

ORIGINAL RESEARCH

Open access

Graph Neural Network for Predicting Fragmented Care Episodes Using Provider Networks, Encounter Sequences, Referral Patterns, Cross-Setting Utilization, and Unresolved Care Gap Indicators

Patrick O'Connor^{1*}, Sean Murphy¹

Abstract

Care fragmentation is a silent driver of adverse outcomes because it distributes responsibility for a patient's care across multiple providers, settings, and transitions. It is often recognized only after duplication, missed follow-up, or contradictory management has already occurred. Fragmentation unfolds through networks of providers, referrals, encounters, and unresolved care tasks. Conventional prediction models often treat these events as independent patient-level variables rather than as relational patterns embedded in a care system. This article develops a conceptual graph neural network model for predicting the likelihood of fragmented care episodes. The proposed model uses provider networks, encounter sequences, referral patterns, cross-setting utilization, and unresolved care gap indicators as an integrated graph representation. A heterogeneous graph is constructed with patient and provider nodes, while edges represent encounters, referrals, shared patients, cross-setting transitions, and unresolved care gaps. A graph neural network learns patient embeddings from this relational structure and produces a fragmentation risk score for each patient-period. Conceptually, the model would identify patients at elevated risk of fragmentation, such as those moving across many disconnected providers, receiving repeated referrals, or accumulating open follow-up tasks. These signals could support earlier care coordination before avoidable gaps widen into adverse outcomes. The relational lens of graph neural networks could transform fragmentation prediction from a static risk scoring task into a network-aware early-warning system. Such a model would be expected to help health systems detect fragmented care episodes before they become operationally and clinically consequential.

Keywords Graph neural networks, Care coordination, Provider networks, Care fragmentation, Referral patterns, Healthcare utilization

*Correspondence:

Patrick O'Connor
patrick.oconnor@gmail.com

¹ Department of Health Informatics and AI Systems, Faculty of Medicine, Trinity College Dublin, Dublin, Ireland

Introduction

Care fragmentation refers to the dispersion of a patient's care across multiple clinicians, organizations, and settings without sufficient coordination, continuity, or shared accountability. It is clinically important because fragmented ambulatory care has been associated with hospitalization, emergency department use, medication burden, and broader failures of continuity across complex patient journeys [1-3]. The burden is especially relevant for patients with chronic

illness or high service use, for whom repeated handoffs and disconnected care plans can increase the risk of duplication, missed follow-up, and preventable deterioration [4, 5]. In this context, fragmentation is not merely a descriptive property of care delivery but a measurable operational risk that may precede avoidable utilization and patient harm [6, 7].

Current predictive approaches to utilization and care coordination often rely on patient-level summaries, such as prior visits, comorbidity burden, or recent acute care use, without explicitly representing the relational pathways through which care becomes fragmented. Deep learning models for electronic health records have shown that longitudinal clinical data can support future risk prediction, but many such architectures are not designed to encode provider-sharing networks, referral chains, or cross-setting care transitions as first-class model objects [8, 9]. Process-oriented healthcare analytics further emphasizes that patient journeys are sequential and event-based, yet sequence models alone may underrepresent the network context connecting patients, clinicians, and care venues [10]. As a result, conventional models may identify patients who are high risk in general while failing to explain whether risk arises from disconnected provider networks, incomplete referrals, or unresolved coordination gaps.

Graph neural networks offer an alternative modeling paradigm because they learn from nodes, edges, and neighborhoods rather than only from independent tabular records. In healthcare, graph-based learning has been applied to representation learning, medication recommendation, laboratory imputation, clinical event prediction, and disease risk modelling, demonstrating the relevance of graph structures for complex clinical data [11–14]. Broader graph neural network methods, including graph convolutional networks, attention-based architectures, and inductive neighbor aggregation, provide mechanisms for propagating information across connected entities while preserving relational structure [15, 16]. These properties make GNNs conceptually well matched to care fragmentation, where risk emerges from patterns of shared patients, referral dependencies, disconnected providers, and repeated transitions across care settings.

This article proposes a conceptual GNN model for predicting fragmented care episodes by combining provider networks, encounter sequences, referral patterns, cross-setting utilization, and unresolved care gap indicators into a unified graph representation. The model would treat fragmentation as a dynamic, relational phenomenon rather than as a static count of visits or providers. Patient nodes would be connected to providers, settings, referrals, and gap indicators, allowing message passing to summarize how coordination risk accumulates across the care network. The intended output is a patient-period fragmentation probability that could support care coordination teams, population health programs, and value-based care organizations seeking earlier intervention.

