

RESEARCH

Open Access



Explainable AI for infection prevention and control: modeling CPE acquisition and patient outcomes in an Irish hospital with transformers

Minh-Khoi Pham^{1,2*}, Tai Tan Mai^{1,2}, Martin Crane^{1,2}, Rob Brennan^{2,3}, Marie E. Ward⁴, Una Geary⁴, Declan Byrne⁴, Brian O'Connell⁴, Colm Bergin^{4,6}, Donncha Creagh⁵, Nick McDonald⁶ and Marija Bezbradica^{1,2}

Abstract

Background Carbapenemase-Producing Enterobacteriace (CPE) poses a critical concern for infection prevention and control in hospitals. However, predictive modeling of previously highlighted CPE-associated risks such as readmission, mortality, and extended length of stay (LOS) remains underexplored, particularly with modern deep learning approaches. This study introduces an eXplainable AI (XAI) modeling framework to investigate CPE impact on patient outcomes from Electronic Medical Records (EMR) data of an Irish hospital.

Methods We analyzed an inpatient dataset from an Irish acute hospital (2018–2022), incorporating diagnostic codes, ward transitions, patient demographics, infection-related variables and contact network features. Several Transformer-based architectures (e.g., TabTransformer, TabNet) were benchmarked alongside traditional machine learning models. Clinical outcomes were predicted, and XAI techniques were applied to interpret model decisions.

Results Our framework successfully demonstrated the utility of Transformer-based models, with TabTransformer consistently outperforming baselines across multiple clinical prediction tasks, especially for CPE acquisition (Area Under Receiver Operating Characteristic and sensitivity). We found infection-related features, including historical hospital exposure, admission context, and network centrality measures, to be highly influential in predicting patient outcomes and CPE acquisition risk. Explainability analyses revealed that features like "Area of Residence", "Admission Ward" and prior admissions are key risk factors. Network variables like "Ward PageRank" also ranked highly, reflecting the potential value of structural exposure information.

Conclusion This study presents a robust and explainable AI framework for analyzing complex EMR data to identify key risk factors and predict CPE-related outcomes. Our findings underscore the superior performance of the Transformer models and highlight the importance of diverse clinical and network features. The transparent interpretability offered by our XAI approach provides actionable insights for infection prevention and control, paving the way for more targeted interventions and ultimately enhancing patient safety within acute healthcare settings.

*Correspondence:
Minh-Khoi Pham
minhkhoi.pham@adapptcentre.ie

Full list of author information is available at the end of the article



© The Author(s) 2025. **Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

Keywords Electronic medical records, Deep learning, Explainable AI, Transformers

Background

Infection Prevention and Control (IPC) is a cornerstone of modern healthcare, aiming to reduce the spread of healthcare-associated infections (HCAIs) and mitigate their significant impact on patients, healthcare professionals, and hospital systems. Among the most urgent threats in this domain are Carbapenemase-Producing Enterobacterales (CPE)—a class of multidrug-resistant organisms that diminish treatment efficacy in clinical environments [1, 2]. The World Health Organization formally recognizes HCAIs as a critical patient safety concern through its International Classification for Patient Safety (ICPS) [3]. In Ireland, CPE has emerged as a national health priority, with case numbers rising steadily over the past decade [4–7]. This growing burden underscores the need for smarter surveillance, screening strategies, and data-driven models to support timely IPC interventions [7, 8].

This work is part of the ARK-Virus project, an interdisciplinary initiative focused on advancing IPC practices across Irish acute hospitals [9]. By combining expertise in microbiology, human ergonomics, informatics, and data science, the project seeks to develop risk-aware, evidence-based frameworks to support healthcare decision-making [10, 11]. Within this framework, our work focuses on using AI technologies to strengthen IPC for CPE infections, leveraging longitudinal data from Electronic Medical Records (EMRs) and applying machine learning (ML) to investigate how transmission risk, screening strategies, and patient characteristics intersect with clinical outcomes [12].

CPE infections are known to associate with poor patient outcomes, including prolonged hospital stays, higher in-hospital mortality, and increased risk of readmission [13, 14]. In this study, we aim to quantify these impacts more systematically by modeling three key clinical endpoints: 30-day readmission, in-hospital mortality, and length of stay (LOS). By leveraging AI, we aim to gain actionable insights into the correlation and causation between CPE infection and the patient outcomes, informing both clinical risk stratification and hospital resource planning.

Recent advances in AI have significantly contributed to IPC, especially in the surveillance and prediction of HCAIs, antimicrobial resistance (AMR), and timely intervention planning [15–17]. Numerous models have been developed to forecast colonization risks of multidrug-resistant organisms at hospital admission by leveraging features such as prior antibiotic use, comorbidities, invasive procedures, and long-term care exposure [18, 19]. Contact network-based features—derived from

ward co-location and patient trajectories have also been explored to model in-hospital transmission risk [20, 21], while large-scale EMR-based models like ARGai 1.0 have demonstrated the utility of Transformer-based architectures in tasks ranging from gene expression classification to pathogen detection [22]. However, specific attention to CPE remains limited, especially within the Irish healthcare context. Prior studies on CPE typically focus on individual risk factors [23–27], but rarely integrate contact data or interpretability frameworks. We summarize the known CPE risk factors reported in the literature in Table 1.

Transformer-based models such as Med-BERT [29], BEHRT [30], Hierarchical BEHRT [31], CORE-BEHT [32], and EHR-Mamba [33] have shown strong performance for EMR-based patient outcome prediction. These models leverage specialized embeddings for age, visit gaps, and diagnoses, often pretrained on large-scale EMR corpora and fine-tuned for downstream clinical tasks. Tabular deep models like Convolutional Networks [34], TabNet [35], TabTransformer [36], ResNet [37], and TabPFN [38] have similarly demonstrated effectiveness across general tabular datasets, but not quite well studied in healthcare domain. In parallel, explainability has become a critical component in model adoption. While tools like SHAP and LIME are commonly used for global interpretation [39–41], techniques like DeepXplainer [42] and recommendations by Wani et al. [43] highlight the need for interpretable models in clinical environments. Recently, gradient-based approaches such as Integrated Gradients have been used to gather finer insight into deep model behavior [44], even for medical models [33]. Our study builds upon this work by benchmarking recent tabular and EMR-specific models for CPE outcome prediction, embedding contact network features, and applying integrated explainable AI (XAI) tools to support actionable insights.

To model these phenomena, we use a real-world anonymized EMR dataset from an Irish acute hospital spanning four years of inpatient admissions. We employ a mix of machine learning models—ranging from traditional gradient-boosted trees to explainable deep learning approaches such as ResNet, TabNet and Transformer-based architectures—to capture complex temporal and contextual patterns that influence CPE risk and outcomes. More importantly, we introduce **contact network-derived features** based on patient room overlaps and exposure histories to model potential transmission pathways within the hospital. Given the critical nature of care in such a setting and known challenges of deep model transparency, we emphasize explainability

Table 1 Common risk factors for CPE acquisition identified in literature

Risk Factor	References
Prior antimicrobial use (carbapenems, cephalosporins, fluoroquinolones)	[23, 28]
Recent hospitalization (within past year)	[23, 24]
Prolonged hospital stay (> 20 days)	[26, 28]
Stay in high-risk wards (e.g., ICU, transplant units)	[18, 23]
Use of invasive devices (e.g., catheters, ventilation)	[19, 26]
Surgical interventions	[23]
Comorbidities (e.g., diabetes, CKD, malignancy)	[19, 25]
Cross-border healthcare exposure or patient transfers	[24, 28]
Residence in long-term care facilities	[18, 21]
Severity of illness or immunosuppression	[19, 25]

throughout, applying multiple XAI techniques to understand feature contributions and to build trust in model-driven IPC strategies.

Objective. This work presents a domain-centered, explainable AI modeling framework for investigating CPE impact on clinical outcomes in an Irish hospital. Rather than proposing novel architectures, we aim to systematically apply and evaluate existing state-of-the-art Transformer-based models in a real-world healthcare setting, with a focus on IPC-relevant insights.

Research Questions. This study is guided by three central research questions:

- **(RQ1)** How do infection-related features impact on future patient outcomes such as 30-day readmission, in-hospital mortality, and LOS?
- **(RQ2)** Which patient features are the most predictive of CPE infection risk?
- **(RQ3)** Can deep learning and XAI methods identify distinctive patterns among CPE-positive patients at the cohort level, and how can these tools support personalized analysis at the individual level?

Contributions. The main contributions of this study are:

- We benchmark a suite of traditional, deep, and Transformer-based models on a rare and newly gathered CPE dataset, offering comparative performance insights.
- We incorporate contact network-derived features to model in-hospital exposure, supporting nuanced analysis of CPE acquisition dynamics.
- We utilize Integrated Gradients to study both global and individual-level patterns of infection and risk. With this, we generate clinically actionable findings on the potential operational effectiveness of screening and isolation strategies, directly supporting IPC optimization.

Dataset

Dataset description

We utilize anonymized inpatient records from a national acute hospital in Ireland, collected between January 2018 and February 2022 as part of the ARK-Virus project [9]. The dataset is derived from multiple sources including the national Hospital Inpatient Enquiry (HIPE) system and Patient Admission System from the hospital discharge database of an Irish acute hospital [45, 46], and includes nearly one million bed-day records.

Each entry contains structured metadata on patient demographics, admission and discharge details, ward transfers, and International Classification of Diseases (ICD-10) codes for diagnoses and procedures [47]. Additionally, the dataset contains results from CPE screening tests conducted during hospitalization. Most CPE-positive cases in our dataset are identified through these routine screenings, which primarily detect colonization (organism presence without symptoms) rather than active infection. As shown in Fig. 1, CPE-positive patient cohort seems to exhibit slightly higher numbers of readmissions, ward transitions, diagnoses, and longer lengths of stay compared to CPE-negative patients. Full summary statistics of the data and a comprehensive overview of the feature types included is presented in Appendix.

