

ORIGINAL RESEARCH

Open access

Graph-Based Machine Learning Model for Predicting Care Coordination Failures Using Referral Networks, Follow-Up Completion Status, Patient Message Logs, Specialty Access Delays, and Provider Communication Patterns

Ali Hassan^{1*}, Noor Siddiqui¹, Bilal Khan², Sana Malik¹

Abstract

Care coordination failures include missed referrals, lost follow-ups, fragmented communication, and incomplete transitions between primary and specialty care. These failures can delay diagnosis, weaken continuity, and increase avoidable utilisation. Existing detection approaches often depend on manual review, retrospective audits, or simple rule-based flags. Such approaches are poorly suited to capture the relational complexity of patient, provider, referral, messaging, and encounter networks. This article develops a conceptual graph-based machine learning model for predicting care coordination failures. The model represents patient–provider referral networks enriched with follow-up status, patient message activity, specialty access delays, and provider communication patterns. The proposed approach uses a heterogeneous graph neural network in which patients and providers are nodes. Referral, encounter, and messaging relationships are represented as edges, while node and edge features encode follow-up adherence, message frequency, wait-time signals, and communication context. Conceptually, the model would identify high-risk referral edges that combine delayed access, incomplete follow-up, weak messaging activity, or limited provider communication. These predictions would support coordinator review before a referral becomes a documented care gap. A graph-based model could shift care coordination from reactive tracking toward predictive prevention. By identifying fragile referral relationships early, it could support more timely outreach and safer continuity of care.

Keywords Electronic health records, Referral networks, Graph neural networks, Care coordination, Patient portals, Specialty access

*Correspondence:

Ali Hassan

ali.hassan@gmail.com

¹ Department of Health Information Systems and Analytics, Faculty of Medicine, Aga Khan University, Karachi, Pakistan

² Department of Digital Health Engineering, Faculty of Engineering, Qatar University, Doha, Qatar

Introduction

Care coordination failures occur when patients move across clinical settings without reliable completion of referrals, follow-up visits, consult documentation, or provider-to-provider communication. These breakdowns are especially consequential in primary-to-specialty care, where referral leakage and missed follow-up may delay diagnosis or treatment and increase fragmented utilisation [1, 2]. Patient-sharing network studies show that care delivery is distributed across many clinicians and organisations, making coordination failures difficult to observe from any single clinical workstation [3].

Fragmented readmissions and outpatient care fragmentation further illustrate how discontinuity can shape discharge destination, acute care use, and longitudinal care quality [4, 5].

Current coordination workflows often rely on manual tracking, spreadsheet-based registries, staff memory, and retrospective audit reports. Although electronic health records contain referral orders, appointment status, consult notes, message logs, and encounter histories, these data are commonly stored as disconnected events rather than analysed as a dynamic coordination process [6, 7]. Patient portal and secure message studies show that communication traces can be classified and modelled, suggesting that routine EHR communication data could support automated detection of care gaps [8, 9]. However, rule-based alerts remain limited because they usually evaluate single events rather than relational patterns across patients, clinicians, specialties, and time.

Care coordination is inherently relational because a patient's outcome depends not only on individual risk but also on the structure and timing of interactions among referring clinicians, specialists, care teams, and patients. Referral networks and patient-sharing networks naturally form graphs in which nodes represent patients or providers and edges represent referrals, shared encounters, documentation, or communication [1, 3]. Graph neural networks are well matched to this structure because they can learn from both attributes and topology, including neighbourhood patterns that may indicate weak handoffs or poorly connected care teams [10, 11]. Recent graph-based EHR models and clinical recommender systems demonstrate that heterogeneous clinical graphs can encode complex relationships for prediction and decision support [12, 13].

