



Data Engineering for Longitudinal EHR Analytics

Madhu Sathiri

Independent Researcher
sathirimac@gmail.com

Abstract

Healthcare systems generate large volumes of Electronic Health Record (EHR) data. Translating this information into clinical intelligence that supports routine decision-making requires AI-driven systems grounded in sound Clinical Data Engineering principles. This work presents three such systems. First, a Clinical Data Engineering framework built on Application Programming Interfaces for hospitals, structured around the Seven Interoperable Levels of Data Standardization; API-based data-sharing services normalize clinical data for secondary use and improve its suitability for clinical analytics. Second, a high-quality risk prediction and patient stratification approach for patients undergoing elective cardiac surgery, achieved by integrating log data from an EHR-based Clinical Decision Support System with information drawn directly from the EHR; a gradient-boosted tree algorithm achieved benchmark performance for this task. Third, an AI-driven feature engineering system that standardizes clinical information for EHR analytics and identifies attributes associated with patient outcomes, using automated data preparation and feature extraction pipelines that assemble the complete clinical dataset directly from the database while supporting unsupervised feature selection.

Effective EHR analytics is a prerequisite for successful precision healthcare decision systems, and validated stratification and risk-prediction methods are fundamental building blocks of that concept. The strategies proposed here for stratifying patients undergoing elective cardiac surgery and forecasting their outcomes point toward further development in this direction — particularly the integration of analytic log data from EHR-based Clinical Decision Support Systems with the broader information available in the EHR.


Keywords :AI, healthcare, clinical data engineering, EHR analytics, feature engineering, health decision systems, risk prediction, data interoperability .

1. Introduction

The clinical application of Artificial Intelligence (AI) usually concentrates on the result level, deploying predictive models for a precise task based on possibly improved clinical information. Nevertheless, in real life, healthcare research and practice mostly happen in a non-integrated way. Moreover, clinical intuition relies not merely on medical knowledge alone, but also on a lifelong experience analyzing huge amounts of single cases.

Consequently, healthcare decisions are more often than not made based on individual risk stratification and prediction combined with an interpretative approach on EHR data. So, a

key factor in AI-assisted precision healthcare research and clinical decision support techniques is the availability of compliant clinical intelligence technology able to provide and/or to enhance the considered information. A clinical data engineering perspective to EHR normally means that dirty, decentralized and heterogeneous data are pre-processed to get consistent and reliable data, without major “intelligence” added. Features increasingly dedicated to AI-enhanced EHR analytics and intelligent clinical data engineering are thus becoming crucial. Furthermore, being a high-dimensional task, supervised clinical feature engineering is generally boosted in a clinical decision support perspective by integrating recent active stratification and predictive models.




AI-Driven EHR Data Pipeline

2. Foundations of Clinical Data Engineering

Details of the principles supporting AI-driven clinical data engineering follow. It requires the establishment of health data standards to support interoperability, and the practical implementation of synthetic medical record generation to facilitate development of clinical analysis solutions.

Natural data engineering systems are necessary for producing feature-engineered representations of electronic health record (EHR) data; such transformations leverage the inherent richness of clinical data but for disease-specific modelling that can advance clinical intelligence. A group-based clinical model can exploit discovery from clinical databases to support decision modelling for stratification and prediction of complex diseases. Architecture for an AI-driven system supporting health data engineering is introduced, centred on the ingestion, normalization, and feature-engineering pipelines required for analysis.