Background

Care fragmentation and its measurement

Care fragmentation has been conceptualized as distinct from continuity and coordination, because it captures the extent to which care is spread across multiple providers and settings rather than the quality of coordination alone [1]. Empirical studies have measured fragmentation using ambulatory dispersion, provider counts, care concentration indices, and related claims- or EHR-derived indicators, linking these measures to hospitalization, emergency department use, medication burden, and adverse device-related outcomes [2, 3, 5, 17]. These measures are valuable because they convert an otherwise diffuse care delivery problem into observable features that can be tracked over time. However, many indices remain patient-centric summaries and do not fully encode whether the providers involved are connected, whether referrals close successfully, or whether unresolved care gaps persist across settings [6, 7].

Provider networks and shared patients

Provider networks are commonly constructed from administrative claims or EHR data by linking clinicians who share patients, participate in referral chains, or jointly contribute to longitudinal episodes of care. Patient-sharing network studies have shown that administrative data can reveal meaningful structures of clinical collaboration, informal care

teams, and organizational boundaries that are not always visible in formal rosters [18, 19]. Network density, provider centrality, and community structure can indicate whether a patient's care occurs within a cohesive local network or across loosely connected providers [20, 21]. For fragmentation prediction, these network properties are important because the same number of providers may imply different risks depending on whether they operate within an integrated care network or across disconnected organizational clusters [22].

Encounter sequences and referral patterns

Fragmentation unfolds through time, so the sequence of encounters, referrals, missed appointments, and setting transitions is central to understanding when care becomes uncoordinated. Closed-loop referral research shows that referral completion requires more than ordering a specialist visit; it also depends on communication, tracking, follow-up, and confirmation that the intended care actually occurred [23, 24]. Patient journey modelling and process mining approaches further demonstrate that healthcare pathways can be represented as ordered event sequences, making it possible to distinguish routine transitions from irregular or potentially unsafe care trajectories [10]. A graph-based model can use referral direction, repeated referrals, referral loops, and temporal spacing between encounters as edge-level signals that may indicate emerging fragmentation.

Cross-setting utilization and care gaps

Cross-setting utilization is a key indicator of fragmentation because patients may move between primary care, specialty care, emergency departments, inpatient units, and home-based care without consistent information transfer. Studies of fragmentation have linked dispersed outpatient care to emergency department visits, hospitalizations, and acute care utilization, suggesting that setting transitions can serve as both drivers and consequences of fragmented care [2, 3, 7]. Unresolved care gaps, such as open referrals, missed follow-ups, incomplete medication reconciliation, or overdue preventive services, add a task-based dimension to fragmentation because they identify care processes that remain unfinished. Integrating these care gap indicators into a graph would allow the model to represent not only where care occurred but also whether necessary coordination actions were completed [4, 17].

Graph neural networks for healthcare

Graph neural networks have become increasingly relevant in healthcare because many clinical problems involve relationships among patients, diagnoses, medications, providers, procedures, and health system entities. Early healthcare GNN applications demonstrated graph-based attention for clinical representation learning and graph-augmented memory for medication recommendation, while later work extended graph convolution to medication recommendation, laboratory imputation, and specialist procedure order prediction [11-13, 25]. Surveys of clinical GNNs and electronic health record knowledge graphs show that graph methods can support risk prediction, phenotyping, recommendation, and clinical knowledge integration when relational structure is available [26, 27]. For care fragmentation analytics, the same principles can be adapted from disease or medication graphs to operational graphs that represent patient-provider relationships, referral pathways, and unresolved coordination tasks.

Model Development Overview

High-level predictive pipeline

The proposed predictive pipeline begins by extracting patient, provider, encounter, referral, utilization, and care gap data from claims, EHR, scheduling, and care management systems. These data are transformed into a dynamic heterogeneous graph in which each patient-period is represented by the patient's evolving neighborhood of providers, settings, referrals, and unresolved tasks. A graph neural network encodes the graph into patient embeddings, and a downstream classifier outputs the conceptual probability that the patient's next care period could become fragmented. This design extends deep learning for healthcare utilization prediction by preserving the relational and temporal context that standard patient-level feature vectors may omit [8, 9].