This MDL article proposes a graph-based machine learning model that could predict care coordination failures by learning from referral network topology and enriched coordination features. The model would integrate follow-up completion status, patient message activity, specialty access delay, provider communication patterns, and shared encounter relationships as dynamic graph signals. Rather than producing only patient-level risk, the proposed model focuses on edge-level referral risk so that coordinators can identify the specific relationship or referral instance most likely to fail. This approach would support targeted, proactive intervention while preserving a conceptual, model-oriented framing rather than reporting experimental results.

Background

Care coordination failures: definitions and consequences

Care coordination failures include incomplete referrals, missed follow-up appointments, lack of consult closure, poor information transfer, and unmanaged transitions after discharge. These failures can emerge when a referral is ordered but not scheduled, when a specialist visit occurs without feedback to the referring clinician, or when a patient disengages during a long wait period [2, 14]. Studies of care fragmentation and readmission pathways show that discontinuity is not merely administrative; it can affect acute utilisation, discharge destination, and the ability of teams to maintain shared situational awareness [4, 5]. Predictive modelling of these failures therefore requires attention to both clinical events and the relational pathways through which care is organised.

Referral networks and provider communication in EHRs

EHR systems capture many elements of the referral lifecycle, including referral orders, appointment scheduling, visit completion, consult notes, and closed-loop communication. Referral automation and closed-loop referral implementation studies show that these data can support structured tracking, but also reveal adaptation challenges when workflows cross organisational and specialty boundaries [6, 7]. Provider communication can also be inferred from shared encounters, co-documentation, concurrent EHR use, and communication patterns among care team members [15, 16]. These traces provide a basis for representing coordination as a provider–patient network rather than as isolated transactions.

Patient engagement signals: portals and messages

Patient portal messages, secure texts, reminder responses, and related communication logs can serve as indicators of engagement and coordination activity. Prior work has shown that patient portal messages can be classified using rule-based, machine learning, and convolutional neural approaches, demonstrating that message content and metadata contain clinically meaningful signals [8, 9]. Qualitative and topic-modelling studies further suggest that secure messaging reflects chronic care coordination, patient concerns, administrative needs, and ongoing navigation of care processes [17, 18]. In a graph model, message frequency, recency, directionality, and unresolved message threads could enrich patient-provider edges with dynamic engagement information.

Specialty access delays as a driver of failure

Specialty access delay is a plausible driver of coordination failure because longer referral-to-visit intervals can increase the chance that patients disengage, symptoms change, insurance status shifts, or communication becomes stale. Referral-loop studies and access-oriented specialty care tools show that completion depends on more than the referral order itself; scheduling, specialist capacity, and feedback loops shape whether the referral becomes clinically effective [2, 6]. Network-based analyses of physician accessibility also indicate that structural position within the specialist workforce can influence access and coordination burden [19]. A predictive graph model should therefore encode access delay as an edge-level feature rather than treating it as a background operational variable.

Graph neural networks for healthcare application

Graph neural networks extend machine learning to relational data by learning representations from nodes, edges, attributes, and neighbourhood structure. In healthcare, graph convolutional and graph representation models have been used for disease prediction, EHR structure learning, medication recommendation, clinical recommendation, and dynamic representation learning [10-12, 20]. Reviews of graph machine learning and medical knowledge graph construction show growing interest in heterogeneous clinical graphs that integrate diagnoses, medications, encounters, providers, and temporal relationships [21, 22]. These methods provide a foundation for modelling care coordination as a dynamic, multi-relational system.

Model Development Overview

High-level predictive pipeline

The proposed pipeline would begin with EHR data extraction, followed by construction of a heterogeneous coordination graph, computation of node and edge features, graph neural network inference, and generation of edge-level referral risk scores. Referral orders, encounter records, patient portal messages, provider identifiers, and scheduling timestamps would be transformed into relational features rather than treated as independent tabular observations [6, 8, 16]. The model would then estimate whether a referral edge is likely to become a coordination failure, while patient-level risk could be summarised from a patient's active outgoing referrals. This design follows the broader movement from static EHR prediction toward relational and graph-based clinical representation learning [10, 12, 23].