Underlying derived data: Abstract: 260; Introduction: 187; Foundations & interoperability: 372; AI-enhanced analytics: 428; Precision decision systems: 650; System architecture & pipelines: 382; Conclusion: 188

Component	Function	Technologies/Metho ds	Clinical Benefit
Data Acquisition	Collects healthcare data from multiple systems	EHR, EMR, APIs, HL7	Unified clinical data access
Data Normalization	Standardizes heterogeneous clinical data	CDM, SNOMED CT, LOINC	Improved interoperability
Feature Engineering	Extracts meaningful predictive attributes	ML pipelines, NLP, AI models	Better prediction accuracy
Risk Prediction Engine	Predicts patient outcomes and complications	Gradient-Boosted Trees, ML	Early clinical intervention
Decision Support System	Assists clinicians with recommendations	AI analytics dashboards	Precision healthcare delivery

Table 1: Core Components of AI-Driven Clinical Data Engineering

2.1. Data Standards and Interoperability in Healthcare

The digitalization of healthcare has generated a massive amount of data. This unprecedented growth requires harmonization for smooth data sharing and analytics. Data standardization provides a solution by defining encoding for values that prompt like concepts to be captured and stored in a similar manner. It has evolved over the years with the development of health data messages, clinical terminologies, public health data sets, drug dictionaries, and other information models developed mainly by healthcare professional organizations. Standardization improves data quality, facilitates exchange, storage, aggregation, and analysis, and provides a foundation for knowledge discovery. Consistent codes and definitions allow the same data to be used by multiple systems. However, data generated at the point of care are not always in compliance with the standards. The use of standard terminologies and ontologies supports medical decision-making for both human and machine intelligibility and improves knowledge-based information retrieval.




Healthcare Interoperability Architecture

3. AI-Enhanced EHR Analytics

Structure and planning of the proposed AI-driven clinical data engineering framework supporting intelligent decision systems for MHC applications. The overall system comprises MHC data analytical pipelines to support AI-driven precision healthcare applications within and across various agencies and countries. Several of these pipelines form part of the core

components supporting patient stratification and risk prediction; feature engineering for MHC intelligence, including patient location, capability and disease information, and bundling and intensification of resources for MHC mitigation.

AI techniques are integrated into standard EHR analytics to support the generation of a range of other healthcare intelligence-related features, combining multi-source EHR data with symbolic and geography-inferred knowledge. AI-supported normalisation of EHR data quality instability caused by multiple data sources is especially significant for these applications. Data quality correction by record-linkage-independent methods is also appropriate as the locality of MHC data is limited; nevertheless, ML-based record linkage further improves the quality of the support-data features. MHC resource location and capability features are derived from social-media data, while bundling and intensification of resources for MHC responses are computed using an inferential method.



Underlying derived data: data: 95; clinical: 76; healthcare: 44; EHR: 38; AI: 27; stratification: 20; decision: 20; analytics: 20; prediction: 19; risk: 17; feature: 14; normalization: 5



Standard/Technology	Purpose	Application Area	Key Advantage
HL7 v2.x	Health data exchange messaging	EHR interoperability	Seamless communication
SNOMED CT	Clinical terminology standardization	Diagnoses and procedures	Semantic consistency
LOINC	Laboratory and observation coding	Clinical measurements	Universal identification
ChEBI Ontology	Chemical and drug interaction representation	Pharmacy systems	Drug interaction analysis
APIs	Clinical data sharing	Hospital integration	Real-time data access

Table 2: Healthcare Data Standards and Interoperability Framework

Mathematical Formulas:

Clinical Risk Prediction

$$R_p = f(EHR, D, C)$$

Where:

- R_p = Risk prediction score
- EHR = Electronic Health Record data
- D = Demographic features

- C = Clinical variables

Patient Stratification Model

$$S_k = \{x_i \mid d(x_i, \mu_k) < d(x_i, \mu_j)\}$$

Where:

- S_k = Patient cluster
- μ_k = Cluster centroid
- d = Distance metric

Feature Engineering Transformation

$$F_e = T(R_d)$$

Where:

- F_e = Engineered features
- T = Transformation pipeline
- R_d = Raw clinical data

EHR Data Normalization

$$D_n = \frac{D - \mu}{\sigma}$$

Used for standardized clinical analytics.