Core graph elements

The graph contains patient nodes with demographic, morbidity, utilization, and prior care pattern features, provider nodes with specialty, organizational affiliation, location, and panel-level attributes, and setting nodes representing primary care, emergency care, inpatient care, specialty clinics, and ancillary services. Edges capture patient-provider encounters, provider-provider shared-patient relationships, referral orders, referral completions, cross-setting transitions, and unresolved care gap flags. Provider networks derived from shared patients and referrals provide the structural layer, while encounter sequences and care gaps provide the temporal and operational layers [19, 20, 22]. This graph design allows the model to distinguish a patient with many coordinated providers inside one connected network from a patient with a similar provider count spread across disconnected clusters.

Design principles

The model is designed to be relationally explicit, temporally aware, scalable to large health system populations, and interpretable for care coordination teams. Relational explicitness means that patient risk is learned from the arrangement of providers, referrals, and settings rather than only from aggregate counts; temporal awareness means that recent encounters, new referrals, and unresolved gaps can exert greater influence than older events. Scalability is important because fragmentation prediction would be most useful when applied across attributed panels, multi-payer populations, or large care networks [6, 20, 21]. Interpretability is equally important because care coordinators require understandable explanations linking a high-risk score to specific providers, referrals, or unresolved tasks rather than opaque model outputs [25, 26].

Figure 1 illustrates the proposed left-to-right operational workflow for transforming provider networks, encounter sequences, referral pathways, cross-setting utilization, and unresolved care gaps into interpretable graph-based fragmentation risk signals for care coordination action.




Figure 1. Operational Workflow of a Graph Neural Network for Predicting Fragmented Care Episodes Across Provider Networks, Referral Pathways, Care Settings, and Unresolved Care Gaps

Data Sources and Graph Construction

Patient and provider node features

Patient node features would include demographics, chronic disease burden, prior utilization summaries, medication complexity, missed visit history, emergency department use, hospitalization history, and prior fragmentation indicators. Provider node features would include specialty, practice location, organizational affiliation, average panel complexity, observed patient-sharing patterns, and role within referral or shared-care networks. Prior research on care fragmentation

and patient-sharing networks supports the feasibility of extracting these constructs from administrative and EHR-derived data sources [2, 6, 18, 19]. The model should treat these features as contextual signals attached to graph entities rather than as sufficient predictors by themselves, because fragmentation risk depends on how patients and providers are connected over time.

Encounter, referral, and cross-setting edges

Encounter edges connect patients to providers using visit timestamps, encounter types, care settings, and visit sequence information, while referral edges connect referring and receiving providers or link patients to ordered services. Cross-setting transition edges represent movement between outpatient, emergency, inpatient, specialist, home-based, and ancillary care environments, reflecting the operational pathways through which care may become dispersed. Closed-loop referral studies show that referral ordering, completion, and communication status are crucial for distinguishing coordinated transitions from incomplete or fragmented transitions [23, 24]. By encoding these transitions as typed, timestamped graph edges, the model could learn whether repeated ED-to-specialist transitions, unresolved referrals, or disconnected provider handoffs are expected to precede fragmented care episodes [3, 7].

Unresolved care gap indicators as edge attributes

Unresolved care gap indicators can be encoded as edge attributes, temporal flags, or dedicated subgraph elements attached to patients, providers, and care management tasks. Examples include open referrals, missed follow-up visits, overdue screening tasks, incomplete diagnostic workups, unresolved medication reconciliation, or lack of confirmation that a recommended service occurred. This approach aligns with the view that fragmentation is partly a failure of coordination and closure, not simply a high number of visits or clinicians [1, 4, 23]. When care gap indicators are integrated into the graph, the model can represent whether a patient's network is merely complex or whether it contains unfinished tasks that should trigger care coordination attention.

Table 1 summarizes how patient, provider, encounter, referral, care setting, and unresolved care gap data are represented within the proposed heterogeneous graph neural network architecture.