Figure 1 illustrates the hierarchical pipeline through which heterogeneous EHR data are transformed into a graph-based representation and used to generate edge-level referral risk predictions for proactive care coordination.




Figure 1. Hierarchical Architecture of a Graph-Based Machine Learning Pipeline for Predicting Care Coordination Failures

Core input data and graph conceptualisation

The graph would include patient nodes, primary care provider nodes, specialist nodes, and possibly clinic or department nodes when organisational structure is relevant. Directed referral edges would connect referring clinicians to receiving specialists through the patient context, while shared encounter and messaging edges would represent communication and engagement patterns [1, 15]. Patient attributes could include follow-up status, prior missed appointments, portal activity, and care gap history, while provider attributes could include specialty, location, and coordination responsiveness [14, 18]. This heterogeneous formulation would allow the model to distinguish between clinical risk, access burden, communication activity, and structural network position.

Design principles

The model should be relationally explicit, clinically interpretable, scalable across a health-system referral network, and capable of handling delayed or missing referral outcomes. Because care coordination failures unfold over time, the model should avoid treating unresolved referrals as simple negatives and should instead distinguish pending, completed, lapsed, and censored states [2, 7]. It should support edge-level explanations that identify why a specific referral appears fragile, drawing on interpretable graph features such as access delay, weak messaging, and prior incomplete referrals [13, 24]. Finally, the model should be designed for coordinator action rather than retrospective reporting, so predictions are useful only when they can guide outreach, scheduling assistance, or provider communication.

Table 1 presents a conceptual decomposition of the major signal domains that jointly contribute to edge-level coordination failure risk within the proposed graph framework.

Table 1. Conceptual Differentiation of Signal Domains Contributing to Edge-Level Care Coordination Failure Risk

Signal Domain	Graph Representation Level	Core Features	Mechanistic Role in Failure Risk	Interaction with Other Domains	Implication for Model Learning
Follow-Up Adherence	Edge (referral-specific)	Completion status, missed visits, elapsed time	Direct indicator of coordination breakdown or delay	Amplified by access delay and low engagement	High predictive weight for imminent failure
Patient Engagement (Messaging)	Edge (patient–provider)	Message frequency, recency, unresolved threads, directionality	Reflects patient activation and responsiveness during referral process	Moderates effect of access delay; low engagement increases vulnerability	Enables temporal sensitivity and dynamic risk updates
Specialty Access Delay	Edge (referral-specific)	Referral-to-visit interval, scheduling lag, specialty capacity proxies	Structural driver of disengagement and referral attrition	Interacts with engagement and provider communication	Captures system-level constraints beyond patient behavior
Provider Communication Patterns	Node + Edge (provider–provider, provider–patient)	Shared encounters, co-documentation, communication density	Indicates strength of coordination pathways and information flow	Mitigates risk from delays and incomplete follow-up	Enhances relational interpretability via network topology
Historical Referral Patterns	Edge + Node (provider-level aggregation)	Prior completion rates, referral volume, collaboration history	Encodes learned coordination reliability between providers	Influences baseline risk independent of patient factors	Supports generalization across similar referral pathways
Network Topology	Global + Local graph structure	Connectivity, centrality, clustering, neighborhood density	Reveals structural fragility or robustness of coordination pathways	Integrates all domains into relational context	Enables GNN to learn non-linear interaction effects

Graph Construction from Care Coordination Data

Node identification and feature extraction

Patient nodes would be defined from the EHR master patient index, while provider nodes would be derived from referring clinician identifiers, receiving specialist identifiers, and care team membership. Patient features would include demographics, comorbidity burden, prior no-show or follow-up history, portal communication patterns, and previous care gap indicators [14, 17]. Provider features would include specialty, practice location, referral volume, responsiveness to messages, and shared-care connectivity within patient-sharing networks [1, 3, 16]. These features would support learning from both individual risk profiles and the structural context in which referrals are placed.

