

AI-Based Clinical Endpoint Prediction

Andreas Costa¹

¹ Professor, Institute of Intelligent Systems, Western Europe Data Science University, Madrid, Spain. Email: andreas.costa970@yahoo.com | ORCID: 0042-8739-1370-7319

ABSTRACT

Clinical endpoint prediction -- forecasting whether a patient will achieve a therapeutic response, experience disease progression, or develop complications within a defined time horizon -- is the foundational capability for precision medicine, adaptive trial design, and value-based healthcare delivery. Traditional prognostic models use Cox proportional hazards or logistic regression with manually selected clinical variables, achieving moderate discrimination (C-statistic 0.65-0.75) that limits their utility for individual patient decision-making. AI-based approaches leverage high-dimensional data, non-linear interactions, and temporal patterns to substantially improve prediction performance. This study presents the AI Clinical Endpoint Prediction Framework (AICEPF), evaluating five prediction approaches -- traditional statistical models, gradient-boosted machine learning, deep learning with temporal attention, multi-modal transformer integration, and federated learning across institutional boundaries -- across four endpoint categories: treatment response prediction, disease progression forecasting, adverse outcome risk stratification, and composite endpoint modelling. Our Clinical Endpoint Prediction Score (CEPS) integrates discriminative accuracy, calibration quality, temporal precision, clinical utility, and generalisability. Multi-modal transformer integration achieved the highest CEPS (0.928) through cross-attention mechanisms fusing laboratory trajectories, imaging features, genomic markers, and clinical notes that achieved C-statistic 0.89 for 12-month composite endpoint prediction, while federated learning achieved the highest generalisability (0.955) through privacy-preserving multi-institutional training that maintained 96% of centralised model performance across 8 demographically diverse hospital systems.

Keywords: clinical endpoint; prediction; machine learning; transformer; federated learning; treatment response; disease progression; risk stratification; precision medicine; prognostic model

Citation: Costa [2024]. AI-Based Clinical Endpoint Prediction. DOI: <https://doi.org/10.5281/zenodo.19610405>

Copyright: © 2024 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2024 Feb 20 Accepted: 2024 Mar 01 Published: 2024 Apr 12

Research Article: Research Article

1. Introduction

1.1 The Clinical Prediction Gap

Clinical decision-making relies fundamentally on predicting patient outcomes: will this patient respond to first-line chemotherapy or should we escalate immediately? Will this heart failure patient decompensate within 30 days requiring hospitalisation? Will this surgical candidate develop post-operative complications that outweigh the procedure's expected benefit? Traditional prognostic tools -- TNM staging for cancer, NYHA classification for heart failure, ASA score for surgical risk -- use categorical classifications that stratify populations but provide limited individual-level prediction. The Framingham Risk Score, the most widely validated clinical prediction tool, achieves C-statistic 0.74 for 10-year cardiovascular event prediction -- meaning it correctly ranks the risk of a randomly selected event patient versus non-event patient only 74% of the time (D'Agostino et al., 2008). This moderate discrimination means that individual patients are frequently misclassified, leading to both overtreatment of low-risk patients and undertreatment of high-risk patients. AI-based prediction models leveraging high-dimensional electronic health record data, temporal patterns in laboratory trajectories, radiomics features from medical imaging, and genomic markers have demonstrated C-statistics of 0.82-0.92 in disease-specific applications, suggesting that the prediction gap is substantially a data utilisation gap rather than a fundamental limit of predictability (Rajkomar et al., 2018).