Table 1. Input Structure and Representation Learning Design for a Graph Neural Network Model of Fragmented Care Episodes

Model Component	Manuscript-Specific Data Element	Graph Representation	Representation Learning Role	Practical Care Coordination Relevance
Patient profile	Demographics, chronic disease burden, prior utilization, medication complexity, missed visit history, prior fragmentation indicators	Patient node attributes	Provides clinical and utilization context for each patient-period embedding	Helps distinguish general clinical complexity from relational fragmentation risk
Provider profile	Specialty, practice location, organizational affiliation, panel complexity, network position, referral role	Provider node attributes	Encodes the characteristics of clinicians and care teams surrounding the patient	Helps identify whether care is concentrated within a coherent team or dispersed across disconnected providers

Care setting profile	Primary care, specialty care, emergency department, inpatient unit, home-based care, ancillary services	Setting nodes or setting attributes on encounters	Preserves where care occurs and whether the patient is moving across high-risk transition points	Supports recognition of cross-setting care instability and handoff burden
Encounter history	Visit timestamps, encounter types, provider contacts, setting transitions, sequence of care events	Timestamped patient-provider and patient-setting edges	Allows the model to learn temporal movement through the care system	Identifies rapidly expanding or unstable care patterns before they become visible as adverse outcomes
Referral activity	Referral order, referring provider, receiving provider, service requested, completion status, communication status	Directed referral edges and referral task nodes	Represents intended care transitions and whether they are completed or unresolved	Supports early identification of referral-closure fragmentation
Shared-patient relationships	Providers connected through shared patients, informal care teams, referral relationships, organizational clusters	Provider-provider edges	Allows the model to learn whether providers belong to connected or disconnected care communities	Differentiates multi-provider coordinated care from fragmented multi-provider care
Cross-setting transitions	Movement between outpatient, emergency, inpatient, specialist, home health, and ancillary settings	Directed transition edges or setting-sequence edges	Encodes the operational pathway of care movement over time	Identifies patients moving through settings without stable coordination or ownership
Unresolved care gaps	Open referrals, missed follow-ups, overdue screening, incomplete diagnostic workups, medication reconciliation gaps	Edge attributes, task nodes, or care gap flags	Adds task-level evidence of unfinished coordination work	Converts the model output into actionable care management priorities
Temporal snapshots	Monthly or quarterly patient-period graph states	Sequential graph snapshots	Allows the model to learn how fragmentation risk evolves over time	Supports periodic population health review and proactive outreach
Sparse-network indicators	Limited observed encounters, out-of-network care, incomplete external data, new patient status	Uncertainty features or missingness indicators	Prevents overconfident predictions when graph evidence is incomplete	Alerts users that a low observed fragmentation score may reflect missing data rather than true coordination
Relational message passing	Patient, provider, setting, referral, and care gap neighborhoods	Heterogeneous graph neural network layers	Aggregates information from clinically and operationally meaningful neighbors	Produces a network-aware patient representation rather than a static tabular risk profile

Patient-period embedding	Learned summary of the patient's care network during a defined prediction window	Final patient embedding vector	Combines clinical context, provider relationships, referral pathways, transitions, and care gaps	Forms the basis for a fragmentation risk score that care teams can interpret and act upon
--------------------------	--	--------------------------------	--	---

Graph Neural Network Architecture

Graph construction and heterogeneity handling

The proposed architecture uses a heterogeneous graph with patient, provider, setting, referral, and care gap entities represented as distinct node or edge types. A relational GNN would be appropriate because encounter, referral, shared-patient, transition, and unresolved-gap edges carry different meanings and should not be collapsed into a single undifferentiated connection. Healthcare GNN applications have already shown that heterogeneous clinical relationships can be encoded for medication recommendation, clinical representation learning, and specialist procedure prediction, supporting the conceptual extension to care coordination graphs [11-13, 25]. In this model, relation-specific transformations would allow the network to learn different message-passing rules for referrals, shared encounters, cross-setting transitions, and care gap signals.

Message passing and node embedding

During message passing, patient nodes aggregate information from neighboring providers, settings, referrals, and unresolved care gap indicators, while provider nodes aggregate signals from their shared patients and referral partners. Attention or weighted aggregation mechanisms could allow the model to emphasize the most informative edges, such as recent incomplete referrals, care gaps involving high-risk specialties, or provider clusters that are weakly connected to the patient's primary care network. Graph attention and graph convolution methods provide general mechanisms for learning such neighborhood-based representations, and clinical GNN surveys indicate that these approaches can be adapted to diverse EHR and health system prediction tasks [15, 16, 26, 28]. The resulting patient embedding would encode not only individual clinical risk but also the structure and recent evolution of the patient's care network.