Referral edges and weighting

Directed referral edges would be constructed at the time a referral order is placed, linking a referring provider to a specialist through a patient-specific referral instance. Edge attributes would include referral type, order timestamp, target specialty, priority, scheduling status, completion state, and elapsed time since order [2, 6]. Referral completion could be represented conceptually as completed, pending, lapsed, cancelled, or unresolved, avoiding premature labels for active referrals [7]. Edge weights could reflect urgency, specialty access delay, historical completion patterns, and the strength of prior collaboration between the referring and receiving clinicians [19, 25].

Supplementary edges: messaging and encounters

Supplementary edges would enrich the referral graph with communication and engagement signals that are not captured by the referral order alone. Patient–provider message edges could be weighted by frequency, recency, directionality, and unresolved content patterns, while encounter edges could represent shared visits, co-management, or repeated contact over time [8, 9, 18]. Provider–provider communication edges could be inferred from shared patients, co-documentation, concurrent EHR work, or documented communication during treatment planning [15, 16]. These additional relationships would help the graph neural network distinguish an isolated referral from one embedded in an active coordination pathway.

Graph Neural Network Architecture

Model selection and rationale

A heterogeneous graph transformer or relational graph convolutional network would be appropriate because the coordination graph contains multiple node types and multiple edge types. These architectures can represent patients, providers, referrals, encounters, and messages without collapsing them into a single homogeneous network [10, 12]. Attention mechanisms would be especially useful because they could learn whether referral history, access delay, message activity, or provider connectivity contributes most to a specific predicted failure [11, 13]. This model family aligns with recent healthcare graph learning studies that use graph structure to improve clinical prediction and recommendation tasks [20, 22].

Input representation and message passing

The model would project patient, provider, and edge features into a shared latent space while preserving type-specific transformations for referrals, encounters, and messages. During message passing, each node would aggregate information from its neighbourhood, allowing patient embeddings to reflect provider connectivity and allowing provider embeddings to reflect shared coordination histories [10, 12]. Referral-edge embeddings would combine the referring provider, receiving specialist, patient context, access delay, follow-up status, and communication features before producing an edge-level failure logit [13, 24]. This structure would allow the model to reason over both local referral details and broader network signals.

Output: edge-level coordination failure probability

The primary model output would be a conceptual probability that a specific referral edge will result in a coordination failure, such as an uncompleted specialist visit, absent closed-loop documentation, or delayed follow-up. A patient-level coordination risk score could be aggregated from all active outgoing referrals, while provider- or clinic-level summaries could identify recurring coordination bottlenecks [4, 25]. The output should be interpreted as a prioritisation signal rather than a definitive judgment, because unresolved referrals may still be completed with appropriate support [2, 7]. In clinical use, the most important contribution would be identifying which referral relationship requires attention and which graph features explain the concern.

Incorporating Temporal and Dynamic Features

Temporal edge features and ageing

Referral edges should include temporal features that represent elapsed time since referral order, expected specialty access delay, and ageing of unresolved coordination tasks. Dynamic graph methods for EHR data suggest that clinical relationships should be updated as new events occur rather than frozen at the time of initial prediction [12, 26]. Specialty access and workforce network studies further support treating delay as a structural property of the care system, not merely as a patient-level attribute [19]. In this model, temporal ageing would help distinguish a new referral awaiting routine scheduling from an older referral becoming increasingly vulnerable to failure.

Dynamic update of follow-up and message activity

The graph would be updated as follow-up visits are completed, patient messages are sent, referral statuses change, and provider communication occurs. Deep learning approaches using EHR data have shown that longitudinal clinical records can support scalable predictive modelling when repeated events are transformed into structured learning signals [23, 27]. Patient portal message studies also indicate that communication activity can be meaningfully classified and incorporated into coordination models [8, 9, 18]. A nightly or otherwise periodic re-evaluation process could refresh risk estimates as the coordination state evolves.