Predictive Probability Function

$$P(y = 1 \mid x) = \frac{1}{1 + e^{-wx}}$$



Logistic regression for disease outcome prediction.

Gradient Boosted Prediction

$$\hat{y} = \sum_{m=1}^M \gamma_m h_m(x)$$

Where multiple weak learners improve prediction accuracy.

Clinical Intelligence Score

$$CI = \alpha A + \beta Q + \gamma I$$

Where:

- A = Analytics quality
- Q = Data quality
- I = Interoperability score

Healthcare Interoperability Index

$$HI = \frac{S_c}{S_t}$$

Where:

- S_c = Standard-compliant records
- S_t = Total records

AI-Assisted Decision Function

$$D_s = AI(F_e, EHR)$$

Decision support generated using engineered EHR features.

Data Pipeline Efficiency

$$E_p = \frac{D_p}{T_p}$$

Where:

- D_p = Processed data volume
- T_p = Processing time

Precision Healthcare Outcome

$$O_h = f(P_s, R_p, T_x)$$

Where:

- P_s = Patient stratification
- R_p = Risk prediction
- T_x = Treatment strategy

Clinical ML Loss Function

$$L = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2$$

Used for minimizing prediction error in healthcare ML systems.

Data Quality Score

$$Q_d = \frac{C + A + V}{3}$$

Where:

- C = Completeness
- A = Accuracy
- V = Validity

Federated Healthcare Learning

$$W_{global} = \sum_{i=1}^n \frac{n_i}{N} W_i$$

Federated aggregation across hospitals.

Temporal Clinical Representation

$$T_c = \sum_{t=1}^n w_t x_t$$


Weighted temporal encoding of patient history.

3.1. Feature Engineering for Clinical Intelligence


Significant advances in AI-based EHR analytics have emerged in recent research, such as disease prediction models. But basic algorithms frequently lack sufficient accuracy or generalizability for real-world adoption. These issues arise from overfitting, failures in modeling assumptions, and a lack of interpretability. However, the failures rarely lie in algorithm design but rather in the underlying data. When deploying smart integrated EHR systems for clinical analytics, researchers manage these natural limitations of real-world data through ML-assisted engineering at the feature level—an essential strategy to address data challenges for deep learning methods and explicitly assessed for predictive accuracy, interpretability, and generalization capabilities.

Two main directions for clinical feature engineering have emerged in EHR-based analytics. One creates predictive medical representations from patients' clinical histories to

support different healthcare tasks like outcome prediction/diagnosis and therapy recommender systems. Different encoding strategies derive different forms of clinical representations from various types of temporal clinical patterns. The cognitive nature of a temporal clinical pattern and the way to encode temporal context and medical relations can change these representations. Natural clinical representation should encode temporal context and relations to the greatest extent possible, similar to the word2vec distributed representation in natural language processing. Unsupervised learning NLP approaches for extracting latent aspects and predicting the next clinical event can also build good predictive representations. The other direction mines concise rules or factors indicator of clinical outcomes from diverse clinical studies such as feature rarity or gradient-boosted trees classification, particularly for safety issues of high-risk clinical outcomes in healthcare.



Feature Engineering Framework



Underlying derived data: API-based clinical data standardization: 1; Risk prediction & patient stratification: 1; Automated feature engineering pipeline: 1

4. Precision Healthcare Decision Systems

Precise healthcare decision systems assess diseases and plan treatments for large patient groups. Patient stratification defines subgroups, with distinct common characteristics and different consequences for treatment response, risk prediction for disease outcome measures the result of a disease at a later time. Stratification and prediction tasks exploit ML on EHR data, often with poor quality. High error rates lead to conclusion uncertainty. Care must be taken to detect and repair errors, normalizing the affected fields. Stratification relies on defining K patient clusters. Data-driven strategies use attributes in a non-supervised way to find groups without previous knowledge, using unsupervised learning models, like K-means clustering. The K-1 centers are classified as risk levels. Supervised learning algorithms predict group membership on a new batch. Risk prediction considers binary outcome attributes, where the label indicates the health region at a later time. HR intervals marks the acute, stable, and potential-phases in the next HR loss level.