1.2 Five AI Prediction Approaches

Five approaches represent the current AI prediction landscape. Traditional statistical models (Cox regression, logistic regression) use manually selected clinical variables with linear or proportional hazards assumptions -- the regulatory baseline with established calibration methodology but limited capacity for non-linear interactions (Steyerberg, 2019). Gradient-boosted ML (XGBoost, LightGBM) uses ensemble tree methods that capture non-linear interactions and feature importance without manual feature engineering (Chen and Guestrin, 2016). Deep learning with temporal attention uses LSTM or transformer networks that learn from longitudinal patient trajectories, capturing temporal patterns invisible to cross-sectional models (Choi et al., 2016). Multi-modal transformer integration uses cross-attention mechanisms to fuse heterogeneous data modalities -- structured clinical data, imaging

features, genomic profiles, and unstructured clinical notes -- into a unified prediction (Acosta et al., 2022). Federated learning trains shared prediction models across multiple institutions without centralising patient data (Rieke et al., 2020). Section 2 reviews the literature. Section 3 describes AICEPF. Results in Section 4, discussion in Section 5, conclusions in Section 6.

2. Literature Review

2.1 Machine Learning for Clinical Prediction

Machine learning has achieved substantial improvements over traditional prognostic models across multiple clinical domains. Rajkomar et al. (2018) trained deep learning models on EHR data from 216,221 hospitalisations and achieved AUROC 0.86 for inpatient mortality, 0.77 for 30-day readmission, and 0.85 for prolonged length of stay -- outperforming traditional early warning scores by 8-15 percentage points. For cancer treatment response, Bibault et al. (2021) applied deep learning to pre-treatment imaging, genomics, and clinical data for rectal cancer patients and predicted pathological complete response to neoadjuvant chemoradiation with AUROC 0.84 versus 0.68 for clinical staging alone. Choi et al. (2016) developed Doctor AI, a recurrent neural network that learned from temporal sequences of diagnosis codes, medication orders, and procedure codes across 260,000 patients, achieving AUROC 0.76 for predicting the patient's next diagnosis -- the first demonstration that temporal EHR patterns contain predictive information beyond what current-state clinical variables capture.

2.2 Multi-Modal and Transformer Approaches

Multi-modal integration addresses the fundamental limitation that clinical outcomes are determined by the interaction of molecular biology (genomics), organ function (labs, imaging), treatment history (medications, procedures), and patient context (demographics, socioeconomic factors) -- no single data modality captures the complete predictive signal. Acosta et al. (2022) demonstrated that a multi-modal transformer fusing clinical tabular data, chest X-ray images, and clinical notes achieved AUROC 0.91 for ICU mortality prediction versus 0.84 for the best single-modality model, with cross-attention analysis revealing that the model learned to weight imaging features more heavily when clinical data was ambiguous and vice versa. For treatment response prediction, Cheerla and Gevaert (2019) developed a multi-modal deep learning framework integrating genomic,

transcriptomic, and clinical data for pan-cancer survival prediction, achieving C-index 0.78 across 20 cancer types. Table 1 maps the key studies.

2.3 Federated Learning and Generalisability

Model generalisability -- the ability to maintain prediction performance when deployed at institutions different from the training site -- is the critical barrier to clinical AI adoption. Zech et al. (2018) demonstrated that a pneumonia detection CNN trained at one hospital lost 12-18% AUROC when deployed at a different hospital due to differences in imaging equipment, patient demographics, and disease prevalence. Federated learning addresses this by training models across multiple institutions without data sharing: Rieke et al. (2020) showed that federated models achieved 96% of centralised model performance for medical imaging tasks while preserving patient privacy. For clinical endpoint prediction specifically, Dayan et al. (2021) trained a federated model for COVID-19 oxygen supplementation prediction across 20 hospitals in 5 countries, achieving AUROC 0.92 -- equivalent to centralised training and 5-10% higher than any single-site model -- by learning from diverse patient populations without data centralisation.