Handling censored and in-progress referrals

Many referrals are unresolved at prediction time, so the model should not treat all pending referrals as failures or successes. A masked or survival-aware learning objective could conceptually separate completed referrals, lapsed referrals, cancelled referrals, and referrals still within a reasonable completion window [2, 7]. This approach would be especially important for specialties with long wait times, where incomplete status may reflect access constraints rather than patient disengagement [6, 19]. Handling censoring carefully would reduce misleading labels and make the model more appropriate for real-world care coordination workflows.

Model Interpretability for Care Coordinators

Edge-level explanations via attention and XAI

Interpretability should focus on explaining why a specific referral edge appears at risk, rather than only reporting a global patient score. Attention weights in heterogeneous graph models could indicate whether the prediction was shaped by access delay, weak patient messaging, prior missed follow-up, provider network position, or limited communication between clinicians [10, 11, 13]. Explainable graph methods and interpretable risk models could further identify influential node and edge features in clinical contexts where coordinators need transparent reasons for action [24]. Such explanations should remain clinically modest, presenting model-derived signals as support for review rather than as causal claims.

Translating predictions into actionable worklists

Predictions should be translated into plain-language worklist entries that tell coordinators which referral needs attention and why. For example, a flagged referral could indicate that the specialty wait is prolonged, the patient has not responded to portal messages, and the referring and receiving providers have limited prior coordination history [15-17]. Referral automation and closed-loop referral studies show that workflow integration is central to whether digital coordination tools become useful in practice [6, 7]. The model should therefore produce explanations that support outreach, rescheduling, specialist office contact, or escalation to the referring clinician.

Clinical Deployment and Proactive Intervention

Integration with care coordination dashboards

The model would feed into existing care coordination dashboards, population health platforms, or referral management work queues. High-risk referrals could be surfaced as daily review items, with explanation fields summarising access delay, follow-up status, messaging activity, and provider communication context [6, 7]. Network-based studies of provider collaboration and patient sharing suggest that dashboard views could also aggregate risks by clinic, specialty, or referral pathway to reveal recurrent bottlenecks [1, 16, 25]. The deployment goal would be to make graph-derived risk actionable within routine coordinator workflows.

Triggering proactive coordinated care actions

A high-risk referral signal could trigger patient navigator outreach, specialist office re-contact, appointment rescheduling, message follow-up, or escalation to the referring clinician. Because fragmented care and readmission pathways are associated with poorer continuity, proactive intervention should aim to strengthen handoffs before failures become visible only through retrospective review [4, 5]. The model could also support institutional monitoring of specialties with persistent access delays or referral leakage patterns [19, 25]. These actions would require governance to ensure that predictions augment coordinator judgment rather than replace clinical responsibility.

Evaluation Strategy

Predictive performance metrics

Evaluation should compare the graph model with non-graph baselines, such as logistic regression or gradient-boosted models using patient and referral attributes alone. Appropriate conceptual metrics would include discrimination, precision-recall behaviour for rare coordination failures, calibration, and usefulness of edge-level prioritisation, without assuming any specific performance values [11, 23, 28]. Because graph models learn from relational structure, evaluation should test whether referral, message, encounter, and provider-network edges add value beyond tabular features [10, 12, 13]. Interpretability should also be assessed by whether coordinators understand and trust the reasons attached to high-risk referrals.

Table 2 analytically contrasts the proposed graph-based approach with traditional modelling strategies, highlighting the conceptual advantages of relational learning for coordination failure prediction.