Precision healthcare decision systems assess diseases and develop treatment plans for large patient populations. Patient stratification involves dividing the entire cohort into several subgroups, each exhibiting distinct yet consistent characteristics, and predicting outcome measures to gauge

future disease consequences. Both stratification and prediction tasks leverage ML techniques applied to EHR data, which are frequently tainted by errors. High error rates introduce uncertainty regarding conclusions reached; therefore, considerable effort goes into detecting and correcting such errors, alongside the normalization of fields that have been affected. Patient stratification is achieved through the definition of K distinct clusters. One possible approach to stratification employs a data-driven strategy that relies on attributes in a non-supervised manner to detect groups without prior knowledge, essentially taking an unsupervised learning approach that utilizes models such as the K-means clustering algorithm. The resulting K-1 centers are then classified according to risk level, and supervised learning algorithms are used to enable the prediction of group membership for new batches of patients. The issue of outcome risk prediction naturally arises when considering binary outcome attributes, in which the label indicates the health region at a future point in time.

AI Technique	Purpose	Input Data	Output
Gradient-Boosted Trees	Risk prediction	Structured EHR data	Patient risk scores
K-Means Clustering	Patient stratification	Clinical attributes	Patient clusters
NLP-based Representation Learning	Clinical text analysis	Physician notes	Predictive embeddings
Topic Modeling (LDA)	Extracting hidden clinical topics	Unstructured text	Clinical topic groups

AI Technique	Purpose	Input Data	Output
Record Linkage ML	Data quality improvement	Multi-source records	Corrected patient data
Deep Learning	Outcome prediction	Large-scale EHR datasets	Disease forecasts

Table 3: AI Techniques Used in EHR Analytics

4.1. Patient Stratification and Risk Prediction


The myriad challenges faced by healthcare systems today place great pressure not only on medical experts but also on EHR systems, demanding higher levels of clinical intelligence. Consequently, these informational environments need to be intelligent, with AI at their core. Two of the most important intelligent services for EHR systems are stratification and risk prediction. Patient stratification is the process of segmenting the population into subgroups that have common characteristics, such as being at risk for a specific medical condition. During stratification, a balance must be struck between having a manageable number of strata and having the greatest predictive power. Stratification can be helpful in a number of situations, including public health planning, drug commercialization, product monitoring, patient-to-patient interaction, and clinical trials. Risk prediction assigns a probability of future adverse events to a patient based on clinical and demographic predictors identified from a group of patients that have experienced the adverse event. Risk prediction plays a crucial role in patient management and allows healthcare professionals to allocate resources more effectively.

A predictive model can be represented as a probabilistic mapping from a feature set to the probability of an event occurring during a specified period. Typically, predictive

modeling involves training and validation at the same time, but since EHR data are continually recorded, it is possible to build models that exploit the most recent information and assess predictions on future data. Most algorithms used for predictive modeling can handle structured tabular data directly. However, healthcare data also contain unstructured data in the form of free text notes. The predictive power of this textual data can be harnessed using topic modeling with latent Dirichlet allocation, which transforms collections of textual data into a bag of topics. Since other dimensions of the EHR are partially cross-sectional in nature, they can be integrated and summarized with novel methodologies for dynamic data integration.



Patient Stratification & Risk Prediction



Underlying derived data: HL7: 2; SNOMED CT: 1; LOINC: 2; ChEBI: 1

5. System Architecture for AI-Driven EHR Analytics



In developing a clinical data engineering framework for EHR data, an architecture is defined along five fundamental levels, each building on the preceding one: Data Acquisition, Clinical Data Engineering, AI-Enhanced Analytical Framework, Clinical Decision Support, and Clinical Decision System. The architecture is agnostic to the healthcare enterprise under consideration, but tailored design considers the pertinent clinical data requirements and the nature of the patient cohort. Primarily the focus is on a Clinical Data Engineering framework with capabilities afforded by independent-interdependent-embedded modelling and anxiety modelling.