Cohort construction and preprocessing

We constructed our study cohort by applying a multi-stage filtering and preprocessing pipeline:

1. Anonymized data was shared on all acute inpatient admissions between January 1, 2018 and February 28, 2022.
2. Data on adult patients (age ≥ 18) with hospital stays of at least 48 hours and at least one prior admission.
3. CPE status was assigned based on the presence of a positive screening test during the admission. In-hospital mortality was identified using the discharge status field.
4. To simulate prospective evaluation, we performed a chronological data split, allocating 90% of episodes to the training set and 10% to the test set.

To reduce diagnostic code sparsity while retaining clinical meaning, we applied code generalization: diagnosis codes were truncated to the first three digits of the ICD-10 hierarchy, and procedure codes to the first five digits. This approach aligns with prior research on medical code abstraction [48–50]. Missing entries were encoded as “None”, and duplicate rows or empty columns were removed to ensure data integrity.

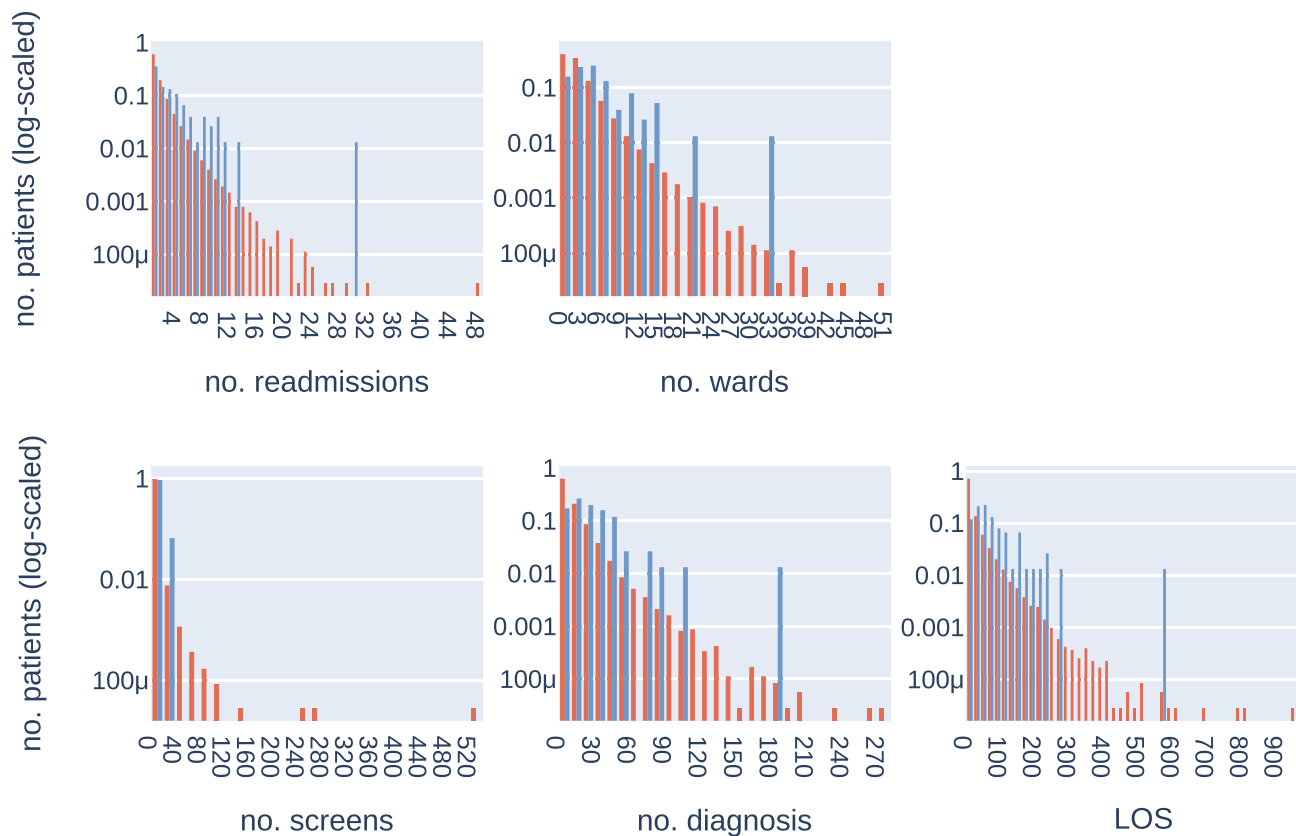


Fig. 1 Distribution of the two studied patient cohorts based on CPE screening status, whether it is positive (CPE+) or negative (CPE-), ■ CPE-, ■ CPE+

Feature engineering

Past episodes features

Patient history plays a critical role in personalized risk modeling. Studies show that incorporating visit-level sequences of clinical events enhances predictive performance [51–53]. While our dataset lacks fine-grained timestamps, we extract cumulative historical features (e.g., such as the total number of diagnoses, comorbidities, ward transfers, and prior admissions) to summarize past medical activities. This approach is expected to maintain the predictive utility of longitudinal data while being compatible with both static tabular models and Transformer-based architectures.

Ward transitions features

During a patient’s hospital stay, following diagnosis, they may be transferred between different medical wards to receive specialist care or for bed management reasons. According to the research we gathered in Table 1, ward information can contain useful information that are indicative of CPE infection. To capture this, we compute the frequency of admissions to both admission and discharge wards, which the model encodes as numerical and categorical features, respectively. These features are presented in the Appendix as “No. Acute Department

visits”, “No. Emergency Department visits”, “Admission Ward”, and “Discharge Ward”

CPE-related features

CPE screening data plays a central role in our analysis. We extracted variables indicating whether a patient underwent CPE testing during an admission, as well as the corresponding results (positive/negative). In some cases, diagnosis codes also contained CPE-related entries. Screening is conducted in accordance with Irish Health Service Executive [54], typically for surveillance or diagnostic purposes. These screening variables serve as outcome labels in our CPE acquisition modeling (in Sect. “CPE infection forecast”), and as input features in our patient outcome prediction tasks (in Sect. “Patient Outcome forecast”). Their presence enables us to evaluate their influence on model predicting the clinical outcomes.

Through out the paper, for the ease of writing, we use the term “CPE-positive” as a patient who has had at least one positive result on CPE screening, regardless of whether this reflects carriage or active infection. We note that infection status can change over time (e.g., clearance of infection or subsequent acquisition), while screening captures both colonization and infection events. This

definition is adopted for consistency and interpretability within our modeling framework

Contact network features

To capture potential transmission dynamics within the hospital, we constructed daily patient contact networks at the ward level. Two patients were considered *in contact* if they occupied the same ward on the same day. Any patient who shared a ward with a CPE-positive case was marked as *exposed*.

From these networks, we computed a set of graph-based features for each patient admission, including node degree, closeness centrality, and daily contact counts (formal definitions of these terms can be found in Appendix). These features were aggregated over the hospital stay. While prior studies have used finer spatial granularity (e.g., room- or building-level) [20], we limited our analysis to ward-level data due to availability constraints. We report the statistics of these features in the Appendix.

Methodology

Overview of the framework

Our approach integrates structured EMR data with contact network-derived features to forecast both clinical outcomes and infection risks. As shown in Fig. 2, we combine information from multiple sources including demographics, past episode summaries, ICD-10 codes (and their textual descriptions), and in-hospital patient

contact networks. Each feature type is encoded separately according to its data modality.

For numerical features, we use multi-layer perceptron (MLP) layers to project them into the shared embedding space. Categorical variables are encoded via lookup tables to obtain dense vector embeddings. All projected representations are aligned into a common hidden dimension and concatenated before model fusion.

To preserve diagnostic history, we incorporate all diagnosis and procedure codes from both the index admission and prior hospitalizations. Prior work has demonstrated that leveraging longitudinal data in this way significantly improves predictive accuracy [12, 23, 24].

Clinical codes are among the most informative inputs for modeling patient trajectories and predicting outcomes [55]. However, their high dimensionality and sparsity introduce significant challenges. To mitigate this, we perform code generalization (as described in Sect. “XAI techniques” and utilize a publicly available pretrained embeddings [56]. By doing this, the list of billing codes patients received during the admission are transformed into a list of feature vectors. These vectors are then, similarly, projected into the shared embedding space by another MLP layer. For the procedure codes, which lack pretrained resources, we train embedding layers from random initialization.

To further address sparsity and enhance generalization, we enrich the clinical code representations with semantic