Table 2. Analytical Comparison of Graph-Based Versus Traditional Approaches for Care Coordination Failure Prediction

Dimension	Traditional Rule-Based / Tabular Models	Graph-Based Model (Proposed)	Conceptual Advantage of Graph Approach
Unit of Prediction	Patient-level or event-level	Edge-level (referral-specific)	Enables precise identification of failing relationships rather than aggregate risk
Data Representation	Independent records (tabular features)	Heterogeneous graph (nodes + edges + topology)	Captures relational dependencies across patients, providers, and interactions
Handling of Communication	Often excluded or simplified	Explicit messaging and provider communication edges	Incorporates dynamic coordination signals absent in tabular models

Temporal Dynamics	Limited (static snapshots)	Dynamic graph updates with temporal edge features	Reflects evolving coordination states and referral ageing
Interpretability	Feature importance (global or local)	Edge-level explanations via attention and relational context	Provides clinically meaningful reasons tied to specific coordination pathways
Sensitivity to Network Effects	Minimal	High (neighborhood aggregation, topology-aware learning)	Detects structural fragility such as weak provider connectivity
Handling of Incomplete Referrals	Often binary classification	Multi-state modelling (pending, completed, lapsed, censored)	Reduces misclassification of ongoing coordination processes
Scalability Across Systems	Moderate	High with graph representation learning	Generalizes across specialties and provider networks
Clinical Actionability	Limited (alerts often generic)	High (targeted referral-level worklists with explanations)	Directly supports coordinator workflows and interventions

Temporal validation and network generalizability

Temporal validation should train the model conceptually on earlier referral periods and evaluate it on later periods to reflect prospective deployment. Network generalizability should be examined across specialties, clinics, provider groups, and patient-sharing communities because referral behaviour may differ substantially across organisational settings [3, 25, 29]. Dynamic and multimodal graph modelling studies highlight the importance of validating whether learned relationships remain stable as care patterns and EHR use change over time [12, 26, 28]. This strategy would reduce the risk that the model merely memorises historical referral habits within one narrow network.

Prospective impact assessment

A prospective pilot could assess whether model-generated worklists improve coordination processes, coordinator usability, and timely referral completion compared with usual tracking workflows. Such an assessment should examine whether alerts identify actionable referrals, whether coordinators can interpret the explanations, and whether the tool changes communication between patients, primary care teams, and specialists [6, 15, 17]. Studies of referral loops, care fragmentation, and closed-loop referral implementation suggest that impact depends on both prediction quality and the surrounding workflow design [2, 4, 7]. The evaluation should therefore include qualitative feedback and safety monitoring alongside predictive assessment.

Limitations

Graph sparsity and cold-start problems

Graph sparsity may limit prediction for patients with few prior encounters, new providers, or specialists with limited referral history. In these cases, the model may have insufficient neighbourhood information and should present lower-confidence estimates rather than overstate risk [10, 21]. Cold-start problems may be partly mitigated by using patient attributes, specialty type, clinic location, and health-system-level access features when local graph context is sparse [19, 22]. Even so, sparse graph regions would remain a significant deployment challenge.

Data availability and missingness

The model depends on accurate capture of referrals, appointments, message logs, provider identifiers, and specialty access intervals. Messaging may occur outside the EHR, consult notes may be delayed, and referral completion may be difficult to infer when patients receive external specialty care [6-8]. Missingness may also vary by patient portal use, language, digital access, and clinic workflow, which could bias graph features if not carefully monitored [17, 18]. Data quality assessment would therefore be necessary before the model could be responsibly evaluated or deployed.

Conclusion

A graph-based machine learning model for care coordination failure prediction would represent referrals, encounters, messages, patients, and providers as an interconnected clinical network. This structure would allow the model to focus on fragile referral relationships rather than treating care gaps as isolated patient-level events.

The central strength of the proposed approach is its ability to integrate multiple coordination signals into an edge-level risk assessment. Referral topology, follow-up status, patient messaging, specialty access delay, and provider communication patterns would jointly inform which care transitions may require proactive attention.

Important challenges remain, including sparse graph regions, incomplete communication data, uncertain referral outcomes, and the organisational effort required to act on predictive insights. The model would need careful workflow integration so that coordinators receive useful, understandable, and clinically appropriate alerts.