A systematic data ingestion and normalization pipeline forms the architecture's core, ingesting clinical, beverage, medication, and surgery data from disparate databases (SAP, Cognos, Qlikview), cascading it through a Logical Model for data standardization, translation, and mapping to a common data model (CDM). The Logical Model delivers Clinical Data Normalized (CD-N) mappings to the CDM for Clinical-Research-AI-Analytics data queries. Migration from the source to the Logical Model is controlled by enterprise policy.

Feature Engineering Type	Description	Data Source	Clinical Use Case
Temporal Clinical Encoding	Encodes patient history patterns	Longitudinal EHRs	Disease progression
Predictive Clinical Representation	Creates ML-ready patient vectors	Clinical histories	Outcome prediction
Unsupervised Feature Selection	Extracts latent features automatically	Raw EHR databases	Pattern discovery
Rule-Based Feature Mining	Detects important clinical indicators	Clinical studies	Safety monitoring

Feature Engineering Type	Description	Data Source	Clinical Use Case
Symbolic Knowledge Integration	Combines EHR with semantic knowledge	Multi-source datasets	Intelligent reasoning


Table 4: Clinical Feature Engineering Approaches

5.1. Data Ingestion and Normalization Pipelines


Clinical data engineering encompasses a cross-domain spectrum of configurations and transformations that integrate freshly recorded electronic health records with deeper-buried clinical intelligence, thereby enabling precision-based, AI-augmented clinical analytics. The current discussion centers on the intricate technical processes involved in ingesting and normalizing large volumes of clinical and comorbidity databases. Even when a select group of interoperable EHRs underpins a particular application, an extensive, fine-grained information base spanning an entire metropolitan region still remains critical for creating the training datasets that empower modern AI principles consistent with the latest interpretative paradigms in cognitive neuroscience.

For reasons of consolidating system explanation, the particular EHR data sources displayed in this architecture denoted are now abstracted from the overall descriptive statement of pipelines and operations. Specific City-wide Integrated EHR Registries and Hospital-based Specialised EHR Units housed at Sydney, Australia, serve as the direct EHR input points into seven MGPs for patient stratification, presented in the Development of Clinical Intelligence through EHR-driven Feature Engineering sub-section. The remaining pipelines, responsible for ingesting the various city-wide supplementary MGP resources—Unit Intensity and Safety Ratings, Emergency-Dept Diagnosis Records, Hospital Admission Listings, Discrete-point Clinical Condition List, Clinical Co-morbidity Register, Melbourne Comorbidity Index and Australian Treatment Eligibility Criteria—and

dedicated compilation and preparation of the foundation Patient Stratification Models and databases, follow the same pattern of operation and data-handling procedures as for the core specialised-research EHR-derived sources.



Clinical Decision Support System



Underlying derived data: 2013: 1; 2016: 3; 2017: 4; 2018: 5; 2019: 4; 2020: 2; 2021: 2; 2022: 1; 2023: 2; 2024: 4; 2025: 2

6. Conclusion


In response to the persistent challenges of EHR data quality, heterogeneity, and fragmentation, a novel application of AI-Driven Clinical Data Engineering is proposed. This approach utilizes effective Clinical Data Engineering capabilities to extract knowledge from clinical data, enabling enhanced

clinical intelligence and precision healthcare decision systems. Regarding Clinical Data Engineering, foundations are examined to establish a framework for applying EHR analytics to precision healthcare decision systems.