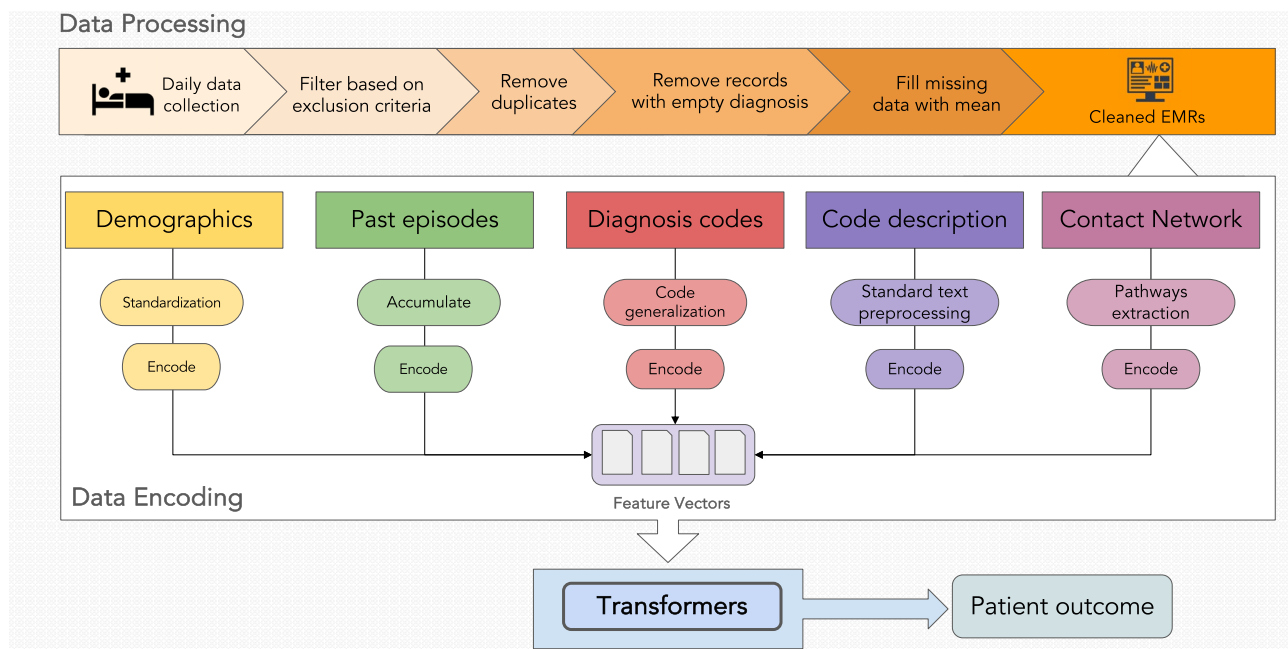


Fig. 2 Overview of the proposed framework for processing and modeling structured EMR data. The pipeline integrates demographic data, ward history, diagnosis codes, and contact network features into transformer-based predictive models

embeddings derived from their textual descriptions. Specifically, we use a pretrained BioBERT model¹ to generate 768-dimensional semantic vectors for each ICD code description. These vectors are then linearly projected to match the model's shared input dimension, allowing for integration with other feature types.

The encoded features are processed through a model-specific backbone—such as ResNet, attention-based models, or Transformer variants—and subsequently projected into a unified embedding space. These combined representations are then passed through a MLP head for final prediction. We benchmark several recent deep learning models tailored for structured or EMR data, each utilizing its own encoding and feature fusion strategy, as detailed in Sect. “[Explainable AI](#)”.

XAI techniques

To deepen our understanding of model behavior and uncover patterns in the data, we incorporate a range of XAI techniques, selected to be compatible with the deep learning architectures used in this study. In particular, we focus on **Integrated Gradients (IG)** [44], a theoretically grounded method for attributing prediction outcomes to input features. Despite its strong theoretical foundations, IG remains underutilized in clinical tabular modeling, as noted by Allgaier et al. [57].

Unlike SHAP, which relies on sampling an exponentially large space of possible feature combinations to approximate feature contributions [58], IG computes gradients along a path from a baseline input to the actual input. This makes it both computationally efficient and well-suited for interpreting predictions from differentiable deep models. IG also enables flexible use cases, from estimating global feature importance across populations to individualized explanations at the patient level.

To address our research questions—**identifying key predictors of patient outcomes**—we calculate feature attributions using IG across multiple benchmarked models (TabNet, TabTransformer and ResNet) and cross-validation folds. For each model, we compute the IG values relative to a baseline input, and then average the resulting scores across all samples. Feature importance rankings are derived by aggregating these scores across the results and ranked. We report and visualize the top contributing features per outcome, providing insights into which variables consistently influence model decisions in Sect. “[Explainable AI](#)”.

Experimental designs

Validation method and performance metrics

Many researchers have studied suitable evaluation metrics for classification tasks within the healthcare domain

[59]. They have emphasized the importance of not relying solely on a subset of metrics, as doing so could potentially yield misleading results when implementing models in clinical settings. To comprehensively assess model performance, we employ multiple evaluation metrics, including Sensitivity, Specificity, Area Under the Receiver Operating Characteristic Curve (AUROC) and Area Under the Precision Recall Curve (AUPRC).

Specificity places a stronger emphasis on correctly classifying negative samples, while sensitivity aims to minimize the misclassification of positive instances, making it a critical metric in medical studies [59]. *AUROC* and *AUPRC* has gained popularity in healthcare ML due to their favorable properties when dealing with imbalanced classes. They are used to distinguish between positive and negative classes across different threshold settings regardless of class distribution, making it an effective balance measurement.

To assess the robustness of our models, we employ cross-validation techniques, particularly stratified 5-fold cross-validation. This allows consistent distribution of classes, ensuring stability in the benchmarking results. Following the guidance of [59], we avoid sharing data from the same patients across different folds to prevent introducing bias during the parameter tuning phase. We report the mean and standard deviation of model performance across these folds, and predictions made on the test set.

Baseline comparison

We benchmark a broad set of Transformer-based architectures alongside traditional machine learning models tailored for tabular and EMR-based patient data. These models range from established tree-based ensembles to recently proposed neural architectures. The comparison results are summarized in Table 2.

XGBoost [60], **LightGBM** [61], and **CatBoost** [62]. These gradient boosting frameworks serve as strong baselines for structured data. Patient records are represented using frequency-based vectors of categorical medical codes combined with aggregate statistics of numerical features.

ResNet. We implement a ResNet-based model following the tabular adaptation proposed in [63], which demonstrates that residual connections in MLP architectures can achieve competitive performance. In their benchmarks, this approach sometimes outperforms more complex attention-based models.

TabNet [35]. TabNet uses sequential attention to select salient features at each decision step, enabling both interpretability and efficient learning. Its encoder comprises feature transformers, attentive transformers, and a feature-masking mechanism. The decoder reconstructs inputs via a sequence of feature transformer blocks.

¹ we leveraged model from <http://huggingface.com/nlpie/tiny-biobert>

Table 2 Benchmark performance of baseline and transformer-based models on three clinical outcome prediction tasks using AUROC / AUPRC

Model	30-day Readmission		In-Hospital Mortality		Length of Stay
	AUROC ↑	AUPRC ↑	AUROC ↑	AUPRC ↑	RMSE ↓
XGBoost	0.824 (0.010)	0.432 (0.020)	0.702 (0.032)	<u>0.024</u> (0.005)	20.328 (2.306)
CatBoost	0.827 (0.005)	0.446 (0.015)	0.714 (0.015)	<u>0.024</u> (0.028)	16.066 (0.170)
LightGBM	0.814 (0.055)	0.438 (0.065)	0.713 (0.028)	<u>0.024</u> (0.001)	16.703 (0.286)
TabNet [35]	0.827 (0.005)	0.440 (0.009)	0.732 (0.006)	<u>0.024</u> (0.001)	54.766 (26.21)
ResNet [37]	<u>0.830</u> (0.001)	<u>0.457</u> (0.006)	0.697 (0.006)	0.023 (0.001)	16.161 (0.086)
TabPFN [38]	0.828 (0.002)	0.455 (0.004)	0.700 (0.010)	<u>0.024</u> (0.003)	<u>15.713</u> (0.018)
TabTransformer [36]	0.834 (0.004)	0.461 (0.007)	<u>0.728</u> (0.022)	<u>0.024</u> (0.003)	15.463 (0.109)
CoreBEHRT [32]	0.678 (0.001)	0.309 (0.003)	0.619 (0.030)	0.017 (0.001)	46.82 (0.509)

The results are trained using cross-validation and benchmarked on our data testset

Sparse attention via *sparsemax* encourages the model to focus on a subset of input dimensions, providing interpretable insights into feature relevance.

TabTransformer [36]. TabTransformer leverages self-attention to generate contextual embeddings for categorical features before combining them with continuous inputs in a downstream MLP. Its column-wise Transformer encoder is particularly effective in modeling high-cardinality categorical data, which is common in EMRs.

TabPFN [38]. TabPFN is a few-shot tabular classifier that emulates Bayesian inference using a Transformer trained entirely on synthetic data. It directly outputs calibrated class probabilities without requiring fine-tuning. Though originally designed for small-scale datasets, we include it as a fast and lightweight baseline to assess how well strong priors generalize to real-world healthcare data, despite its limit of 500 features and 10,000 training samples.

CoreBEHRT [32]. CoreBEHRT builds on the BEHRT architecture by modeling visit-level sequences using age embeddings, time gaps, and discretized medical values. While powerful for temporal EMR data, our application context is limited by coarse-grained (bed-day level) timestamps and the absence of lab or medication data. Consequently, CoreBEHRT only utilizes clinical code sequences, which may restrict its effectiveness compared to other models with richer feature sets.

Handling class imbalance

Patient clinical outcomes, such as 30-day readmission or mortality or CPE-related outcomes, are often rare, leading to highly imbalanced datasets. To mitigate this issue, we employed a range of training strategies:

- **Focal Loss** [64] is used as the primary classification loss, which down-weights easy negatives and focuses learning on hard, minority-class examples.
- **Random Oversampling** is applied during training, as shown to be effective in our previous work [12],

ensuring minority-class samples are adequately represented.

- **Balanced Batch Sampling** is used to maintain class parity within mini-batches, ensuring that gradients are not dominated by the majority class, especially helpful for deep neural networks.

These methods help prevent the model from becoming biased toward negative (non-readmission or non-mortality) outcomes and improve sensitivity to rare events. Evaluation across folds confirms that these strategies contribute to more stable training and improved detection of high-risk patients.

Results

Quantitative results

Patient outcome forecast

We evaluate a broad set of baseline and Transformer-based models on three clinically relevant outcome prediction tasks: 30-day readmission, in-hospital mortality, and length of stay (LOS). Results are summarized in Table 2. All models were trained using cross-validation and evaluated on a held-out test set, ensuring consistent comparisons across architectures.

As can be seen from the table, Transformer-based models (apart from CoreBEHRT) consistently outperform traditional gradient boosting methods (XGBoost, CatBoost, LightGBM) across nearly all tasks. In particular, **TabTransformer** achieves the best overall performance—achieving the highest AUROC and AUPRC for 30-day readmission and the lowest RMSE for length-of-stay prediction. These results highlight its ability to effectively capture complex tabular dependencies in EMRs.

TabNet also demonstrates competitive performance, particularly on the in-hospital mortality task, where its attention-driven feature selection appears well-suited for distinguishing high-risk patients. **ResNet** and **TabPFN** follow closely behind, with **TabPFN** showing impressive results on LOS prediction despite its data constraints

(limited to 10,000 samples and 500 features), indicating strong potential in resource-constrained settings.

In contrast, **CoreBEHRT** underperforms across all tasks, particularly on readmission and LOS prediction. This may be due to the model's focusing exclusively on sequential ICD codes—without incorporating aggregated hospital-level features such as ward transfers, demographics, or contact networks. These findings reinforce the importance of integrating both structured clinical codes and broader contextual features when modeling real-world hospital outcomes. One of the reasons for this may be the relatively coarse granularity of the ward-based contact measures and perhaps room or bed-based measures would have more impact, if that data was available.

CPE infection forecast

In addition to forecasting patient outcome, we evaluated the feasibility of predicting CPE acquisition risk at admission by training models on screened patients, using only input features unrelated to HCAs. Specifically, we excluded any downstream infection or outcome labels and constructed the prediction task using pre-admission features (e.g., patient history, current admission ward) and in-hospital contact network metrics derived from ward co-location. This setup simulates a real-world scenario in which hospitals must assess CPE carriage risk early using available admission and contact information, before any lab-confirmed diagnosis is made, before any clinical confirmation. We further leverage explainability methods to assess the influence of contact-based versus pre-admission features, identifying whether the model learns transmission-relevant signals from network structure.

Table 3 Benchmark performance of baseline and transformer-based models on CPE risk estimation task using AUROC / AUPRC

Model	AUROC	AUPRC	Sensitivity	Specificity
XGBoost	0.755 (0.043)	0.036 (0.019)	0.075 (0.081)	0.960 (0.070)
CatBoost	0.645 (0.078)	0.027 (0.025)	0.028 (0.020)	<u>0.996</u> (0.002)
LightGBM	0.719 (0.092)	0.030 (0.021)	0.034 (0.028)	0.996 (0.005)
TabNet [35]	0.558 (0.083)	0.111 (0.055)	0.125 (0.031)	0.996 (0.004)
ResNet [37]	<u>0.773</u> (0.030)	0.046 (0.030)	<u>0.150</u> (0.125)	0.969 (0.039)
TabPFN [38]	0.719 (0.051)	0.022 (0.010)	0.000 (0.000)	1.000 (0.000)
TabTransformer [36]	0.781 (0.057)	<u>0.080</u> (0.017)	0.425 (0.211)	0.909 (0.071)
CoreBEHRT [32]	0.522 (0.127)	0.015 (0.008)	0.000 (0.000)	1.000 (0.000)

All models evaluated on the same held-out test set using cross-validation setting

As shown in Table 3, Transformer-based models such as TabTransformer and ResNet outperform traditional boosting methods across AUROC and sensitivity, with TabTransformer achieving the highest AUROC (0.781) and notably strong sensitivity (0.425), whereas most other models fall below 0.2, highlighting its robustness in identifying true positives in a highly imbalanced setting.

While AUPRC remains low due to extreme class imbalance, TabNet demonstrates a competitive edge in precision (AUPRC=0.111), suggesting that attention-based tabular models may offer utility in identifying a smaller pool of high-risk individuals. Notably, CoreBEHRT underperforms due to its lack of access to the full feature set, reinforcing the need for aggregated representations beyond sequential clinical codes. These results demonstrate that early risk estimation of CPE is feasible using contact-informed features, paving the way for targeted screening and timely IPC intervention.

Explainable AI

Global feature importance

Patient outcome tasks

To address our first research question—“*How do infection-related features impact on patient outcomes such as 30-day readmission, in-hospital mortality, and future LOS?*”, we analyze feature importance across all outcome tasks using IG.

Figures 3 and 4 present the rankings of input features based on IG attribution scores, aggregated across five cross-validation folds and three model architectures (TabTransformer, ResNet, TabNet). Features are ranked by their median importance across models and folds, where lower ranks indicate greater predictive influence.

Firstly, we analyze the attribution scores of CPE-related clinical codes in Fig. 3 to assess their contribution to outcome prediction. On average, these ICD-10-CM codes rank relatively low—typically beyond position 650 out of approximately 1,000 total codes—indicating limited standalone impact on the model's predictions. Among the outcomes, the length of stay (“NEXT_LOS”) task shows slightly higher variability, with a few instances where CPE-related codes receive notably higher importance rankings. This suggests that while these features may influence predictions in specific contexts, their global contribution remains limited.

While patient contact network features were included to provide contextual information about intra-hospital dynamics, their overall contribution to outcome prediction appears moderate. For 30-day readmission, in-hospital mortality, and length of stay, variables such as “Ward Closeness Centrality”, “Ward PageRank”, “Ward Degree Centrality”, and “Episode Isolated” generally rank between 15 and 25

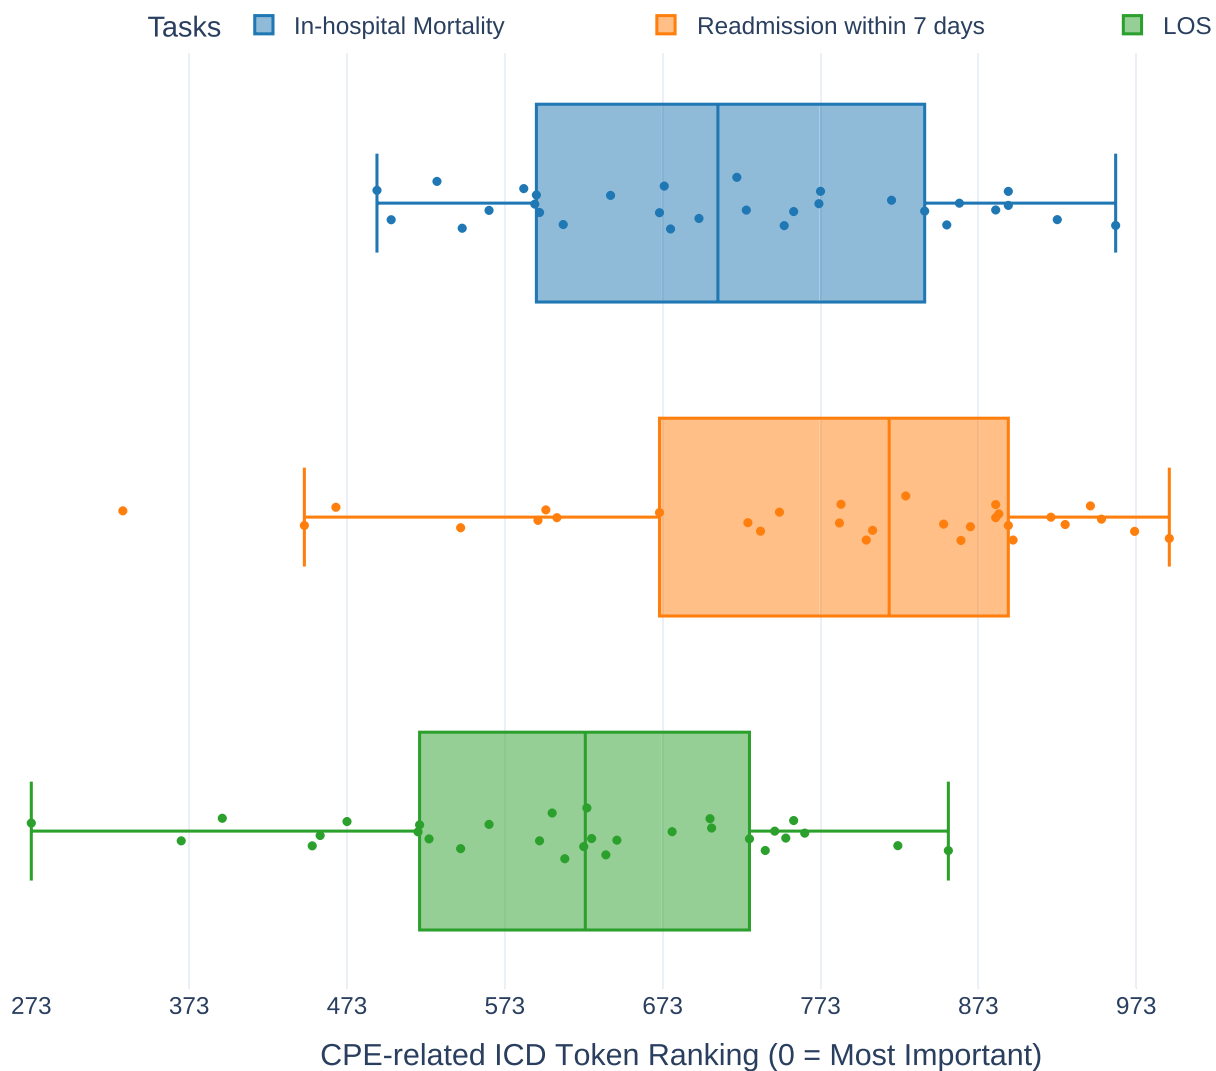


Fig. 3 Impact of CPE-related ICD-10 CM diagnosis codes on patient outcome modeling using IG

out of over 40 features, suggesting a mid-range importance across tasks.

These results indicate that contact network metrics offer supplementary value to the models, particularly in capturing structural aspects of patient flow. However, they are not among the most consistently influential features across the dataset. The wide interquartile ranges observed for some of these variables suggest that their relevance may vary by context or patient subgroup—potentially becoming more prominent in localized settings. Overall, their contribution appears to complement, rather than dominate, the predictive utility of other clinical and demographic features.

CPE infection forecasting

To address our second research question—“*Which patient features are the most predictive of CPE infection risk?*”—we assess which features contributed most to CPE acquisition prediction using the same method

and report in (Fig. 4d). The top-ranked features included “Area of Residence” and “Admission Ward”, suggesting that the model found contextual and ward-specific entry points informative, as found in the list of CPE risk factors in Table 1. “Ward PageRank”, a network centrality measure, also ranked highly, indicating that ward-level contact structure may contain useful signals for identifying potential acquisition risk.

Additionally, several historical features—such as “No. Previous Readmissions”, “Previous LOS Total”, and “No. Previous Diagnoses”—consistently appeared among the top contributors. This reflects the model’s use of longitudinal patient information when estimating risk. While these patterns highlight potentially important data signals, they represent model-driven associations and should not be interpreted as causal or clinical risk factors without further validation.

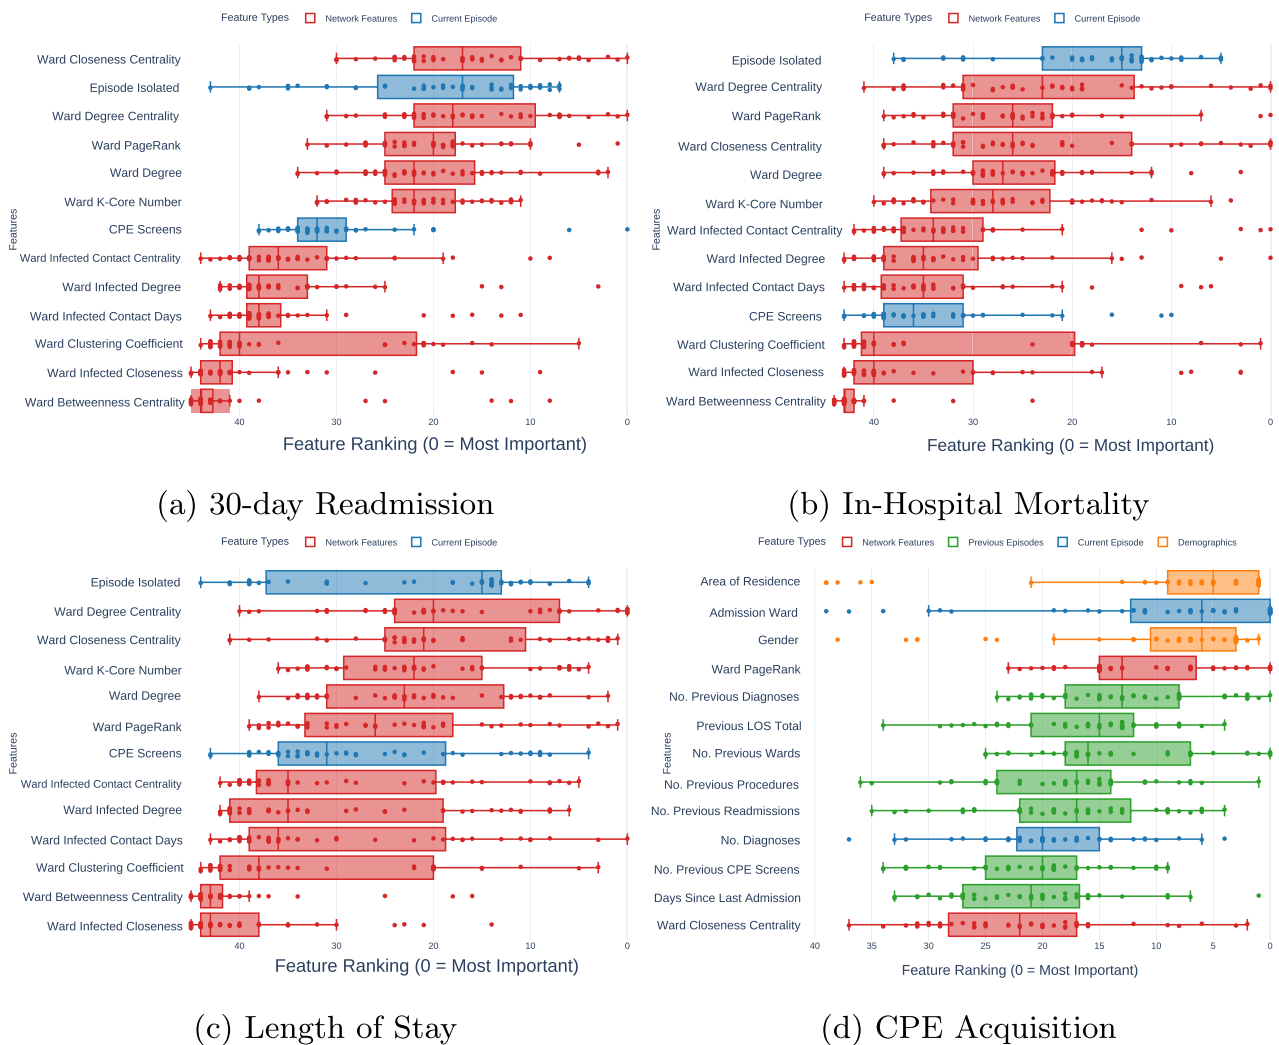


Fig. 4 Feature importance rankings based on integrated gradients for each prediction task. Shown in (a), (b) and (c) are network-related and infection-related features (definitions can be found in Appendix). Whereas, (d) demonstrates top predictive features of CPE acquisition from our benchmarked models. Lower ranks indicate higher importance. ■ Demographics, ■ Network features, ■ Past episode features, ■ Current episode features

Embedding projection

Beyond population-level analysis, we explore cohort-level representations using embedding visualizations of both clinical codes and patient episodes, shown in Figs. 5 and 6.

We project patient-level embeddings from the best-performing model in Table 3 (TabTransformer) using t-SNE. As shown in Fig. 5, episodes involving CPE-positive patients are visibly clustered together on one side of the embedding space. This occurs despite CPE status not being used as an explicit training label, indicating that the model implicitly learns meaningful clinical patterns related to CPE carriage. Such clustering highlights the model’s ability to recognize between patient cohorts, proving its training capability.

Figure 6 illustrates a t-SNE projection of pretrained embeddings for the ten most frequent ICD code groups

in our dataset, following grouping conventions from [56]. Codes are aggregated by their highest-level categories and projected into 2D space using t-SNE. The resulting plot reveals a clear separation between different diagnostic categories, suggesting that the embeddings capture semantic relationships among codes. The “infectious and parasitic” group, shown in black, forms a contiguous cluster, though some codes appear closer to other categories—indicating shared or overlapping characteristics. This visual structure supports the use of pretrained embeddings in representing clinical semantics and inter-code dependencies.

Together, these visualizations suggest that deep learning models can capture latent representations that reflect real-world clinical structure and infection risks. Embedding projections offer a complementary perspective to traditional attribution methods by revealing




Fig. 5 t-SNE projection of learned patient episode embeddings from TabTransformer. Colors denote CPE status: ■ CPE-, ■ CPE+

global patterns in patient populations and code semantics. These findings support the model's ability to detect infection-related characteristics even in the face of class imbalance, and demonstrate their value for visualizing patient flow and infection risk at scale—both essential for hospital IPC decision-making.

Personalized explanation

In a recent study, Duell et al. proposed a modified version of IG tailored to capture temporal dynamics in patient records, enabling the analysis of evolving patient characteristics over time [65]. In contrast, our approach treats each hospitalization episode independently, as our features are aggregated as static data rather than modeled as temporal sequences. To illustrate the potential impact of CPE on frequent readmissions, we present a case study of a representative patient in Fig. 7. While this individual example suggests that CPE-related risks may contribute to multiple short-term readmissions, it should be emphasized that this observation cannot be generalized due to the limited number of CPE-positive cases in the dataset.

Figure 7 visualizes the attention dynamics across sequential episodes for the selected patient, highlighting how the model assigns importance values to different

ICD code groups over time. The heatmap's color intensity indicates the contribution of each group to the predicted risk of readmission for the following episode.

In this sample, we identify three distinct hospitalization phases in the patient's trajectory: episodes t_0-t_3 , t_4 , and t_5 . In the first phase, increased IG values to CPE-related codes coincides with a cluster of closely spaced readmissions, suggesting possible CPE acquisition and microbiological diagnosis. Additionally, in most phases, diagnoses related to the nervous system and digestive appear more prominently. This example demonstrates the potential of attribution methods to support personalized risk profiling and retrospective investigation of recurrent readmission drivers.

Discussion

This study presents a practical application of deep tabular learning and Transformer-based models for assessing CPE-associated risks using real-world EMR data. Rather than proposing novel architectures, our contribution lies in developing a flexible explainable AI (XAI) framework that integrates diverse patient features to support infection prevention and control (IPC) efforts.




Fig. 6 t-SNE projection of pretrained embeddings for the top 10 most frequent high-level ICD code groups, based on [56]

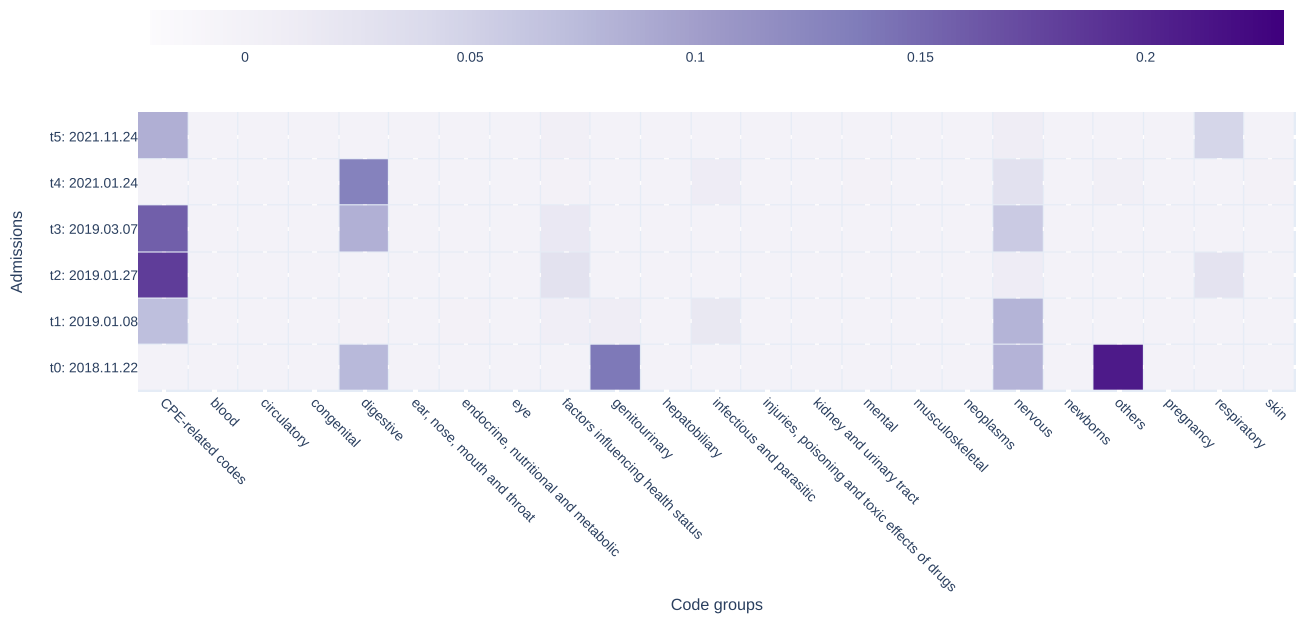


Fig. 7 Heatmap illustrating the importance of features for a randomly selected patient with multiple readmissions

Transformer-based models, particularly TabTransformer, consistently outperformed traditional baselines in AUROC across all predictive tasks. However, performance on precision-recall metrics remained modest due to the dataset's inherent class imbalance. Since this imbalance is inherent to many rare event risks like CPE, this may limit practical applications without new ways to aggregate data such as multi-microbial event studies such as TREE [66]. For the CPE acquisition task, TabTransformer also achieved the highest sensitivity score, suggesting its potential utility in identifying high-risk patients who may otherwise be missed under rule-based screening criteria. While not a diagnostic tool, such a model could assist IPC staff in prioritizing patients for early screening, potentially based on admission information ("Area of Residence", "Admission Ward"), and historical indicators such as number of prior admissions and diagnoses, thereby reducing manual workload and enabling more proactive surveillance.

To address data imbalance, we applied targeted mitigation strategies, including focal loss and balanced sampling, alongside the introduction of contact network-derived features. These variables provide spatial and contextual insight by capturing patient movement and potential exposure pathways within the hospital. Although these network features did not consistently rank among the top predictors across all tasks, their presence in mid-tier importance and variability across folds suggest that they may offer complementary value in specific scenarios.

Our flexible use of Integrated Gradients for interpretability allows deeper inspection of model predictions. Feature attribution analyses revealed that CPE-related clinical codes had limited contribution to outcome prediction—likely a reflection of the small number of confirmed cases in the dataset. Nonetheless, several literature-supported risk factors for CPE acquisition (e.g., prior admissions, length of stay, admission ward) were ranked highly in the acquisition task, demonstrating that the model can recover meaningful associations from daily EMR data. Additionally, embedding visualizations suggest that the model learns distinguishable representations for CPE-positive patients, further supporting its potential for cohort-level stratification.

Finally, in practice, this framework can be integrated as a risk-scoring module into hospital IPC systems. Given EMR and patient location data, the system outputs: (1) impact of CPE prevalence on patient outcomes (2) risk of CPE acquisition and (3) correlated features behind each prediction. Moreover, an interpretability dashboard can flag patients who may benefit from targeted screening or preemptive isolation—especially those falling outside national criteria. For example, a patient with frequent short admissions across multiple wards and high contact

network exposure might be flagged for screening, even if they do not meet criteria such as previous inpatient care in the last twelve months or ICU admission. This bridges the gap between prediction and action, aligning model capabilities with IPC goals.

Limitations & future works

The short timeframe of our dataset further exacerbates data imbalance, with CPE-positive cases representing only 0.2% of all recorded episodes. Data imbalance is a major concern when using AI technologies to deal with real-world AMR datasets since it can cause ML models to output inconsistent and unreliable results [67]. The insufficient sample size hinders the DL model's ability to effectively learn and draw robust conclusions about HAI-related risks. Fitzpatrick et. al [15] also points out the requirement for a high-quality representative dataset to develop accurate models for HCAI surveillance. In our dataset, the quality of the clinical coding is highly dependent on the information in the discharge summary that were collected, which we found to be inconsistent among different coders, but effects can be mitigated by generalizing the codes. Moving forward, it is crucial to address this challenge by prioritizing the collection of more extensive and high-quality datasets. This will enable the development of more accurate and reliable models, ultimately improving clinical risk management and patient outcomes. Exploring aggregation across related infection events (e.g., multi-microbial studies) could also help improve signal for model training. Nonetheless, we believe that our proposed model and XAI techniques hold promise for application in other desired patient groups.

A promising direction to address these limitations lies in the use of autoregressive longitudinal modeling. Our recent work, NOSOS [68], demonstrates that foundation models designed for token-by-token generation can learn temporal dynamics in patient health trajectories more robustly, even in imbalanced settings. These models treat EMR data as a generative sequence and have shown resilience to missing data and rare outcomes. Integrating such generative approaches with infection prevention and control tasks may enable more accurate unsupervised forecasting of HCAs over time, particularly in resource-constrained or low-prevalence settings.

Another avenue worth exploring is the refinement of TabPFN-based models [38]. Although TabPFN performed surprisingly well under the small-data regime, it is currently limited by constraints on input dimensionality and sample size. With recent developments in scaling few-shot learning for tabular data, future work could investigate whether adapting or retraining TabPFNs on larger, healthcare-specific priors can further improve

their generalization and performance in imbalanced infection datasets.

Another important consideration is data privacy and generalizability. Our analysis is constrained to a single Irish hospital, and patient privacy protections prevented linking patient data across hospital systems. This limitation introduces a risk of mislabeling—for instance, if patients were readmitted elsewhere without corresponding records. Therefore, future work should explore multi-institutional datasets, such as MIMIC-III/IV [69, 70], to validate our findings in broader settings and ensure more robust generalization across clinical contexts. Moreover, while AUROC values were strong across tasks, the relatively modest precision-recall performance highlights a limitation for practical applications, where high precision is critical for identifying small groups of high-risk patients. Precision-oriented methods such as calibrated thresholds, cost-sensitive learning, or ensemble strategies should be explored to better support real-world clinical decision-making.

Conclusion

This study presents an explainable AI framework for investigating CPE-related risks and outcomes using structured EMR data from an Irish acute hospital. We benchmarked a suite of traditional and Transformer-based models, finding TabTransformer to consistently outperform baselines across multiple clinical prediction tasks, particularly in AUROC and sensitivity for CPE acquisition. The integration of diverse patient features—including ward history, contact networks, and clinical codes—allowed for a comprehensive analysis of infection-related outcomes. Throughout, explainability analysis was conducted via the flexible use of Integrated Gradients and embedding visualizations, enabling transparent model interpretation. Our study addressed four key research questions:

(RQ1) Our findings show that infection-related features, including historical hospital exposure and admission context, are moderately associated with patient outcomes such as readmission, mortality, and LOS. Contact network variables—designed to capture in-hospital exposure—ranked mid-tier in feature importance, suggesting contextual relevance in certain scenarios but not dominating model predictions. CPE diagnosis codes, due to data scarcity, had limited impact on outcome forecasting. **(RQ2)** For predicting CPE acquisition risk, the most influential features included “Area of Residence”, “Admission Ward”, and historical indicators such as number of prior admissions and diagnoses. These features align with known CPE risk factors in the literature, indicating that the model successfully recovers associations from routine EMR data. Network centrality measures like “Ward PageRank” also ranked highly,

reflecting the potential value of structural exposure information.

(RQ3) Our results show that explainable AI techniques, particularly Integrated Gradients and embedding projections, enable multi-level insights into the model’s handling of CPE-positive patients. At the cohort level, CPE-positive episodes form distinguishable clusters in the embedding space, suggesting the model can capture latent characteristics associated with this subgroup. At the individual level, attribution visualizations reveal how diagnostic histories and contact-related features contribute to predictions, offering interpretable signals that may support more tailored IPC interventions. While these patterns are not clinically validated, they highlight the potential of XAI to support case-level audit and decision-making in real-world hospital workflows.

Overall, our framework offers a practical, interpretable pipeline for infection risk estimation that could be deployed in hospital IPC systems to guide early screening, risk scoring, and resource allocation.

Abbreviations

AI	Artificial intelligence
DL	Machine Learning
DL	Deep Learning
CPE	Carbapenemase-Producing Enterobacteriaceae
HCAI	Healthcare-Associated Infection
EMRs	Electronic Medical Records
XAI	Explainable AI
AUROC	Average area under the receiver operating characteristic curve
AMR	Antimicrobial resistance
IG	Integrated Gradients
HIPE	Hospital Inpatient Enquiry
ICD-10CM	International Classification of Diseases Codes version 10 Clinical Modification
IPC	Infection Prevention and Control

Supplementary Information

The online version contains supplementary material available at <https://doi.org/10.1186/s12911-025-03214-1>.

Supplementary Material 1

Acknowledgements

Not applicable.

Author contributions

MKP: Involved in all aspects of this study. MKP, MC and MB: Conceptualization. MEW, UG, DB, BO, CB, DC: Data acquisition and interpretation. MKP, TTM, MC, RB, MEW, UG, DB, BO, CB, DC, NM, MB: revision of the manuscript. MKP, MEW, UG, DB, BO, CB, DC: Result validation and interpretation. All authors read and approved the final manuscript.

Funding

This research was conducted with the financial support of Taighde Éireann-Research Ireland under Grant Agreement No. 13/RC/2106_P2 at ADAPT, the Research Ireland Centre for AI-Driven Digital Content Technology at DCU funded through the Research Ireland Research Centres Programme. For the purpose of Open Access, the author has applied a CC BY public copyright licence to any author-accepted manuscript version arising from this submission.

Data availability

The datasets generated or analysed during the current study are not publicly available due to restrictions imposed by institutional ethics approvals and data protection regulations but are available from the corresponding author on reasonable request. The code and reproducibility guidelines supporting the analyses are publicly available at <https://github.com/kaylode/carbapen>.

Declarations

Ethics approval and consent to participate

The study was conducted using a real-world anonymized dataset that was shared as part of the ARK risk management research project. Our study was approved by the St James Hospital's institutional review board (TUH/SJH REC Project ID: 0291) and by Dublin City University Research Ethics Committee, Dublin, Ireland (DCUREC/2021/118). The requirement for informed consent was waived by the Ethics Committee of St James Hospital's institutional review board because of the retrospective nature of the study and was also considered unnecessary according to national regulations at the time of the research. All methods were carried out in accordance with relevant guidelines and regulations, and in line with the ethical principles of the WMA Declaration of Helsinki (as amended in 2024), with strict safeguards to ensure patient confidentiality and privacy.

Consent for publication

Not applicable.

Competing interests

The authors declare no competing interests.

Author details

¹School of Computing, Dublin City University, Dublin, Ireland

²ADAPT Centre, Dublin, Ireland

³University College Dublin, Dublin, Ireland

⁴St James's Hospital, Dublin, Ireland

⁵Health Service Executive (HSE), Dublin, Ireland

⁶Trinity College Dublin, Dublin, Ireland

Received: 5 August 2024 / Accepted: 18 September 2025

Published online: 22 October 2025

References

- Codjoe FS, Donkor ES. Carbapenem resistance: a review. *Med Sci*. 2017;6(11):1. <https://doi.org/10.3390/medsci610001>.
- Haque M, Sartelli M, McKimm J, et al. Health care-associated infections—an overview. *Infect Drug Resist*. 2018;2321–33. <https://doi.org/10.2147/IDR.S177247>.
- Safety WP, Organization WH, et al. Conceptual framework for the international classification for patient safety version 1.1: final technical report January 2009. Technical report. World Health Organization. 2010. <https://www.who.int/publications/i/item/WHO-HER-PSP-2010.2>.
- Vellinga A, Brennan W, Humphreys H, et al. Initial impact of a national programme to contain the spread of carbapenemase-producing enterobacteriales in Ireland. *J Hosp Infect*. 2021;109:107–14. <https://doi.org/10.1016/j.jhin.2021.01.001>.
- O'Connell NH, Gasior S, Slevin B, et al. Microbial epidemiology and clinical risk factors of carbapenemase-producing enterobacteriales amongst Irish patients from first detection in 2009 until 2020. *Infect Prev Pract*. 2022;4(3):100230. <https://doi.org/10.1016/j.infpip.2022.100230>.
- Humphreys H, Cormican M, Brennan W, et al. Reflections on a national public health emergency response to carbapenemase-producing enterobacteriales (cpe). *Epidemiol Infect*. 2022;150:69. <https://doi.org/10.1017/S0950268822000218>.
- Fahy S, O'Connor J, O'Brien D, et al. Carbapenemase screening in an Irish tertiary referral hospital: best practice, or can we do better? *Infect Prev Pract*. 2020;2(4):100100. <https://doi.org/10.1016/j.infpip.2020.100100>.
- Kardas-Sloma L, Fournier S, Dupont J-C, et al. Cost-effectiveness of strategies to control the spread of carbapenemase-producing enterobacteriales in hospitals: a modelling study. *Antimicrob Resist Infect Control*. 2022;11(1):117. <https://doi.org/10.1186/s13756-022-01149-0>.
- Ward ME, Geary U, Brennan R, et al. A systems approach to managing the risk of healthcare acquired infection in an acute hospital setting supported by human factors ergonomics, data science, data governance and ai. *Ergonomics*. 2024;1–19. <https://doi.org/10.1080/00140139.2024.2396527>.
- McDonald N, McKenna L, Vining R, et al. Evaluation of an access-risk-knowledge (ark) platform for governance of risk and change in complex socio-technical systems. *Int J Environ Res Pub Health Public Health*. 2021;18(23):12572. <https://doi.org/10.3390/ijerph182312572>.
- Crotti Junior A, Basereh M, Abgaz Y, et al. The ARK platform: enabling risk management through semantic web technologies. *CEUR-WS*. 2020. <https://ceur-ws.org/Vol-2807/paperM.pdf>
- Pham M-K, Mai TT, Crane M, et al. Forecasting patient early readmission from Irish hospital discharge records using conventional machine learning models. *Diagnostics*. 2024;14(21):2405. <https://doi.org/10.3390/diagnostics14212405>.
- Corson A, Byock I, Dale CR. Hospital-acquired infections and readmissions: let's refocus on the person. *Am J Med Qual*. 2020;35(1):89–90. <https://doi.org/10.1177/1062860619863105>.
- Emerson CB, Eyzaguirre LM, Albrecht JS, et al. Healthcare-associated infection and hospital readmission. *Infect Control Hosp Epidemiol*. 2012;33(6):539–44. <https://doi.org/10.1086/665725>.
- Fitzpatrick F, Doherty A, Lacey G. Using artificial intelligence in infection prevention. *Curr Treat Options Infect Dis*. 2020;12:135–44. <https://doi.org/10.1007/s40506-020-00216-7>.
- Yang L, Lu S, Zhou L. The implications of artificial intelligence on infection prevention and control: current progress and future perspectives. *China CDC Wkly*. 2024;6(35):901. <https://doi.org/10.46234/ccdcw2024.192>.
- Abu-El-Ruz R, AbuHaweeleh MN, Hamdan A, et al. Artificial intelligence in bacterial infections control: a scoping review. *Antibiotics*. 2025;14(3):256. <https://doi.org/10.3390/antibiotics14030256>.
- Caglayan C, Barnes SL, Pineles LL, et al. A data-driven framework for identifying intensive care unit admissions colonized with multidrug-resistant organisms. *Front Public Health*. 2022;10. <https://doi.org/10.3389/fpubh.2022.853757>.
- Wang Y, Wang G, Zhao Y, et al. A deep learning model for predicting multidrug-resistant organism infection in critically ill patients. *J Intensive Care*. 2023;11(1):49. <https://doi.org/10.1186/s40560-023-00695-y>.
- Myall A, Price JR, Peach RL, et al. Prediction of hospital-onset COVID-19 infections using dynamic networks of patient contact: an international retrospective cohort study. *Lancet Digit Health*. 2022;4(8):573–83. [https://doi.org/10.1016/S2589-7500\(22\)00093-0](https://doi.org/10.1016/S2589-7500(22)00093-0). Accessed 2023-10-12.
- Gouareb R, Bornet A, Proios D, et al. Detection of patients at risk of enterobacteriaceae infection using graph neural networks: a retrospective study. *medRxiv*. 2023;2023–26. <https://doi.org/10.1101/2023.06.01.23290386>.
- Nayak DSK, Das R, Sahoo SK, et al. Argai 1.0: a gan augmented in silico approach for identifying resistant genes and strains in e. coli using vision transformer. *Comput Biol Chem*. 2025;115:108342. <https://doi.org/10.1016/j.cmpbiolchem.2025.108342>.
- Segagni Lusignani L, Presterl E, Zatorska B, et al. Infection control and risk factors for acquisition of carbapenemase-producing enterobacteriaceae. A 5 year (2011–2016) case-control study. *Antimicrob Resist Infect Control*. 2020;9:1–10. <https://doi.org/10.1186/s13756-019-0668-2>.
- Goodman KE, Simner PJ, Klein EY, et al. Predicting probability of perirectal colonization with carbapenem-resistant enterobacteriaceae (cre) and other carbapenem-resistant organisms (cros) at hospital unit admission. *Infect Control Hosp Epidemiol*. 2019;40(5):541–50. <https://doi.org/10.1017/ice.2019.42>.
- Freire MP, Rinaldi M, Terrabuio DRB, et al. Prediction models for carbapenem-resistant enterobacteriales carriage at liver transplantation: a multicenter retrospective study. *Transpl Infect Dis*. 2022;24(6):13920. <https://doi.org/10.1111/tid.13920>.
- Liang Q, Ding S, Chen J, et al. Prediction of carbapenem-resistant gram-negative bacterial bloodstream infection in intensive care unit based on machine learning. *Bmc Med Inf Decis*. 2024;24(1):123. <https://doi.org/10.1186/s12911-024-02504-4>.
- Liang Q, Zhao Q, Xu X, et al. Early prediction of carbapenem-resistant gram-negative bacterial carriage in intensive care units using machine learning. *J Glob Antimicrob Resist*. 2022;29:225–31. <https://doi.org/10.1016/j.jgar.2022.03.019>.
- Disease Prevention, E.C. Control: risk assessment on the spread of carbapenemase-producing Enterobacteriaceae (CPE) through patient transfer between healthcare facilities, with special emphasis on cross-border transfer. Sweden: European Centre for Disease Prevention and Control Solna. 2011. <https://www>

- wecdc.europa.eu/sites/default/files/media/en/publications/Publications/110913_Risk_assessment_resistant_CPE.pdf
29. Rasmusy L, Xiang Y, Xie Z, et al. Med-bert: pretrained contextualized embeddings on large-scale structured electronic health records for disease prediction. *NPJ Digit Med.* 2021;4(1):86. <https://doi.org/10.1038/s41746-021-00455-y>.
 30. Li Y, Rao S, Solares JRA, et al. BEHRT: transformer for electronic health records. *Sci Rep.* 2020;10:7155. <https://doi.org/10.1038/s41598-020-62922-y>. Accessed 2023-07-12.
 31. Li Y, Mamouei M, Salimi-Khorshidi G, et al. Hi-BEHT: hierarchical transformer-based model for accurate prediction of clinical events using multimodal longitudinal electronic health records. *IEEE J Biomed Health Inf.* 2023;27(2):1106–17. <https://doi.org/10.1109/JBHI.2022.3224727>. Accessed 2024-04-28.
 32. Odgaard M, Klein KV, Thysen SM, et al. Core-behrt: a carefully optimized and rigorously evaluated behrt. *arXiv preprint arXiv:2404.15201* (2024). <https://doi.org/10.48550/arXiv.2404.15201>.
 33. Fallahpour A, Alinoori M, Ye W, et al. Ehrmamba: towards generalizable and scalable foundation models for electronic health records. *arXiv preprint arXiv:2405.14567*. 2024. <https://doi.org/10.48550/arXiv.2405.14567>.
 34. Bednarski BP, Singh AD, Zhang W, et al. Temporal convolutional networks and data rebalancing for clinical length of stay and mortality prediction. *Sci Rep.* 2022;12(1):21247. <https://doi.org/10.1038/s41598-022-25472-z>.
 35. Arik SÖ, Pfister T. Tabnet: attentive interpretable tabular learning. *Proc AAAI Conf Artif Intell.* 2021;35:6679–87. <https://doi.org/10.1609/aaai.v35i8.16826>.
 36. Huang X, Khetan A, Cvitkovic M, et al. Tabtransformer: tabular data modeling using contextual embeddings. *arXiv preprint arXiv:2012.06678*. 2020. <https://doi.org/10.48550/arXiv.2012.06678>.
 37. Gorishniy Y, Rubachev I, Khrukov V, et al. Revisiting deep learning models for tabular data. *Adv Neural Inf Process Syst.* 2021;34:18932–43. <https://doi.org/10.48550/arXiv.2106.11959>.
 38. Hollmann N, Müller S, Purucker L, et al. Accurate predictions on small data with a tabular foundation model. *Nature.* 2025;637(8045):319–26. Nature Publishing Group. <https://doi.org/10.1038/s41586-024-08328-6>. Accessed 2025-06-16.
 39. Loh HW, Ooi CP, Seoni S, et al. Application of explainable artificial intelligence for healthcare: a systematic review of the last decade (2011–2022). *Comput Met Programs Biomed.* 2022;226:107161. <https://doi.org/10.1016/j.cmpb.2022.107161>.
 40. Di Martino F, Delmastro F. Explainable AI for clinical and remote health applications: a survey on tabular and time series data. *Artif Intell Rev.* 2023;56(6):5261–315. <https://doi.org/10.1007/s10462-022-10304-3>.
 41. Chaddad A, Peng J, Xu J, et al. Survey of explainable ai techniques in healthcare. *Sensors.* 2023;23(2):634. <https://doi.org/10.3390/s23020634>.
 42. Wani NA, Kumar R, Bedi J. Deepexplainer: an interpretable deep learning based approach for lung cancer detection using explainable artificial intelligence. *Comput Met Programs Biomed.* 2024;243:107879. <https://doi.org/10.1016/j.cmpb.2023.107879>.
 43. Wani NA, Kumar R, Bedi J, et al. Explainable AI-driven iomt fusion: unravelling techniques, opportunities, and challenges with explainable ai in healthcare. *Inf Fusion.* 2024;102472. <https://doi.org/10.1016/j.inffus.2024.102472>.
 44. Sundararajan M, Taly A, Yan Q. Axiomatic attribution for deep networks. *International Conference on Machine Learning.* PMLR; 2017, pp. 3319–28. <https://doi.org/10.5555/3305890.3306024>.
 45. O’Loughlin R, Allwright S, Barry J, et al. Using hipec data as a research and planning tool: limitations and opportunities. *Ir J Med Sci.* 2005;174:40–45. <https://doi.org/10.1007/BF03169128>.
 46. Office, H.P. Hospital In-patient enquiry - data dictionary 2022 - V14.0. https://hpo.ie/hipec/hipec_data_dictionary/HIPE_Data_Dictionary_2022_V14.0.pdf
 47. Organization, W.H. et al. ICD-10. International statistical classification of diseases and related health problems: tenth revision 1992, volume 1 = CIM-10. In: *Classification statistique internationale des maladies et des problèmes de santé connexes: dixième Révision 1992, Volume 1.* WHO: World Health Organization; 1992.
 48. Deschepper M, Eeckloo K, Vogelaers D, et al. A hospital wide predictive model for unplanned readmission using hierarchical icd data. *Comput Met Programs Biomed.* 2019;173:177–83. <https://doi.org/10.1016/j.cmpb.2019.02.007>.
 49. Wang WW, Li H, Cui L, et al. Predicting clinical visits using recurrent neural networks and demographic information. 2018 IEEE 22nd International Conference on Computer Supported Cooperative Work in Design ((CSCWD)). 2018, pp. 353–58. <https://doi.org/10.1109/CSCWD.2018.8465194>.
 50. Choi E, Bahadori MT, Schuetz A, et al. Doctor AI: Predicting clinical events via recurrent neural networks. In: *Machine Learning for Healthcare Conference.* PMLR; 2016. p. 301–18. <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5341604/>
 51. Allam A, Feuerriegel S, Rebhan M, et al. Analyzing patient trajectories with artificial intelligence. *J Med Internet Res.* 2021;23(12):29812. <https://doi.org/10.2196/29812>.
 52. Choi E, Schuetz A, Stewart WF, et al. Using recurrent neural network models for early detection of heart failure onset. *J Am Med Inf Assoc.* 2017;24(2):361–70. <https://doi.org/10.1093/jamia/ocw112>.
 53. Pham T, Tran T, Phung D. Predicting healthcare trajectories from medical records: a deep learning approach. *J Educ Chang Biomed Inf.* 2017;69:218–29. <https://doi.org/10.1016/j.jbi.2017.04.001>.
 54. HSE. Management and Control of Carbapenemase Producing Enterobacteriales (CPE) in all Healthcare Settings.
 55. Moloney ED, Bennett K, Silke B. Factors influencing the costs of emergency medical admissions to an Irish teaching hospital. *Eur J Health Econ.* 2006;7:123–28. <https://doi.org/10.1007/s10198-006-0343-4>.
 56. Wu X, Zhao Y, Yang Y, et al. A comparison of representation learning methods for medical concepts in mimic-iv. *medRxiv.* 2022;2022–28.
 57. Allgaier J, Mulansky L, Draelos RL, et al. How does the model make predictions? A systematic literature review on the explainability power of machine learning in healthcare. *Artif Intel Med.* 2023;143:102616. <https://doi.org/10.1016/j.artmed.2023.102616>.
 58. Shapley LS, et al. *A value for n-person games.* Princeton University Press Princeton; 1953. <https://doi.org/10.7249/P0295>.
 59. Hicks SA, Strümke I, Thambawita V, et al. On evaluation metrics for medical applications of artificial intelligence. *Sci Rep.* 2022;12(1):5979. <https://doi.org/10.1038/s41598-022-09954-8>.
 60. Chen T, Guestrin C. Xgboost: a scalable tree boosting system. In: *Proceedings of the 22nd Acm Sigkdd International Conference on Knowledge Discovery and Data Mining.* 2016. p. 785–94. <https://doi.org/10.1145/2939672.2939785>.
 61. Ke G, Meng Q, Finley T, et al. Lightgbm: a highly efficient gradient boosting decision tree. *Adv Neural Inf Process Syst.* 2017;30. <https://doi.org/10.5555/3294996.3295074>.
 62. Prokhorenkova L, Gusev G, Vorobev A, et al. Catboost: unbiased boosting with categorical features. *Adv Neural Inf Process Syst.* 2018;31. <https://doi.org/10.5555/3327757.3327770>.
 63. Dai X, Chalkidis I, Darkner S, et al. Revisiting transformer-based models for long document classification. In: *Findings of the association for computational linguistics: EMNLP 2022.* 2022, pp 7212–30. <https://doi.org/10.18653/v1/2022.findings-emnlp.534>.
 64. Lin T-Y, Goyal P, Girshick R et al. Focal loss for dense object detection. *Proceedings of the IEEE International Conference on Computer Vision.* 2017, pp. 2980–88. <https://doi.org/10.1109/ICCV.2017.324>.
 65. Duell J, Fan X, Fu H, et al. Batch integrated gradients: explanations for temporal electronic health records. *International Conference on Artificial Intelligence in Medicine.* Springer; 2023, pp. 120–24. https://doi.org/10.1007/978-3-031-34344-5_15.
 66. Rock C, Hsu Y-J, Milstone AM, Salinas AB, Maragakis L, Kwon N, Cosgrove SE, Simner PJ, Virk Z, Vorsteg AH, et al. Electronic evaluation of infection risk from room environment and prior room occupants: the transmission from room environment events (tree) measure. *Infect Control Hosp Epidemiol.* 2025;46(2):209–12. <https://doi.org/10.1017/ice.2024.194>.
 67. Ali T, Ahmed S, Aslam M. Artificial intelligence for antimicrobial resistance prediction: challenges and opportunities towards practical implementation. *Antibiotics.* 2023;12(3):523. <https://doi.org/10.3390/antibiotics12030523>.
 68. Pham MK, Mai TT, Thi Phuong Dao T. Zero-shot detection of nosocomial infections using clinical foundation models on structured electronic patient records. *Technix preprint technix.175021755.55836808/v1.* 2025. <https://doi.org/10.36227/technix.175021755.55836808/v1>.
 69. Johnson AE, Pollard TJ, Shen L, et al. Mimic-iii, a freely accessible critical care database. *Sci Data.* 2016;3(1):1–9. <https://doi.org/10.1038/sdata.2016.35>.
 70. Johnson AE, Bulgarelli L, Shen L, et al. Mimic-iv, a freely accessible electronic health record dataset. *Sci Data.* 2023;10(1):1. <https://doi.org/10.1038/s41597-022-01899-x>.

Publisher’s Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.