Future work should pursue multi-site validation within accountable care organisations and integrated delivery systems. Shared benchmarks for care coordination graphs would help the field compare modelling strategies, evaluate fairness, and develop safer predictive tools for preventing care gaps.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 10 Apr 2023

Revised: 12 May 2023

Accepted: 08 Jul 2023

Published online: 20 July 2023

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

- Landon BE, Keating NL, Onnela JP, Zaslavsky AM, Christakis NA, O'Malley AJ. Patient-sharing networks of physicians and health care utilization and spending among Medicare beneficiaries. *JAMA Intern Med.* 2018;178(1):66-73.
- Patel MP, Schettini P, O'Leary CP, Bosworth HB, Anderson JB, Shah KP, et al. Closing the referral loop: an analysis of primary care referrals to specialists in a large health system. *J Gen Intern Med.* 2018;33(5):715-21.
- DuGoff EH, Fernandes-Taylor S, Weissman GE, Huntley JH, Pollack CE. A scoping review of patient-sharing network studies using administrative data. *Transl Behav Med.* 2018;8(4):598-625.
- Edwards ST, Greene L, Chaudhary C, Boothroyd D, Kinoshian B, Zulman DM. Outpatient Care Fragmentation and Acute Care Utilization in Veterans Affairs Home-Based Primary Care. *JAMA Netw Open.* 2022;5(9):e2230036.
<https://doi.org/10.1001/jamanetworkopen.2022.30036>.
- Turbow SD, Ali MK, Culler SD, Rask KJ, Perkins MM, Clevenger CK, et al. Association of Fragmented Readmissions and Electronic Information Sharing With Discharge Destination Among Older Adults. *JAMA Netw Open.* 2023;6(5):e2313592.
<https://doi.org/10.1001/jamanetworkopen.2023.13592>.
- Odisho AY, Lui H, Yerramsetty R, Bautista F, Gleason N, Martin E, et al. Design and development of referrals automation, a SMART on FHIR solution to improve patient access to specialty care. *JAMIA Open.* 2020;3(3):405-12.
- Zehner ME, Kirsch JA, Adsit RT, Gorrilla A, Hayden K, Skora A, et al. Electronic health record closed-loop referral ("eReferral") to a state tobacco quitline: a retrospective case study of primary care implementation challenges and adaptations. *Implement Sci Commun.* 2022;3(1):107.
<https://doi.org/10.1186/s43058-022-00357-4>.
- Cronin RM, Fabbri D, Denny JC, Rosenbloom ST, Jackson GP. A comparison of rule-based and machine learning approaches for classifying patient portal messages. *Int J Med Inform.* 2017;105:110-20.
- Suliman L, Gilmore D, French C, Cronin RM, Jackson GP, Russell M, et al. Classifying patient portal messages using Convolutional Neural Networks. *J Biomed Inform.* 2017;74:59-70.
<https://doi.org/10.1016/j.jbi.2017.08.014>.
- Choi E, Xu Z, Li Y, Dusenberry M, Flores G, Xue Y, et al. Learning the graphical structure of electronic health records with graph convolutional transformer. *Proc AAAI Conf Artif Intell.* 2020;34(1):606-13.
- Sun Z, Yin H, Chen H, Chen T, Cui L, Yang F. Disease Prediction via Graph Neural Networks. *IEEE J Biomed Health Inform.* 2021;25(3):818-26.
<https://doi.org/10.1109/JBHI.2020.3004143>.
- Lee D, Jiang X, Yu H. Harmonized representation learning on dynamic EHR graphs. *J Biomed Inform.* 2020;106:103426.
- Fouladvand S, Gomez FR, Nilforoshan H, Schwede M, Noshad M, Jee O, et al. Graph-based clinical recommender: Predicting specialists procedure orders using graph representation learning. *J Biomed Inform.* 2023;143:104407.
<https://doi.org/10.1016/j.jbi.2023.104407>.