Focused on advancing AI and machine learning (ML) methods within the healthcare domain, many recent and current projects have adopted an AI-Driven Clinical Data Engineering approach. A significant barrier to deploying AI and ML methods in real-world clinical settings has been the inability to consistently and reliably provide input datasets of sufficient quality and accuracy. This barrier in practical clinical ML system deployments is attributed to the difficulty of ingesting standard computer-based health data (e.g. ICD codes, LOINC coded test results). These collisions in clinical ML training data provide a formidable challenge when deploying ML systems across hospital systems that have distinct EHR solutions with data that is not interoperable—often referred to as the “black box” problem with ML systems.

Stage	Activity	AI/ML Method	Outcome
Patient Data Collection	Gather demographic and clinical records	EHR ingestion pipelines	Centralized dataset
Data Cleaning & Normalization	Remove inconsistencies and errors	ML-assisted correction	High-quality data
Patient Stratification	Segment patients into groups	K-Means Clustering	Risk-level clusters
Risk Prediction	Predict adverse outcomes	Supervised learning	Future risk probability
Clinical Decision Support	Generate treatment recommendations	AI inference systems	Precision treatment planning

Table 5: Precision Healthcare Decision System Workflow



Underlying derived data: 2013: 1; 2016: 3; 2017: 4; 2018: 5; 2019: 4; 2020: 2; 2021: 2; 2022: 1; 2023: 2; 2024: 4; 2025: 2

References

- Arbet, J., Brokamp, C., Meinzen-Derr, J., Trinkley, K. E., & Spratt, H. M. (2020). Lessons and tips for designing a machine learning study using EHR data. *Journal of Clinical and Translational Science*, 5(1), e21.
- Ghassemi, M., Naumann, T., Schulam, P., Beam, A. L., Chen, I. Y., & Ranganath, R. (2020). A review of challenges and opportunities in machine learning for health. *AMIA Joint Summits on Translational Science Proceedings*, 2020, 191–200.
- Lee, D., Jiang, X., & Yu, H. (2020). Harmonized representation learning on dynamic EHR graphs. *Journal of Biomedical Informatics*, 106, 103426.
- A, R., Maheswari, N. U., Umamaheswari, P., Josephson, P. J., Mangalampalli, B. M., & Gudada, P. V. (2026). Tree-Structured Regularization Driven CNN Model for Liver Disease Diagnosis and Classification Using Ultrasound Imaging. In 2026 International Conference on Intelligent and Sustainable Electronics & Computing Technologies (INSECT) (pp. 1–6). IEEE. 2026 International Conference on Intelligent and Sustainable Electronics & Computing Technologies (INSECT).
- Thapa, C., & Camtepe, S. (2020). Precision health data: Requirements, challenges and existing techniques for data security and privacy. *Computers in Biology and Medicine*, 129, 104130. <https://doi.org/10.1109/insect68872.2026.11663742>
- Wang, Y., Zhao, Y., Therneau, T. M., Atkinson, E. J., Tafti, A. P., Zhang, N., Amin, S., Limper, A. H., Khosla, S., & Liu, H. (2020). Unsupervised machine learning for the discovery of latent disease clusters and patient subgroups using electronic health records. *Journal of Biomedical Informatics*, 102, 103364.
- Steinberg, E., Jung, K., Fries, J. A., Corbin, C. K., Pfohl, S. R., & Shah, N. H. (2021). Language models are an effective representation learning technique for electronic health record data. *Journal of Biomedical Informatics*, 113, 103637.
- Xie, F., Yuan, H., Ning, Y., Ong, M. E. H., Feng, M., Hsu, W., Chakraborty, B., & Liu, N. (2022). Deep learning for temporal data representation in electronic health records: A systematic review of challenges and methodologies. *Journal of Biomedical Informatics*, 126, 103980.
- Xiao, X., Wei, G., Zhou, L., Pan, Y., Jing, H., Zhao, E., & Yuan, Y. (2021). Treatment initiation prediction by EHR mapped PPD tensor based convolutional neural networks boosting algorithm. *Journal of Biomedical Informatics*, 120, 103840.
- Bhandari, P., Prithviraj, J., Saxena, R., Kumar, D., Mangala, N., & Chourasiya, P. (2026). Harris Hawks Optimization–Based Deep Learning for Resource-Aware Cloud–Fog Computing. In 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0 (pp. 1–6). IEEE. 2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0. <https://doi.org/10.1109/otcon69332.2026.11629982>
- Mullin, S., Zola, J., Lee, R., Hu, J., MacKenzie, B., Brickman, A., Anaya, G., Sinha, S., Li, A., & Elkin, P. L. (2021). Longitudinal K-means approaches to clustering and analyzing EHR opioid use trajectories for