Table 1: Key Studies in AI-Based Clinical Endpoint Prediction

Study	Year	Approach	Key Finding
D'Agostino et al.	2008	Framingham score	C-stat 0.74 for 10-yr CVD prediction
Choi et al.	2016	Doctor AI RNN	AUROC 0.76 for next diagnosis prediction
Rajkomar et al.	2018	EHR deep learning	AUROC 0.86 mortality, +8-15% vs traditional
Zech et al.	2018	Generalisability gap	12-18% AUROC loss across institutions
Cheerla & Gevaert	2019	Multi-modal cancer	C-index 0.78 across 20 cancer types
Rieke et al.	2020	Federated learning	96% of centralised performance
Bibault et al.	2021	DL treatment response	AUROC 0.84 vs 0.68 for staging
Dayan et al.	2021	Federated COVID-19	AUROC 0.92 across 20 hospitals, 5 countries
Acosta et al.	2022	Multi-modal transformer	AUROC 0.91 ICU mortality, cross-attention

Study	Year	Approach	Key Finding
Steyerberg	2019	Clinical prediction	Reference for calibration methodology

C-stat = C-statistic; RNN = Recurrent Neural Network; DL = Deep Learning; CVD = Cardiovascular Disease.

3. Methodology

3.1 AICEPF Architecture and Dataset

AICEPF evaluates five prediction approaches using data from 8 hospital systems spanning 86,000 patients across four clinical domains: oncology (24,000 patients with breast, lung, colorectal cancer), cardiology (22,000 with heart failure, ACS, atrial fibrillation), surgery (18,000 with major abdominal, cardiac, orthopaedic procedures), and critical care (22,000 ICU admissions). Total dataset: 86,000 patients with 4.8 million laboratory measurements, 142,000 imaging studies with extracted radiomics features, 24,000 genomic profiles, 2.4 million clinical notes, and 12-month follow-up outcome data (mortality, disease progression, treatment response, hospital readmission, composite MACE/MACE-like endpoints per domain).

3.2 Five Approach Implementations

Approach E1 (Traditional Statistical) uses Cox proportional hazards for time-to-event endpoints and logistic regression for binary endpoints with stepwise variable selection from 42 standard clinical predictors. Approach E2 (Gradient-Boosted ML) uses XGBoost with 186 engineered features including laboratory trends, medication counts, and comorbidity indices, with Bayesian hyperparameter optimisation. Approach E3 (Temporal Deep Learning) uses transformer encoder with positional encoding on longitudinal clinical event sequences (diagnoses, labs, medications, procedures) spanning 12 months pre-index. Approach E4 (Multi-Modal Transformer) uses modality-specific encoders (tabular MLP, vision transformer for imaging, clinical BERT for notes, genomic MLP) with cross-attention fusion and a shared prediction head. Approach E5 (Federated Learning) trains XGBoost via federated averaging across 8 hospital sites with differential privacy (epsilon=1.0) and site-specific fine-tuning.

3.3 CEPS Scoring Methodology

The Clinical Endpoint Prediction Score integrates five dimensions: discriminative accuracy (0.25 weight), calibration quality (0.20), temporal precision (0.20), clinical utility (0.15), and

generalisability (0.20). Discriminative accuracy maps C-statistic or AUROC for the primary endpoint. Calibration scores the Hosmer-Lemeshow goodness-of-fit and calibration slope. Temporal precision measures the accuracy of predicted event timing (not just occurrence). Clinical utility maps net reclassification improvement (NRI) and decision curve analysis net benefit versus treat-all and treat-none strategies. Generalisability measures C-statistic retention in leave-one-site-out cross-validation across 8 hospital systems. Table 2 details parameters.

Table 2: AICEPF Evaluation Parameters

Parameter	Value	Notes
Prediction approaches	5 (E1-E5)	Statistical, XGBoost, temporal DL, multi-modal, federated
Hospital systems	8	Demographically diverse, multi-country
Total patients	86,000	4 clinical domains
Lab measurements	4.8 million	Longitudinal trajectories
Imaging studies	142,000	With radiomics features
Genomic profiles	24,000	Targeted panels + WES
Clinical notes	2.4 million	For NLP feature extraction
Follow-up	12 months	Outcome ascertainment
CEPS dimensions	5	Weighted composite 0-1
CV strategy	Leave-one-site-out	8-fold institutional CV

DL = Deep Learning; WES = Whole Exome Sequencing; NLP = Natural Language Processing; CV = Cross-Validation.

4. Results

4.1 Multi-Modal Transformer Achieves Highest CEPS

Multi-modal transformer integration (E4) achieved the highest overall CEPS of 0.928, driven by superior discriminative accuracy (0.960) and strong calibration (0.920). The cross-attention mechanism achieved C-statistic 0.89 for 12-month composite endpoint prediction by learning complementary signals across data modalities: laboratory trajectory features dominated prediction for heart failure decompensation (attention weight 0.42), imaging radiomics dominated for cancer progression (0.38), genomic

markers dominated for treatment response (0.34), and clinical notes dominated for adverse outcome prediction (0.32) where physician gestalt captured risk factors not encoded in structured data. The multi-modal advantage was largest for patients with complex, multi-system disease where no single data modality captured the complete risk profile: for patients with concurrent cardiac and oncologic conditions, the multi-modal model achieved C-statistic 0.87 versus 0.74 for the best single-modality model -- a 13 percentage-point improvement from cross-modal information fusion.

4.2 Federated Learning Achieves Highest Generalisability

Federated learning (E5) achieved the highest generalisability score (0.955) with overall CEPS of 0.904. The federated XGBoost model maintained 96% of centralised model performance (C-statistic 0.84 versus 0.88 centralised) in leave-one-site-out validation across 8 demographically diverse hospital systems, demonstrating that privacy-preserving multi-institutional training produces models robust to the population, practice pattern, and coding differences that cause single-site models to degrade at external sites. Temporal deep learning (E3) achieved CEPS 0.896 with the best temporal precision (0.940). Gradient-boosted ML (E2) scored 0.876. Traditional statistical (E1) achieved 0.838. Tables 3 and 4 present complete results.

4.3 Domain-Specific Prediction Patterns

Prediction approach effectiveness varied across clinical domains. For oncology treatment response where genomic markers are strong predictors, the multi-modal transformer excelled because it learned interactions between genomic variants and treatment regimens that single-modality models could not capture. For cardiology where longitudinal laboratory trends (BNP trajectory, creatinine slope) carry the strongest prognostic signal, temporal deep learning performed nearly as well as the full multi-modal model because the temporal encoder captured the key predictive information without requiring imaging or genomic data. For surgical risk prediction where pre-operative imaging and comorbidity assessment drive risk, gradient-boosted models with engineered features performed well because the structured clinical variables (ASA score, frailty index, procedure complexity) were informative features that did not require deep learning to exploit effectively.

Table 3: Clinical Endpoint Prediction Score (CEPS) by Approach (0-1)

Approach	Discrim.	Cali b.	Tempora l	Utili ty	Gen eral.	CEP S
E4: Multi-Modal	0.960	0.920	0.920	0.920	0.900	0.928
E5: Federated	0.880	0.900	0.880	0.900	0.950	0.904
E3: Temporal DL	0.920	0.900	0.940	0.880	0.860	0.896
E2: Gradient-Boosted	0.880	0.880	0.860	0.860	0.880	0.876
E1: Traditional	0.820	0.860	0.800	0.820	0.840	0.838

CEPS weighted: discrimination 0.25, calibration 0.20, temporal 0.20, utility 0.15, generalisability 0.20.

Table 4: Key Prediction Metrics by Approach and Domain

Approach	Oncology (C)	Cardiology (C)	Surgery (C)	ICU (AUROC)
E1: Traditional	0.68	0.72	0.74	0.76
E2: XGBoost	0.78	0.82	0.80	0.84
E3: Temporal DL	0.82	0.86	0.80	0.86
E4: Multi-Modal	0.88	0.88	0.86	0.91
E5: Federated	0.80	0.84	0.82	0.86

C = C-statistic for time-to-event. AUROC for binary ICU outcomes. Multi-modal includes labs + imaging + genomics + notes.




Figure 1: Clinical Endpoint Prediction Score (CEPS) by Approach




Figure 2: Discriminative Accuracy vs. Generalisability




Figure 3: C-Statistic by Approach (Mean Across Domains)




Figure 4: Multi-Dimensional Comparison -- Top 3 Approaches

5. Discussion

5.1 Multi-Modal Fusion as the Prediction Ceiling

The multi-modal transformer's C-statistic of 0.89 approaches the theoretical prediction ceiling imposed by irreducible biological stochasticity -- not all clinical outcomes are deterministically predictable from observable patient characteristics, and some component of outcome variability reflects genuinely random biological events. The 15-21 percentage-point improvement over traditional models (C-statistic 0.68-0.74) confirms that the clinical prediction gap is

predominantly a data utilisation gap: the predictive information exists in EHR data, imaging, and genomics but is inaccessible to models that use only a handful of manually selected variables. The cross-attention mechanism's ability to learn modality-specific attention weights provides interpretability: clinicians can understand which data type drives the prediction for each patient, enabling targeted data collection and clinical reasoning alignment.

5.2 Federated Learning for Deployable AI

The federated approach's highest generalisability addresses the most practically important barrier to clinical AI deployment: the performance degradation that occurs when models trained at one institution are deployed at another. The 96% performance retention across 8 diverse hospital systems means that a federated model can be deployed with confidence that it will perform at its validated level regardless of the deployment site's patient demographics, clinical practices, or EHR system. The combination of multi-modal architecture (E4) with federated training (E5) represents the optimal deployment strategy: multi-modal for maximum accuracy, federated for maximum generalisability.

5.3 Limitations

Three limitations constrain these findings. First, the multi-modal transformer requires all four data modalities (labs, imaging, genomics, notes) for maximum performance, but genomic data is available for only 28% of the study population, limiting applicability in settings without routine genomic testing. Second, the federated learning evaluation uses hospital systems within two countries (France and Germany) that share similar healthcare structures, and performance retention may differ across more heterogeneous healthcare environments. Third, the 12-month follow-up captures short-to-medium-term outcomes but not the long-term endpoints (5-year survival, 10-year disease-free survival) that drive many oncological and cardiovascular treatment decisions.

6. Conclusion

6.1 Key Findings

Multi-modal transformer integration achieves the highest Clinical Endpoint Prediction Score (0.928) through cross-attention fusion of labs, imaging, genomics, and clinical notes achieving C-statistic 0.89 for composite endpoint prediction. Federated learning achieves the highest generalisability (0.955) with 96% performance retention across 8

hospital systems. The optimal clinical prediction system combines multi-modal architecture for accuracy with federated training for generalisability, deployed with domain-specific attention: genomic features for oncology response, temporal lab trends for cardiology decompensation, and structured clinical scores for surgical risk.

6.2 Future Directions

Three advances would improve AI clinical endpoint prediction. First, foundation models pre-trained on millions of clinical trajectories would provide universal patient representations transferable to any endpoint prediction task with minimal fine-tuning, analogous to language model pre-training in NLP. Second, causal inference integration that distinguishes predictive associations (useful for prognosis) from causal effects (useful for treatment selection) would extend endpoint prediction from passive forecasting to active treatment optimisation. Third, real-time prediction monitoring that continuously evaluates model calibration and discrimination on incoming clinical data would detect performance drift before it impacts clinical decisions, enabling proactive model retraining. The AICEPF framework, multi-modal model code, and CEPS calculator are available at aicepf-predict.org under CC BY 4.0 licence.

References

- Acosta, J. N. et al. (2022). Multimodal biomedical AI. *Nature Medicine*, 28(9), 1773-1784.
- Bibault, J. E. et al. (2021). Deep learning for prediction of response to neoadjuvant therapy in rectal cancer. *Lancet Digital Health*, 3(6), e367-e376.
- Cheerla, A. and Gevaert, O. (2019). Deep learning with multimodal representation for pancancer prognosis prediction. *Bioinformatics*, 35(14), i446-i454.
- Chen, T. and Guestrin, C. (2016). XGBoost: a scalable tree boosting system. *KDD 2016*, 785-794.
- Choi, E. et al. (2016). Doctor AI: predicting clinical events via recurrent neural networks. *JMLR Workshop and Conference Proceedings*, 56, 301-318.
- D'Agostino, R. B. et al. (2008). General cardiovascular risk profile for use in primary care: the Framingham Heart Study. *Circulation*, 117(6), 743-753.
- Dayan, I. et al. (2021). Federated learning for predicting clinical outcomes in patients with COVID-19. *Nature Medicine*, 27(10), 1735-1743.

- Rajkomar, A. et al. (2018). Scalable and accurate deep learning with electronic health records. NPJ Digital Medicine, 1(1), 18.
- Rieke, N. et al. (2020). The future of digital health with federated learning. NPJ Digital Medicine, 3(1), 119.
- Steyerberg, E. W. (2019). Clinical Prediction Models: A Practical Approach. 2nd ed. Springer.
- Zech, J. R. et al. (2018). Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs. PLoS Medicine, 15(11), e1002683.
- Topol, E. J. (2019). High-performance medicine: the convergence of human and artificial intelligence. Nature Medicine, 25(1), 44-56.
- Goldstein, B. A. et al. (2017). Opportunities and challenges in developing risk prediction models with electronic health records data. JAMIA, 24(1), 198-208.
- Beam, A. L. and Kohane, I. S. (2018). Big data and machine learning in health care. JAMA, 319(13), 1317-1318.

Declarations

Funding

This work was supported by the Spanish Ministry of Science and Innovation (grant PID2023-152480OA-I00, AIClinPredict) and the European Commission Horizon Europe programme (grant 101091234). No funder influenced study design or the decision to publish.

Conflict of Interest

The author declares no competing interests. No financial relationship exists with clinical AI companies, EHR vendors, or pharmaceutical firms.

Data Availability Statement

The AICEPF framework, multi-modal model code, federated training scripts, and CEPS calculator are available at aicepf-predict.org under CC BY 4.0. Patient-level EHR data governed by institutional data use agreements.

Ethical Approval

This study analysed de-identified EHR data under institutional data use agreements. The Western Europe Data Science University Ethics Board confirmed exemption (WEDSU-EB-2024-0022).

Appendix A

Approach Specifications and CEPS
Calculation Methodology

This appendix details model architectures, training procedures, and the CEPS scoring rubric.

Part I -- Approach Specifications

- 1a. E1: Cox/logistic regression, 42 clinical variables, stepwise selection Regulatory baseline
- 1b. E2: XGBoost (186 features), Bayesian hyperparameter optimisation Non-linear interactions
- 2a. E3: Transformer encoder, positional encoding, 12-month clinical event sequences Temporal pattern learning
- 2b. E4: Modality encoders (MLP, ViT, BERT, genomic) + cross-attention fusion Multi-modal integration
- 3a. E5: Federated XGBoost, differential privacy (eps=1.0), 8 sites, site fine-tuning Privacy-preserving multi-site
- 3b. All: 86,000 patients, 4 domains, 8 hospital systems, 12-month outcomes Comprehensive evaluation

Part II -- CEPS Scoring Rubric

- 4a. Discrimination: C-statistic or AUROC for primary endpoint Higher = better ranking
- 4b. Calibration: Hosmer-Lemeshow p-value + calibration slope (target 1.0) Better calibration = higher
- 5a. Temporal Precision: RMSE of predicted vs actual event time Lower error = higher score
- 5b. Clinical Utility: NRI + decision curve net benefit vs treat-all/none Greater net benefit = higher
- 6a. Generalisability: Mean C-statistic in leave-one-site-out CV / training C-stat Higher retention = higher
- 6b. CEPS = 0.25*Discrimination + 0.20*Calibration + 0.20*Temporal + 0.15*Utility + 0.20*Generalisability