Dantas LF, Fleck JL, Oliveira FL, Hamacher S. No-shows in appointment scheduling: a systematic literature review. *Health Policy*. 2018;122(4):412-21.

Steitz BD, Unertl KM, Levy MA. Characterizing communication patterns among members of the clinical care team to deliver breast cancer treatment. *J Am Med Inform Assoc*. 2020;27(2):236-43.

Li P, Chen B, Rhodes E, Slagle J, Alrifai MW, France D, et al. Measuring Collaboration Through Concurrent Electronic Health Record Usage: Network Analysis Study. *JMIR Med Inform*. 2021;9(9):e28998.
<https://doi.org/10.2196/28998>.

Hefner JL, MacEwan SR, Biltz A, Sieck CJ. Patient portal messaging for care coordination: a qualitative study of perspectives of experienced users with chronic conditions. *BMC Fam Pract*. 2019;20(1):57.

De A, Huang M, Feng T, Yue X, Yao L. Analyzing patient secure messages using a Fast Healthcare Interoperability Resources (FHIR)-based data model: development and topic modeling study. *J Med Internet Res*. 2021;23(7):e26770.

Moen EL, Brooks GA, O'Malley AJ, Schaefer A, Carlos HA, Onega T. Use of a Novel Network-Based Linchpin Score to Characterize Accessibility to the Oncology Physician Workforce in the United States. *JAMA Netw Open*. 2022;5(12):e2245995.
<https://doi.org/10.1001/jamanetworkopen.2022.45995>.

Mao C, Yao L, Luo Y. MedGCN: medication recommendation and lab test imputation via graph convolutional networks. *J Biomed Inform*. 2022;127:104000.

Murali L, Gopakumar G, Viswanathan DM, Nedungadi P. Towards electronic health record-based medical knowledge graph construction, completion, and applications: a literature study. *J Biomed Inform*. 2023;143:104403.

Lu H, Uddin S. Disease prediction using graph machine learning based on electronic health data: a review of approaches and trends. *Healthcare (Basel)*. 2023;11(7):1031.

Ayala Solares JR, Diletta Raimondi FE, Zhu Y, Rahimian F, Canoy D, Tran J, et al. Deep learning for electronic health records: A comparative review of multiple deep neural architectures. *J Biomed Inform*. 2020;101:103337.
<https://doi.org/10.1016/j.jbi.2019.103337>.

Dong X, Wong R, Lyu W, Abell-Hart K, Deng J, Liu Y, et al. An integrated LSTM-HeteroRGNN model for interpretable opioid overdose risk prediction. *Artif Intell Med*. 2023;135:102439.
<https://doi.org/10.1016/j.artmed.2022.102439>.

Flemming R, Sundmacher L. Organization and quality of care in patient-sharing networks. *Health Policy*. 2023;136:104891.

Zhu W, Razavian N. Variationally regularized graph-based representation learning for electronic health records. In: *Proceedings of the Conference on Health, Inference, and Learning*; 2021. p. 1-13.

Rajkomar A, Oren E, Chen K, Dai AM, Hajaj N, Hardt M, et al. Scalable and accurate deep learning with electronic health records. *NPJ Digit Med*. 2018;1:18.
<https://doi.org/10.1038/s41746-018-0029-1>.

Tang S, Tariq A, Dunnmon JA, Sharma U, Elugunti P, Rubin DL, et al. Predicting 30-Day All-Cause Hospital Readmission Using Multimodal Spatiotemporal Graph Neural Networks. *IEEE J Biomed Health Inform*. 2023;27(4):2071-82.
<https://doi.org/10.1109/JBHI.2023.3236888>.

Landon BE, Onnela JP, Meneades L, O'Malley AJ, Keating NL. Assessment of racial disparities in primary care physician specialty referrals. *JAMA Netw Open*. 2021;4(1):e2029238.