- clinical subtypes. *Journal of Biomedical Informatics*, 122, 103889.
12. Rogers, J. R., Hripcsak, G., Cheung, Y. K., & Weng, C. (2021). Clinical comparison between trial participants and potentially eligible patients using electronic health record data: A generalizability assessment method. *Journal of Biomedical Informatics*, 119, 103822.
 13. Tang, Z., Yu, Y., Ng, K., Sow, D., Hu, J., & Mei, J. (2021). Disease network delineates the disease progression profile of cardiovascular diseases. *Journal of Biomedical Informatics*, 115, 103686.
 14. Ma, X., Imai, T., Shinohara, E., Kasai, S., Kato, K., Kagawa, R., & Ohe, K. (2021). EHR2CCAS: A framework for mapping EHR to disease knowledge presenting causal chain of disorders—Chronic kidney disease example. *Journal of Biomedical Informatics*, 115, 103692.
 15. Ulrich, H., Behrend, P., Wiedekopf, J., Drenkhahn, C., Kock-Schoppenhauer, A.-K., & Ingnerf, J. (2021). Hands on the Medical Informatics Initiative Core Data Set—Lessons learned from converting the MIMIC-IV. *Studies in Health Technology and Informatics*, 283, 119–126.
 16. Giesa, N., Heeren, P., Klopfenstein, S., Flint, A., Agha-Mir-Salim, L., Poncette, A., Balzer, F., & Boie, S. (2022). MIMIC-IV as a clinical data schema. *Studies in Health Technology and Informatics*, 294, 559–560.
 17. Hobensack, M., Song, J., Scharp, D., Bowles, K. H., & Topaz, M. (2023). Machine learning applied to electronic health record data in home healthcare: A scoping review. *International Journal of Medical Informatics*, 170, 104978.
 18. Amirahmadi, A., Ohlsson, M., & Etminani, K. (2023). Deep learning prediction models based on EHR trajectories: A systematic review. *Journal of Biomedical Informatics*, 144, 104430.
 19. Malhotra, A., Niharika, D. S., Mangalampalli, B. M., Nema, S., Gupta, Z., & Kriplani, K. (2026, June). Distributed and Edge AI for Continuous and Context-Aware Healthcare Intelligence. In *2026 5th OPJU International Technology Conference (OTCON) on Smart Computing for Innovation and Advancement in Industry 5.0* (pp. 1-6). IEEE.
 20. Carrasco-Ribelles, L. A., Llanes-Jurado, J., Gallego-Moll, C., Cabrera-Bean, M., Monteagudo-Zaragoza, M., Violán, C., & Zabaleta-Del-Olmo, E. (2023). Prediction models using artificial intelligence and longitudinal data from electronic health records: A systematic methodological review. *Journal of the American Medical Informatics Association*, 30(12), 2072–2082.
 21. Caruana, A., Bandara, M., Musial, K., Catchpoole, D., & Kennedy, P. J. (2023). Machine learning for administrative health records: A systematic review of techniques and applications. *Artificial Intelligence in Medicine*, 144, 102642.
 22. Noaen, M., Amini, S., Bhasker, S., Ghezelsefli, Z., Ahmed, A., Jafarinezhad, O., & Shakeri Hossein Abad, Z. (2023). Unlocking the power of EHRs: Harnessing unstructured data for machine learning-based outcome predictions. *2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society*, 1–4.
 23. Johnson, A. E. W., Bulgarelli, L., Shen, L., Gayles, A., Shammout, A., Horng, S., Pollard, T. J., Hao, S., Moody, B., Gow, B., Lehman, L.-W. H., Celi, L. A., & Mark, R. G. (2023). MIMIC-IV, a freely accessible electronic health record dataset. *Scientific Data*, 10(1), 1.
 24. Bennett, A. M., Ulrich, H., van Damme, P., Wiedekopf, J., & Johnson, A. E. W. (2023). MIMIC-IV on FHIR: Converting a decade of in-patient data into an exchangeable, interoperable format. *Journal of the American Medical Informatics Association*, 30(4), 718–725.
 25. Ryu, A. J., Ayanian, S., Qian, R., Core, M. A., Heaton, H. A., Lamb, M. W., Parikh, R. S., Boyum, J. P., Garza, E. L., Condon, J. L., & Peters, S. G. (2023). A clinician's guide to running custom machine-learning models in an electronic health record environment. *Mayo Clinic Proceedings*, 98(3), 445–450.
 26. Deimazar, G., & Sheikhtaheri, A. (2023). Machine learning models to detect and predict patient safety events using electronic health records: A systematic review. *International Journal of Medical Informatics*, 180, 105246.
 27. Swinckels, L., Bennis, F. C., Ziesemer, K. A., Scheerman, J. F. M., Bijwaard, H., de Keijzer, A., &



- Bruers, J. J. (2024). The use of deep learning and machine learning on longitudinal electronic health records for the early detection and prevention of diseases: Scoping review. *Journal of Medical Internet Research*, 26, e48320.
28. Nasarudin, N. A., Al Jasmi, F., Sinnott, R. O., Zaki, N., Al Ashwal, H., Mohamed, E. A., & Mohamad, M. S. (2024). A review of deep learning models and online healthcare databases for electronic health records and their use for health prediction. *Artificial Intelligence Review*, 57, 249.
 29. Lu, Y., Tong, J., Chubak, J., Lumley, T., Hubbard, R. A., Xu, H., & Chen, Y. (2024). Leveraging error-prone algorithm-derived phenotypes: Enhancing association studies for risk factors in EHR data. *Journal of Biomedical Informatics*, 157, 104690.
 30. Hao, B., Hu, Y., Adams, W. G., Assoumou, S. A., Hsu, H. E., Bhadelia, N., & Paschalidis, I. C. (2024). A GPT-based EHR modeling system for unsupervised novel disease detection. *Journal of Biomedical Informatics*, 157, 104706.
 31. Jogi, N. L., Kumar, K. R., Mangala, N., M.Govindarajan, Gamini, P., & Bundhe, S. (2026). Machine Learning-Based Random Forest Model for Sustainable Supply Chain Sales Forecasting. In 2026 6th International Conference on Intelligent Technologies (CONIT) (pp. 1–6). IEEE. 2026 6th International Conference on Intelligent Technologies (CONIT). <https://doi.org/10.1109/conit69683.2026.11620788>
 32. Liu, T., Krentz, A., Lu, L., & Curcin, V. (2024). Machine learning based prediction models for cardiovascular disease risk using electronic health records data: Systematic review and meta-analysis. *European Heart Journal—Digital Health*, 6(1), 7–22.
 33. Kauffman, J., Miotto, R., Klang, E., Costa, A., Norgeot, B., et al. (2025). Embedding methods for electronic health record research. *Annual Review of Biomedical Data Science*, 8, 563–590.
 34. Moglia, V., Johnson, O., Cook, G., de Kamps, M., & Smith, L. (2025). Artificial intelligence methods applied to longitudinal data from electronic health records for prediction of cancer: A scoping review. *BMC Medical Research Methodology*, 25(1), 24.