% per hour of delayed antibiotic treatment. Bedside tools (qSOFA, SIRS, MEWS) often miss cases, generate many false alarms, or fail to detect sepsis before deterioration, which causes alarm fatigue and delayed intervention. Existing AI systems are trained to provide "best-guesses" without a measure of confidence, making them unsafe for deployment in ICUs. This study introduces SepsisGuard, a scalable GRU–Transformer ensemble (43.9M parameters, 59.5 ms latency, 9.3 GFLOPS) trained on 20,000+ ICU patients (PhysioNet 2019) using domain-adversarial learning for cross-hospital generalization and Monte Carlo dropout for uncertainty estimation. Trained on Hospital A and evaluated on a fully held-out Hospital B cohort, SepsisGuard achieves 82.4% early detection at 0.132 false positives per patient-day (efficiency 6.2), outperforming XGBoost, LightGBM, Random Forest, and bedside tools (AUROC 0.9447, $p < 0.05$) with 2.5–7.8× greater efficiency. It also detects sepsis on median about 28 hours before clinical recognition (7.4× mortality escalation if untreated). With it costing ~$36.5k per case in the U.S., the detection rate corresponds to 824 cases per 1,000 (~$30.1M in associated expenditure). Using uncertainty filtering, detection increases to 96.4% among high-confidence predictions, which enables prioritized alerts while routing uncertain cases to clinicians for further review. With its combination of cross-hospital validation, uncertainty estimation, and early detection, SepsisGuard creates a scalable, physician-centered framework for safer AI deployment in critical conditions, with the framework applicable to broader healthcare settings.

Awards Won: