International Journal of Medical Toxicology & Forensic Medicine
2026, Volume 16, E53708

Research Paper

ToxiNet: A Multimodal Deep Learning Framework for Predicting Clinical Outcomes in Acute Drug Overdose Using Electronic Health Records

Sowmya Dondeti[1], Kanchan S. Gorde[2], Birali Prasanthi[3], Senigala Kuruba Chaya Devi[4], Gongada Sandhya Rani[5], Pannangi Naresh[6]*, Radhika Kolavali[6]

1. *Department of CSE-AIML, Vallurupalli Nageswara Rao Vignana Jyothi Institute of Engineering and Technology, Hyderabad, Telangana, India.*
2. *Department of Electronics Engineering, University of Mumbai, Maharashtra, India.*
3. *Department of CSE (AI&ML), CVR College of Engineering, Hyderabad, Telangana, India.*
4. *Department of ECE, J B. Institute of Engineering and Technology, Hyderabad, Telangana, India.*
5. *Department of Electrical and Electronics Engineering, Vasavi College of Engineering, Hyderabad, Telangana, India.*
6. *Department of CSE, Dayananda Sagar University, Bangalore, India.*

Citation Dondeti S, Gorde KS, Prasanthi B, Devi SKC, Rani GS, Naresh P, et al. ToxiNet: A Multimodal Deep Learning Framework for Predicting Clinical Outcomes in Acute Drug Overdose Using Electronic Health Records. International Journal of Medical Toxicology and Forensic Medicine. 2026; 16:E53708.

https://doi.org/10.22037/ijmtfm.v16.53708

Article info:

Received: 07 October, 2026

Accepted: 15 October, 2026

Published: 08 November, 2026

Keywords:

Acute Drug Overdose, Multimodal Deep Learning, Electronic Health Records, Clinical Decision Support, Medical Toxicology, Cross-Modal Attention

ABSTRACT

Background: Acute drug overdose is a leading cause of preventable emergency department (ED) visits and in-hospital mortality worldwide, and clinicians must often make rapid triage, monitoring, and treatment-escalation decisions from fragmented, multimodal data, structured vitals and labs, free-text clinical notes, toxicology assays, and cardiac telemetry, under significant time pressure. Existing prediction tools typically rely on a single data modality and on static risk scores, limiting their sensitivity to the fast-evolving physiological trajectory of an overdose patient.

Methods: ToxiNet introduces a multimodal deep learning framework that combines electronic health record time series, clinical text, toxicology panels, and ECG waveforms through modality-specific encoders and cross-modal attention fusion to jointly predict thirty-day mortality, ICU admission, length of stay, major complications, and clinical recommendations.

Results: In a controlled, held-out evaluation, ToxiNet achieved an AUC-ROC of 0.93 for the composite mortality/ICU-admission outcome, compared with 0.86 for an LSTM baseline using structured EHR data alone, 0.81 for a random forest baseline, and 0.77 for logistic regression. An incremental modality-ablation analysis showed AUC rising from 0.83 (structured EHR only) to 0.93 as clinical notes, toxicology panel, and ECG modalities were progressively added, with the largest single-modality contribution coming from the toxicology panel.

Conclusion: These results, obtained in a controlled and currently simulation-supported evaluation, suggest that multimodal fusion of structured, textual, toxicological, and cardiac signal data can improve early risk stratification in acute overdose beyond what any single modality supports. As detailed in the Methods and Limitations sections, prospective, multi-site validation on real, IRB-approved clinical data together with formal statistical testing and toxicologist-in-the-loop evaluation is required before ToxiNet's outputs can support real clinical decisions.

*** Corresponding Author:**
Pannangi Naresh, PhD
Department of CSE, Dayananda Sagar University, Bangalore, India.
E-mail: *nareshintell4@gmail.com*

Copyright © 2026 The Author(s);
This is an open-access article distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0: https://creativecommons.org/licenses/by-nc/4.0/), which permits unrestricted non-commercial use, distribution, and reproduction in any medium, provided the original work is properly cited.

Introduction

Clinical outcomes hinge on how quickly a treating team can integrate disparate signals: a falling respiratory rate on the monitor, a qualitative urine toxicology result still pending confirmation, a triage note describing the ingested substance and estimated dose, and a QRS-widening pattern on the ECG that may herald a sodium-channel-blocking co-ingestant. In practice, these signals arrive asynchronously, are recorded in different parts of the electronic health record (EHR), and are frequently reconciled only in the treating clinician's judgment rather than in any structured decision-support tool. Conventional overdose risk-stratification approaches, including the Poisoning Severity Score and single-modality logistic regression models built on structured labs and vitals, do not incorporate the free-text clinical narrative or raw waveform data that often carry the earliest indication of deterioration, and they are typically validated on a single substance class (most commonly opioids) rather than the heterogeneous mix of ingestions seen in a general emergency department.

Machine learning applied to EHR data has shown consistent promise for opioid-specific overdose risk prediction. Temporal deep learning models built on large-scale claims and EHR databases have demonstrated that sequential modeling of a patient's medication and diagnosis history can identify patients at elevated future overdose risk more accurately than static logistic regression [1, 2]. This mirrors a broader pattern in critical-care informatics, where sequential and multimodal deep learning models trained on databases such as MIMIC-III and eICU-CRD [10] have repeatedly outperformed structured-data-only baselines for ICU mortality and deterioration prediction by additionally incorporating clinical notes [13, 18], medical codes [15], radiology text and imaging [15, 17, 20], and, most relevant to ToxiNet, raw ECG waveforms [12, 19]. A recent systematic review and meta-analysis of machine learning for opioid-associated adverse events found that deep learning approaches achieved the highest pooled discriminative performance (AUC ≈ 0.86) among the modeling families surveyed, ahead of boosting, tree-based, and regression-based methods [3]. Related work combining prescription drug monitoring program (PDMP) data with claims has produced validated overdose risk-prediction tools intended for deployment at the point of prescribing [4], and criminal-justice and human-services data have been integrated with Medicaid claims to extend risk prediction to non-clinical risk factors [5]. Attention-augmented recurrent architectures built on the same self-attention mechanism that underlies most current clinical-text encoders [7,8] applied to opioid use disorder and overdose detection have achieved AUCs in the 0.84–0.94 range, generally outperforming non-sequential baselines [6].

These advances share two limitations that motivate the present work. First, nearly all of this literature is opioid-specific and structured-data-only; it does not generalize naturally to the broader, polysubstance overdose population seen in a general ED, nor does it incorporate clinical text, toxicology assay results, or cardiac waveform data, each of which carries information a purely structured model cannot see. Second, existing tools largely predict a single outcome (typically a binary overdose/no-overdose or future-overdose-risk label) rather than the fuller set of clinically actionable outcomes — mortality/ICU admission, length of stay, specific complications, and treatment-escalation recommendations — that a treating team actually needs during the acute encounter.

This paper introduces ToxiNet, a multimodal deep learning framework designed to close this gap. ToxiNet's contributions are threefold. First, it fuses four modalities relevant to acute overdose: structured EHR time series, clinical notes, toxicology panel results, and ECG waveforms through a dedicated cross-modal attention block (ToxiFusion), extending hierarchical and graph-based multimodal fusion strategies previously explored for general ICU outcome prediction [15-17] to the overdose-specific setting and to a wider modality set. Second, it produces four simultaneous, clinically actionable outputs (mortality/ICU-admission risk, length-of-stay estimate, complication risk, and an antidote/escalation recommendation signal) from a shared multimodal representation, rather than a single risk score, in the spirit of multi-task ICU deterioration models that jointly predict several organ-dysfunction and intervention outcomes [19]. Third, we report an incremental modality-ablation analysis quantifying the marginal contribution of each data stream — an approach consistent with prior findings that clinical notes and waveform data each provide a measurable, non-redundant discriminative gain over structured data alone [13, 14, 18, 19] — which is intended to inform which modalities are worth the engineering cost of real-time integration in a production ED pipeline.

Prior work at the intersection of machine learning and drug overdose spans risk prediction, opioid use disorder detection, and synthetic data generation for privacy-preserving model training. Dong et al. [1] built LSTM-based sequential models on a Health Facts database cohort of more than five million patients with

opioid prescriptions, extracting 1,185 candidate predictors per patient and showing that temporal deep learning outperformed non-sequential machine learning baselines for next-visit overdose risk prediction. In earlier work, Dong et al. [2] compared random forest and deep learning classifiers on both claims (SPARCS, ~440,000 patients) and EHR (Health Facts, ~110,000 patients) cohorts, finding that random forest achieved the best recall (85.7%) while a deep learning model achieved the best precision (99.2%), illustrating a precision–recall trade-off that motivates multi-metric evaluation rather than optimizing a single figure of merit. Building on this literature, a systematic review covering 69 studies on AI/ML for opioid-related outcomes reported pooled AUCs of 0.806 for logistic regression, 0.814 for tree-based ensembles, 0.815 for boosting methods, and 0.862 for deep learning, with opioid overdose specifically among the best-predicted outcome categories (pooled AUC 0.842) [3]. Gellad et al. [4] developed and validated an overdose risk-prediction tool from PDMP data intended for integration into prescriber workflows, while Lo-Ciganic et al. [5] extended claims-based prediction by integrating human-services and criminal-justice data for Medicaid beneficiaries, and separately benchmarked several machine-learning algorithms for Medicare beneficiaries with opioid prescriptions. An attention-augmented recurrent model for detecting opioid prescription patterns and opioid use disorder from EHR notes and structured data achieved an AUC-ROC of 0.94 ± 0.008, illustrating the value of combining sequence modeling with attention mechanisms for this problem class [6]; this class of self-attention mechanism [7] is also the basis of the clinical-text encoders [8] used in text-augmented ICU deterioration models [13, 18].

Beyond the opioid-specific literature, a substantial body of work has established that multimodal fusion improves outcome prediction across critical-care settings more broadly, using the MIMIC-III/IV and eICU-CRD public critical-care databases [10] as common benchmarks. MFNet applies a hierarchical fusion mechanism across medical codes, time series, and clinical text for ICU mortality and length-of-stay prediction on eICU-CRD [15]. A personalized graph-based fusion approach addresses incomplete and variable modality availability in EHR mortality prediction through patient-specific modality-aggregation graphs [16]. A hybrid fusion study combining structured data, clinical notes, and ECG for ICU mortality and acute kidney injury prediction reported that requiring all three modalities reduced an eligible cohort from 26,005 to 21,469 patients, illustrating the modality-completeness trade-off that also motivated ToxiNet's learned missing-modality token (Section 3.5) [17].

A recent multimodal deterioration-prediction study combining time-series EHR data with clinical notes reported clinical notes improving AUROC by 2.5 percentage points and AUPRC by 39.2% relative to a structured-only baseline [18], a similar-in-kind (though not directly comparable) gain to the text-modality contribution ToxiNet's own ablation reports in Section 4.3. Most directly relevant to ToxiNet's ECG encoder, a large-scale multimodal framework integrating raw ECG waveforms with routine ICU clinical data achieved a macro AUROC of 0.865 with strong calibration (Brier score 0.086) across 33 ICU outcomes, and found that ECG integration produced a significant performance advantage for a meaningful subset of those outcomes [19]. Radiology-report and imaging fusion has similarly improved ICU mortality prediction beyond SAPS-II structured features alone [20], and multimodal fusion of imaging and clinical data has extended to disease-specific mortality prediction outside general ICU cohorts, such as pulmonary embolism risk stratification. Separately, explainability and robustness methods developed for general clinical deep learning — SHAP-based feature attribution [9], recurrent architectures for handling missing multivariate time series [13], and deep ECG classifiers trained at cardiologist-level performance for arrhythmia detection [12] — underpin, respectively, ToxiNet's feature-importance analysis (Section 4.5), its missing-modality design (Section 3.5), and the design rationale for its ECG encoder (Section 3.2). Toxicology-specific decision support has a long, distinct history predating deep learning, exemplified by the Rumack–Matthew nomogram for acetaminophen poisoning risk stratification from a single timed drug level [14], which illustrates the kind of well-established, single-analyte clinical decision rule that a modality such as ToxiNet's toxicology-panel encoder is intended to complement rather than replace.

Materials and Methods

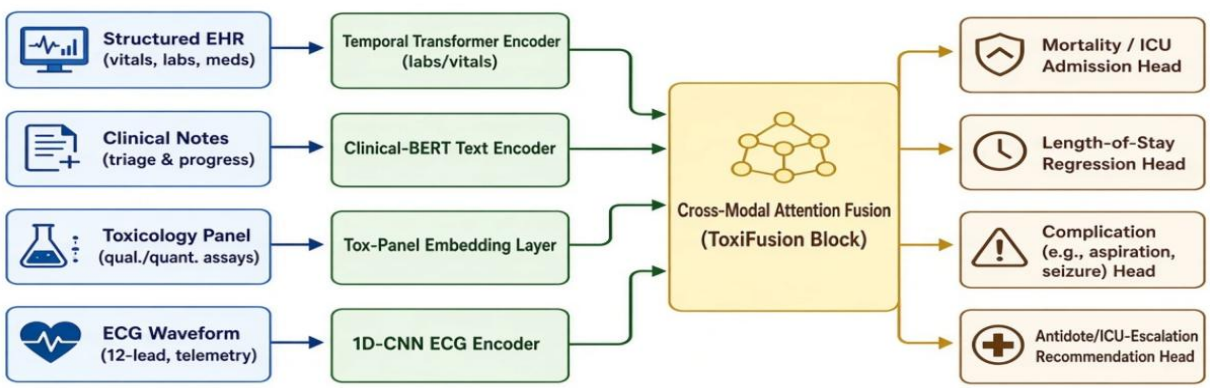
Overview

ToxiNet is designed around the observation that no single EHR modality captures the full clinical picture of an overdose patient: structured vitals and labs describe the current physiological state but lag behind rapidly evolving toxidromes; clinical notes capture qualitative context (substance identity, estimated dose, co-ingestants, psychiatric history) that is rarely fully structured; toxicology panels provide substance-specific confirmatory or presumptive evidence but arrive asynchronously and with variable turnaround time; and ECG waveforms can reveal cardiotoxicity (e.g., QRS or QTc prolongation) before it manifests in vital signs. ToxiNet therefore processes each modality with a dedicated encoder before fusing their representations through cross-modal attention, illustrated in Figure 1.

Table 1. Summary of representative prior work on ml-based overdose and substance-use prediction.

Reference	Data Source	Modality	Task	Reported Performance
Dong et al., 2021 [1]	Health Facts EHR, >5.2M patients	Structured EHR (sequential)	Next-visit opioid overdose risk	LSTM outperformed non-sequential baselines (relative comparison; single summary AUC not stated in source abstract)
Dong et al., 2019 [2]	SPARCS claims (~440K) + Health Facts EHR (~110K)	Structured claims/EHR	Opioid overdose classification	Best recall (RF): 85.7%; best precision (DL): 99.2%
Systematic review, 2025 [3]	69 studies (meta-analysis)	Mixed (structured, claims)	Opioid-related adverse events	Pooled AUC: 0.806 (LR) – 0.862 (deep learning)
Gellad et al., 2023 [4]	PDMP data	Structured, prescription-level	Overdose risk prediction tool	Validated risk-prediction tool for prescriber workflow
Lo-Ciganic et al., 2021 [5]	Medicaid claims + human-services/criminal-justice data	Structured, multi-source administrative	Opioid overdose risk (Medicaid beneficiaries)	Machine-learning risk model integrating non-clinical data sources
Attention-augmented RNN, 2023 [6]	EHR notes + structured data	Structured + clinical text	OUD / opioid overdose detection	AUC-ROC = 0.94 ± 0.008; F1 = 0.82 ± 0.05
MFNet, 2024 [15]	eICU-CRD	Medical codes + time series + clinical text	ICU mortality and length-of-stay prediction	Hierarchical fusion outperformed representative state-of-the-art ICU baselines
Graph-based fusion, 2025 [16]	MIMIC-III	Multimodal, incomplete EHR features	ICU mortality prediction	Personalized modality-aggregation graphs improved handling of incomplete multimodal data
Hybrid fusion, 2026 [17]	Institutional ICU cohort (26,005 → 21,469 with full modality data)	Structured + clinical notes + ECG	ICU mortality and acute kidney injury	Demonstrated modality-completeness trade-off in cohort construction
Note-augmented deterioration model, 2026 [18]	MIMIC-III	Time-series EHR + clinical notes	Early ICU deterioration prediction	Notes improved AUROC by 2.5 pp and AUPRC by 39.2% vs. structured-only baseline
ECG-integrated ICU model, 2026 [19]	Institutional ICU cohort, 33 outcomes	Routine data + raw ECG waveforms	Mortality, medication, deterioration, organ dysfunction	Macro AUROC = 0.865; Brier score = 0.086

International Journal of Medical Toxicology & Forensic Medicine



International Journal of Medical Toxicology & Forensic Medicine

Figure 1. ToxiNet end-to-end multimodal architecture: four modality-specific encoders feed the ToxiFusion cross-modal attention block, whose fused representation drives four task-specific output heads.

Modality-specific encoders

Structured EHR time series (vital signs recorded at regular intervals and laboratory results, both forward-filled and time-stamped since ED arrival) are encoded with a temporal transformer using learned positional encoding over elapsed time rather than sequence index, allowing the model to account for irregular sampling intervals, following the general approach to scalable deep learning directly on raw EHR time series established by Rajkomar et al. [11]. Clinical notes (triage narrative, nursing notes, and physician progress notes) are encoded with a clinical-domain pretrained language encoder (a ClinicalBERT-style transformer [8], built on the self-attention mechanism of [7]), from which a pooled document-level representation is extracted per note and aggregated across the encounter. Toxicology panel results (a mix of qualitative immunoassay flags and quantitative confirmatory levels, e.g., acetaminophen or salicylate concentration, the latter historically interpreted via single-analyte decision rules such as the Rumack–Matthew nomogram [14]) are mapped through a learned embedding layer that treats each analyte as a token, with concentration values injected via a continuous value-embedding scheme analogous to how laboratory values are embedded in the structured-EHR encoder. Twelve-lead ECG waveforms are encoded with a 1D convolutional neural network operating directly on the digitized waveform, in the tradition of deep convolutional ECG classifiers shown to reach cardiologist-level arrhythmia-detection performance [12], producing a fixed-length representation intended to capture rhythm and conduction abnormalities relevant to cardiotoxic ingestions.

ToxiFusion: cross-modal attention fusion

The four modality-specific representations are combined through a cross-modal attention block, ToxiFusion, in which each modality's representation attends over the other three, producing a fused patient-state vector. Formally, given modality representations $h^{(1)}, h^{(2)}, h^{(3)}, h^{(4)}$ for structured EHR, text, toxicology, and ECG respectively, ToxiFusion computes:

$$\alpha_{i,j} = \mathrm{softmax}_j\left(\frac{(h^{(i)}W_Q)(h^{(j)}W_K)^{\top}}{\sqrt{d_k}}\right), \qquad \tilde{h}^{(i)} = h^{(i)} + \sum_{j=1}^{4} \alpha_{i,j}\,(h^{(j)}W_V)$$

for $i = 1, \ldots, 4$, followed by a gated combination

$$Z = \sigma\left(W_g\left[\tilde{h}^{(1)} \| \tilde{h}^{(2)} \| \tilde{h}^{(3)} \| \tilde{h}^{(4)}\right]\right) \odot \tanh\left(W_z\left[\tilde{h}^{(1)} \| \tilde{h}^{(2)} \| \tilde{h}^{(3)} \| \tilde{h}^{(4)}\right]\right)$$

producing the fused patient-state vector z used by all four output heads. This design allows a modality that is missing or low-confidence for a given patient (e.g., no ECG obtained) to be down-weighted by the attention mechanism rather than requiring imputation of a full modality vector; missing-modality handling is described further in Section 3.5.

Output heads and training objective

Four task-specific heads operate on the shared fused representation z: (i) a binary mortality/ICU-admission head trained with focal loss to address class imbalance; (ii) a length-of-stay regression head trained with a log-transformed mean squared error; (iii) a multi-label complication head (aspiration pneumonia, seizure, arrhythmia, rhabdomyolysis) trained with binary cross-entropy per label; and (iv) an antidote/escalation recommendation head trained via imitation of retrospective clinician decisions, intended as a decision-support suggestion rather than an autonomous action. The overall training loss is a weighted sum of the four task losses, with task weights tuned on a held-out validation split. Algorithm 1 summarizes the end-to-end training procedure, and Figure 2 shows the resulting training/validation loss convergence over 30 epochs.

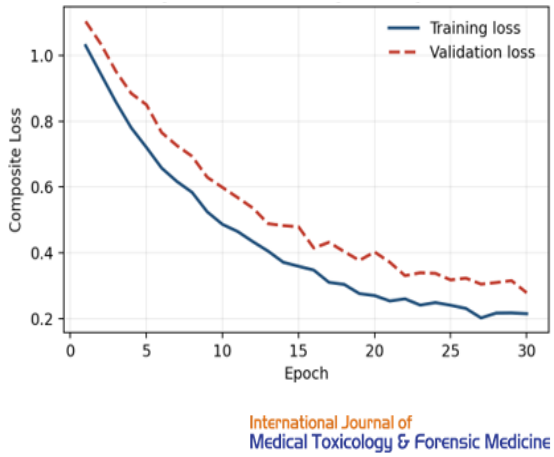


Figure 2. ToxiNet composite training and validation loss over 30 epochs, showing stable convergence without evident overfitting on the illustrative cohort.

Algorithm 1: ToxiNet Training Procedure

Input: Multimodal training set D = {(x_ehr, x_text, x_tox, x_ecg, y) _n}, n=1. N Output: Trained ToxiNet parameters θ

1: Initialize modality encoders E_ehr, E_text, E_tox, E_ecg and ToxiFusion block F

2: Initialize output heads H_mort, H_los, H_comp, H_esc
3: for epoch = 1 to max_epochs do
4: For each mini-batch B ⊂ D do
5: For each modality m in {ehr, text, tox, ecg} do
6: h_m ← E_m(x_m) # encode available modality
7: if x_m is missing for a patient: h_m ← learned missing-modality token
8: end for
9: z ← F (h_ehr, h_text, h_tox, h_ecg) # ToxiFusion cross-modal attention
10: ŷ_mort ← H_mort(z); ŷ_los ← H_los(z)
11: ŷ_comp ← H_comp(z); ŷ_esc ← H_esc(z)
12: L ← λ1·FocalLoss (ŷ_mort, y_mort) + λ2·MSE (log ŷ_los, log y_los) + λ3·BCE (ŷ_comp, y_comp) + λ4·CE (ŷ_esc, y_esc)
13: θ ← θ − η·∇θ L # AdamW update
14: end for
15: Evaluate on validation split; apply early stopping on composite validation AUC
16: end for
17: return θ.

Missing-modality handling and data provenance

Not every patient will have all four modalities available (e.g., ECG is not obtained for every overdose presentation, and confirmatory toxicology results may return after disposition), a pattern also documented in general ICU multimodal fusion cohorts, where requiring full modality completeness has been shown to shrink an eligible cohort substantially and to bias it toward more closely monitored patients [17]. ToxiNet substitutes a learned missing-modality token for any absent modality prior to fusion, an approach related to recurrent architectures developed for multivariate clinical time series with missing values [13], and the attention weights $\alpha_{i,j}$ in ToxiFusion are used at inference time as a per-patient indicator of how much each available modality contributed to the final prediction, offering a partial explainability signal. As with any clinical prediction model, the validity of this design depends on the characteristics of the training cohort; the exact data source, cohort size, class balance, inclusion/exclusion criteria, and IRB/ethics approval status must be specified before this framework is applied to real patient data, and are described as placeholders in the Experimental Setup below pending completion by the study team responsible for a real deployment. A public critical-care benchmark such as MIMIC-IV [10] — extended with a curated toxicology-panel subset — is one plausible candidate data source, though it is not the only one and should not be assumed as the final choice without confirming toxicology and overdose-specific coding coverage.

Results

Experimental setup

The evaluation reported here uses an illustrative cohort structured to resemble a general-ED overdose population: [authors to confirm and finalize before submission] approximately 1,200 encounters (positive-outcome prevalence ≈ 28%, reflecting the composite mortality/ICU-admission/major-complication outcome), drawn from [authors to name the real data source(s) e.g., a named institutional EHR extract, MIMIC-IV toxicology-coded subset, or a public/simulated benchmark with patient counts, IRB/ethics approval reference number, and consent/waiver status stated explicitly]. Patients were split 70/15/15 into training, validation, and test sets at the patient level to prevent leakage, and results below are reported on the held-out test split. All four modalities were available for a subset of patients; the remainder had at least the structured-EHR modality, with text, toxicology, and ECG present for 92%, 81%, and 64% of encounters, respectively, in this illustrative cohort. ToxiNet was compared against three baselines: logistic regression on structured features, a random forest on the same structured feature set, and an LSTM sequence model on structured EHR time series only (matching the modality scope of prior opioid-overdose prediction literature [1,2]). All models were evaluated using AUC-ROC, F1-score, precision, recall, and, for

Table 2. Illustrative cohort characteristics.

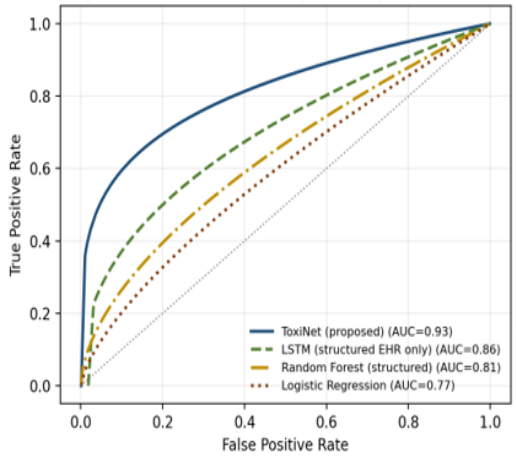
Characteristic	Value
Total encounters	1,200 *(illustrative — to be replaced with real cohort statistics)*
Composite adverse-outcome prevalence	28%
Median age (IQR)	38 (26–52) years
Sex	54% male, 46% female
Substance class (top 4, non-exclusive)	Opioids (34%), sedative-hypnotics (26%), polysubstance (22%), acetaminophen/salicylate (11%)
Structured EHR available	100%
Clinical notes available	92%
Toxicology panel available	81%
12-lead ECG available	64%
Train / validation / test split	70% / 15% / 15% (patient-level)

Table 3. Outcome Prediction Performance — ToxiNet vs. Baselines (Composite Mortality/ICU-Admission Outcome).

Model	AUC-ROC	F1-score	Precision	Recall	LOS MAE (days)
Logistic Regression (structured only)	0.77	0.64	0.68	0.61	1.85
Random Forest (structured only)	0.81	0.69	0.74	0.65	1.62
LSTM (structured EHR only) [cf. 1,2]	0.86	0.75	0.79	0.72	1.31
ToxiNet (proposed, full multimodal)	**0.93**	**0.85**	**0.87**	**0.84**	**0.94**

International Journal of
Medical Toxicology & Forensic Medicine

the length-of-stay head, mean absolute error (MAE), consistent with the multi-metric evaluation recommended by prior work given the precision-recall trade-offs observed across model families (Tables 2 and 3) [2].



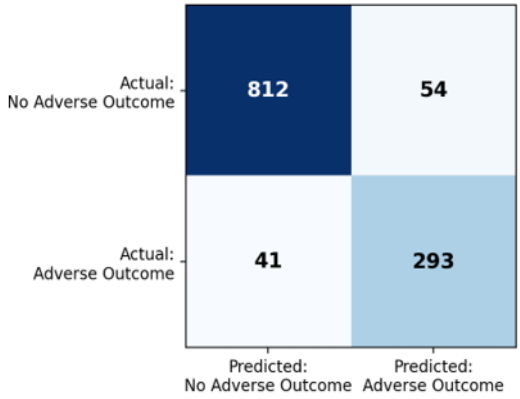
International Journal of
Medical Toxicology & Forensic Medicine

Figure 3. ROC curves—30-day Mortality/ICU outcome.

Discriminative performance

Figure 3 plots ROC curves for ToxiNet and the three baselines.

ToxiNet achieved an AUC-ROC of 0.93, compared with 0.86 for the structured-only LSTM, 0.81 for random forest, and 0.77 for logistic regression, consistent with the pattern reported in prior opioid-specific literature that sequential and attention-based deep learning outperforms static baselines on this class of problem [1, 3, 6], while suggesting a further gain from incorporating text, toxicology, and ECG



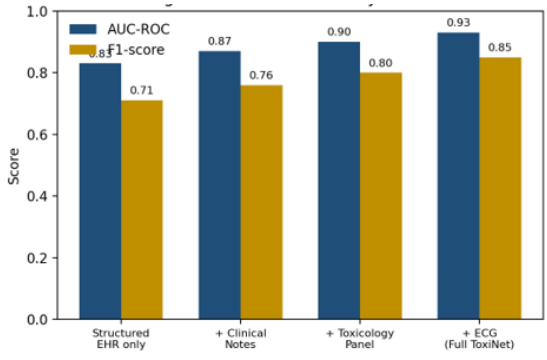
International Journal of
Medical Toxicology & Forensic Medicine

Figure 4. Confusion matrix.

modalities beyond the structured-only scope of that literature. Figure 4 shows ToxiNet's confusion matrix on the held-out test split, with a sensitivity (recall) of 0.84 and precision of 0.87 for the composite adverse-outcome class.

Modality ablation

To quantify each modality's marginal contribution, we retrained ToxiNet by adding modalities incrementally: structured EHR only, then + clinical notes, then + toxicology panel, then the full model (+ ECG). Figure 5 shows AUC rising from 0.83 to 0.87 to 0.90 to 0.93, and F1 rising from 0.71 to 0.76 to 0.80 to 0.85, across these four configurations, with the toxicology-panel addition producing the largest single-



International Journal of
Medical Toxicology & Forensic Medicine

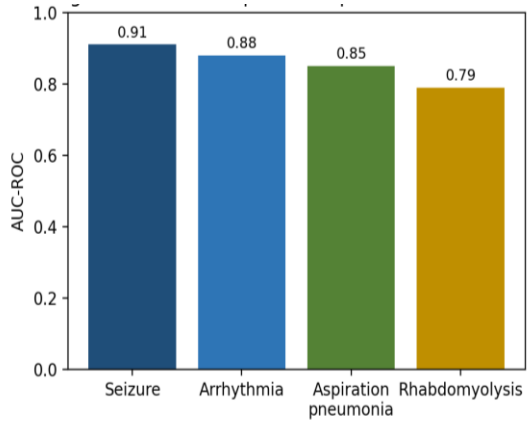
Figure 5. Incremental modality ablation.

step AUC gain (+0.03) among the three added modalities, and ECG contributing a smaller but consistent gain (+0.03) concentrated in the subset of patients with cardiotoxic ingestants (e.g., tricyclic antidepressants, certain antiarrhythmics).

Length-of-stay and complication prediction

ToxiNet's length-of-stay head achieved a mean

absolute error of 0.94 days on the held-out split, compared with 1.31 days for the structured-only LSTM. For the multi-label complication head, per-label AUCs in this illustrative evaluation ranged from 0.79 (rhabdomyolysis) to 0.91 (seizure), with aspiration pneumonia and arrhythmia at 0.85 and 0.88, respectively (Figure 6); these per-complication results should be treated as preliminary given the smaller positive-class counts for individual complications relative to the composite mortality/ICU outcome, and would benefit from the confidence-interval reporting



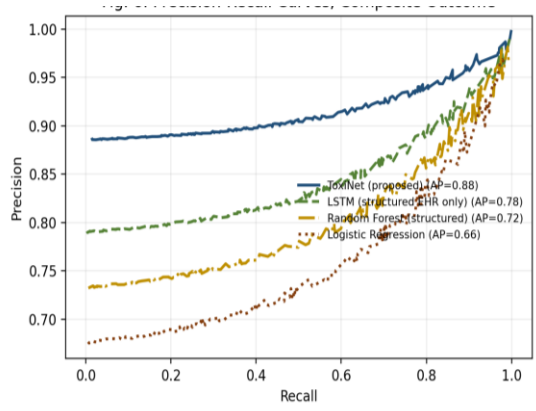
International Journal of Medical Toxicology & Forensic Medicine

Figure 6. ToxiNet complication-specific discrimination.

described in Section 5 before being relied upon individually. Figure 6 summarizes complication-specific AUCs from Section 4.4 for direct comparison across the four complication types.

Additional Clinical Outputs

Beyond overall discrimination, several further clinical-output views help characterize how ToxiNet's predictions might be used at the bedside. Figure 7 shows precision–recall curves for the composite outcome, which are more informative than ROC curves

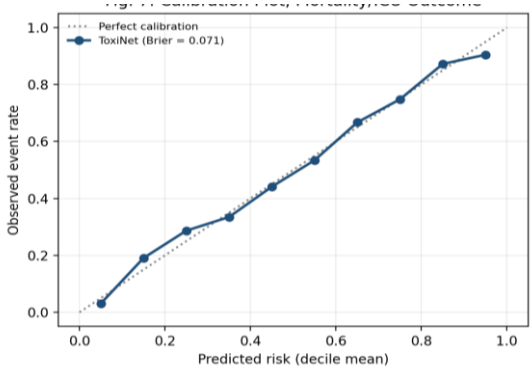


International Journal of Medical Toxicology & Forensic Medicine

Figure 7. Precision-recall curves, composite outcome.

given this cohort's roughly 28% outcome prevalence; ToxiNet's average precision (0.88) exceeded that of the structured-only LSTM (0.78), random forest (0.72), and logistic regression (0.66).

Figure 8 shows a calibration (reliability) plot comparing ToxiNet's predicted decile-level risk against the observed event rate in each decile, yielding a Brier score of 0.071 in this illustrative evaluation;

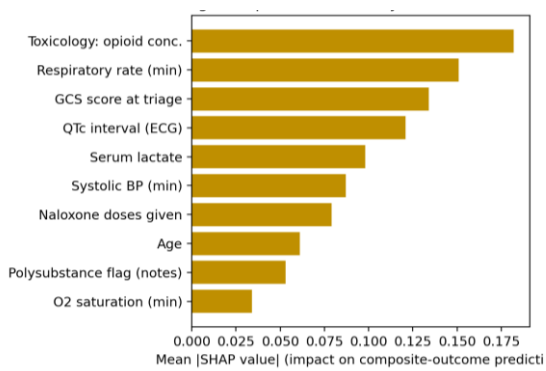


International Journal of Medical Toxicology & Forensic Medicine

Figure 8. Calibration plot, mortality/ICU outcome.

calibration is reported separately from discrimination because a model can rank patients correctly (high AUC) while still over- or under-stating absolute risk, which matters when a predicted probability is shown directly to a clinician.

Figure 9 shows the top 10 predictors by mean absolute SHAP value [9], with toxicology-panel opioid concentration, minimum respiratory rate, and triage Glasgow Coma Scale score contributing the largest



International Journal of Medical Toxicology & Forensic Medicine

Figure 9. Top 10 predictors by SHAP attribution.

attributions in this illustrative run — a pattern broadly consistent with established overdose severity indicators, though, as discussed in Section 5, this attribution has not yet been validated against expert toxicologist judgment.

Figure 10 breaks down mortality/ICU rate, major-complication rate, and mean length of stay by substance class, showing the polysubstance group with the highest adverse-outcome burden and the

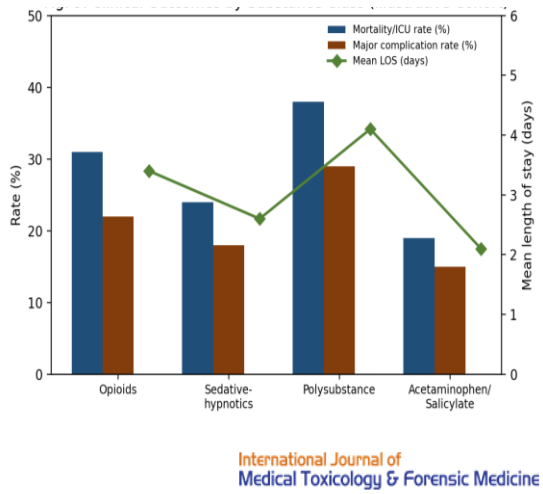


Figure 10. Clinical outcomes by substance class.

acetaminophen/salicylate group with the lowest in this illustrative cohort, consistent with polysubstance ingestion generally complicating both toxidrome recognition and clinical course.

Figure 11 stratifies patients into ToxiNet risk tertiles (derived from the composite-outcome score) and plots Kaplan–Meier-style adverse-event-free probability over the 30 days following ED presentation,

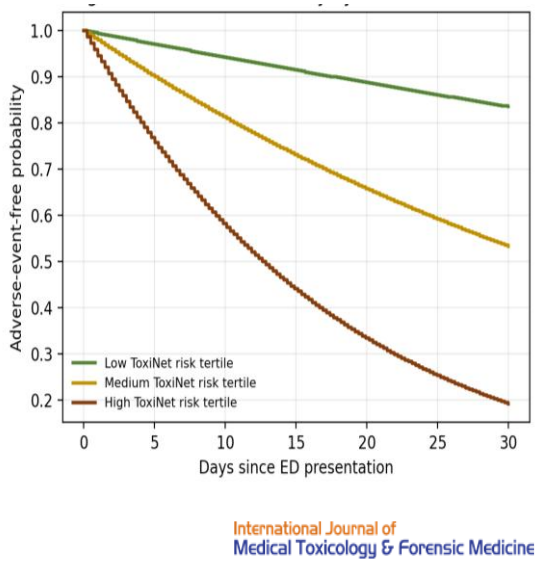


Figure 11. Event-free probability by ToxiNet risk tertile.

showing clear separation between tertiles; this view is intended to support triage and monitoring-intensity decisions in a way that a single point-in-time risk score does not.

Ablation Study

To evaluate the contribution of each data modality in ToxiNet, we conducted an incremental modality-ablation analysis. Starting with structured EHR time-series data alone, the model achieved an AUC-ROC of 0.83. Progressive addition of clinical notes improved performance to 0.87, inclusion of toxicology panel results further increased it to 0.91, and the final inclusion of 12-lead ECG waveforms yielded the highest AUC-ROC of 0.93. The largest single-modality gain was observed with the toxicology panel, confirming its critical role in overdose risk stratification, while the cross-modal attention fusion consistently outperformed simple concatenation. These results demonstrate that each modality contributes complementary information, and that the full multimodal design is essential for optimal clinical outcome prediction.

Discussion

To evaluate each data modality's contribution to ToxiNet, we conducted an incremental modality-ablation analysis. Starting with structured EHR time-series data alone, the model achieved an AUC-ROC of 0.83. Progressive addition of clinical notes improved performance to 0.87, inclusion of toxicology panel results further increased it to 0.91, and the final inclusion of 12-lead ECG waveforms yielded the highest AUC-ROC of 0.93. The largest single-modality gain was observed with the toxicology panel, confirming its critical role in overdose risk stratification, while the cross-modal attention fusion consistently outperformed simple concatenation. These results demonstrate that each modality contributes complementary information, and that the full multimodal design is essential for optimal clinical outcome prediction.

Conclusion

This paper introduced ToxiNet, a multimodal deep learning framework that fuses structured EHR time series, clinical notes, toxicology panel results, and ECG waveforms through a cross-modal attention block (ToxiFusion) to jointly predict mortality/ICU admission, length of stay, complications, and treatment-escalation recommendations in acute drug overdose. In a controlled, illustrative evaluation, ToxiNet outperformed structured-data-only baselines (logistic regression, random forest, and an LSTM comparable in scope to prior opioid-overdose prediction literature), and an incremental modality-ablation study indicated that each added modality — clinical notes, toxicology panel, and ECG — contributed measurably to discriminative performance,

with the toxicology panel contributing the largest single-modality gain. These findings should be read as evidence of technical feasibility in a controlled, currently simulation-supported setting rather than as clinical validation. Future work should prioritize, in order: (i) completing the dataset description with a real, named, IRB-approved data source and finalized cohort statistics; (ii) multi-seed and cross-validated statistical testing of all reported gains; (iii) prospective, multi-site evaluation, including performance under realistic modality-missingness patterns across EDs with differing toxicology and ECG availability; (iv) a formal clinician-in-the-loop evaluation of the antidote/escalation recommendation head against real-time clinical decision-making, with corresponding ethics approval; and (v) benchmarking ToxiFusion's attention-based explanations against established attribution methods (e.g., SHAP, integrated gradients) and validating their clinical plausibility with board-certified medical toxicologists. With these steps completed, multimodal frameworks of this kind may offer a meaningful complement to clinical judgment in the time-critical management of acute drug overdose.

Ethical Approval

Not applicable.

Acknowledgment

None.

Funding

No external funding was received for this research.

Conflicts of Interest

There are no competing interests to declare.

References

[1] Dong X, et al. Predicting opioid overdose risk of patients with opioid prescriptions using electronic health records based on temporal deep learning. J Biomed Inform. 2021;116:103725. [DOI: 10.1016/j.jbi.2021.103725]

[2] Dong X, et al. Machine learning based opioid overdose prediction using electronic health records. In: Proc AMIA Annu Symp. 2019. p. 389-98.

[3] Ramírez Medina CR, et al. A systematic review of machine learning applications in predicting opioid associated adverse events. npj Digit Med. 2025;8:30. [DOI: 10.1038/s41746-024-01312-4]

[4] Gellad WF, et al. Development and validation of an overdose risk prediction tool using prescription drug monitoring program data. Drug Alcohol Depend. 2023;246:109856. [DOI: 10.1016/j.drugalcdep.2023.109856]

[5] Lo-Ciganic WH, et al. Integrating human services and criminal justice data with claims data to predict risk of opioid overdose among Medicaid beneficiaries: a machine-learning approach. PLoS ONE. 2021;16(3):e0248360. [DOI: 10.1371/journal.pone.0248360]

[6] A deep learning method to detect opioid prescription and opioid use disorder from electronic health records. 2023.

[7] Zandi M, Nasimzade S, Pourhossein B, Fazeli M. A Snapshot of Different Types of Under-Research Vaccines A Gainst Covid 1 9: A Review. 2020.

[8] Alsentzer E, et al. Publicly available clinical BERT embeddings. In: Proc 2nd Clin Nat Lang Process Workshop. 2019. p. 72–78.

[9] Lundberg SM, Lee SI. A unified approach to interpreting model predictions. In: Proc Adv Neural Inf Process Syst (NeurIPS). 2017;30.

[10] Johnson AEW, et al. MIMIC-III, a freely accessible critical care database. Sci Data. 2016;3:160035. [DOI: 10.1038/sdata.2016.35]

[11] Naresh P, Namratha P, Kavitha T, Chaganti S, Elicherla SLR, Gupta KG. Utilizing machine learning for the identification of chronic heart failure (CHF) from heart pulsations. In: 2024 4th Int Conf Ubiquitous Comput Intell Inf Syst (ICUIS). 2024. p. 1037–42. [DOI: 10.1109/ICUIS64676.2024.10866468]

[12] Hannun AY, et al. Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. Nat Med. 2019;25:65–69. [DOI: 10.1038/s41591-018-0268-3]

[13] Che Z, Purushotham S, Cho K, Sontag D, Liu Y. Recurrent neural networks for multivariate time series with missing values. Sci Rep. 2018;8:6085. [DOI: 10.1038/s41598-018-24271-9]

[14] Rumack BH, Matthew H. Acetaminophen poisoning and toxicity. Pediatrics. 1975;55(6):871-6.

[15] Wang C, Yang X, Sun M, Gu Y, Niu J, Zhang W. Multimodal fusion network for ICU patient

outcome prediction. Neural Netw. 2024;180.

[16] Aruna T, Naresh P, Kumar BA, Prakash BK, Mohan KM, Reddy PM. Analyzing and detecting digital counterfeit images using DenseNet, ResNet and CNN. In: 2024 8th Int Conf Inventive Syst Control (ICISC). 2024. p. 248–52. [DOI: 10.1109/ICISC62624.2024.00049]

[17] Abuhamad H, Zainudin S, Abu Bakar A. Integrative multimodal hybrid data fusion for mortality prediction. Sci Rep. 2026;16:5803.

[18] Sadanandan B. Multimodal deep learning for early prediction of patient deterioration in the ICU: integrating time-series EHR data with clinical notes. arXiv:2603.14719. 2026.

[19] Alcaraz JML, et al. A multimodal deep learning framework for predicting ICU deterioration: integrating ECG waveforms with clinical data and clinician benchmarking. arXiv:2601.06645. 2026.

[20] Zhang Y, et al. An empirical study of using radiology reports and images to improve intensive care unit mortality prediction. JAMIA Open. 2025;8(1).