Fragmentation prediction head

The patient embedding would be passed to a fully connected prediction head with a sigmoid output representing the conceptual probability that the next patient-period could become fragmented. The prediction target should be defined using prospective indicators such as dispersed provider involvement, incomplete referral closure, duplicate care processes, unresolved follow-up tasks, or cross-setting utilization patterns consistent with fragmentation. The model may also be extended to predict fragmentation subtype, such as duplication-dominant fragmentation, referral-closure fragmentation, cross-setting transition fragmentation, or medication-coordination fragmentation, depending on the operational definition used by the health system. This prediction head should be evaluated as a decision-support signal rather than as an autonomous clinical determination, consistent with the need for interpretable and workflow-aware clinical prediction systems [25, 26].

Handling Temporal and Dynamic Graph Evolution

Temporal snapshots and graph sequence models

Because care fragmentation develops over time, the graph should be represented as a sequence of patient-centered temporal snapshots rather than as a single static structure. Quarterly or monthly snapshots could capture how provider connections, referral pathways, care settings, and open care gaps expand, contract, or shift before a fragmented episode becomes clinically visible. A temporal GNN or recurrent graph encoder would be expected to summarize these evolving

neighborhoods while preserving the relational meaning of encounters, referrals, and care transitions [8, 10, 28]. This approach would allow the model to distinguish a stable multi-provider care team from a rapidly changing and poorly connected care network.

Edge aging and recency weighting

Temporal edge weighting should reflect the assumption that recent encounters, referrals, care gaps, and cross-setting transitions are more informative for near-future fragmentation than remote events. Older edges could remain in the graph as historical context but contribute less to message passing, while recent incomplete referrals, emergency visits, or unresolved follow-up tasks could receive greater attention. This design is consistent with longitudinal healthcare prediction, where sequential information and recent clinical events often shape future utilization risk [8, 9]. In a care fragmentation model, recency weighting would help the GNN prioritize active coordination risks rather than treating all historical contacts as equally relevant.

Handling new patients and sparse provider networks

New patients, infrequent users, and patients receiving much care outside the observed network may have sparse graph neighborhoods that make direct fragmentation prediction more uncertain. For such cases, the model could use inductive graph learning, provider embeddings, setting-level embeddings, or similarity to patients with comparable utilization trajectories as fallback sources of contextual information. GraphSAGE-style inductive learning and broader GNN methods are relevant because they support embedding generation for previously unseen or weakly connected nodes using neighborhood features rather than fixed node identities [15, 16]. The model should therefore express greater uncertainty for sparse cases and should be evaluated for whether sparse-network predictions remain useful to care coordinators.

Model Interpretability for Care Coordination

Explaining fragmentation risk via subgraph importance

For clinical and operational use, the model must explain which subgraph patterns most contributed to a high fragmentation score. Edge importance, node attribution, attention weights, and post hoc GNN explanation methods could highlight specific providers, referral edges, care settings, or unresolved care gaps that shaped the patient embedding. This is important because care coordinators need to know whether the risk arises from disconnected specialists, repeated emergency transitions, incomplete referral closure, or unresolved medication-related tasks [25, 26]. Interpretable graph explanations would therefore convert the model from a black-box risk score into a practical coordination aid.

From prediction to coordination plan

A high fragmentation probability should be translated into a concrete coordination plan that identifies the most actionable relational problem. For example, the explanation might indicate that risk is driven by several unconnected specialists, repeated referrals without documented closure, or an open medication reconciliation task after an acute care transition. Evidence on referral management and closed-loop communication supports the importance of moving from recognition of a fragmented pathway to specific closure actions [23, 24]. The care coordinator could then prioritize consolidating specialty communication, confirming referral completion, assigning a responsible clinician, or closing unresolved follow-up tasks.

Clinical Deployment and Integration Into Care Management

Embedding in population health platforms

The proposed GNN would be most useful if embedded within population health or care management platforms that already support attributed panels, outreach lists, and care coordinator workflows. The model could run periodically on the health system's current graph and surface patients whose relational patterns suggest increasing fragmentation risk. Prior work on fragmentation across payer populations and high-risk care settings indicates that fragmentation signals can vary by population, setting, and network structure, so deployment should allow local calibration and workflow-specific thresholds [3, 6, 7]. The output should include both a patient-level risk signal and a concise explanation linking risk to modifiable providers, referrals, transitions, or care gaps.

Supporting value-based care and network management

Beyond patient-level outreach, the model could support value-based care by identifying structural fragmentation within provider networks and referral ecosystems. Aggregated explanations could reveal specialist groups associated with repeated leakage, primary care panels with many open referrals, or cross-setting pathways that frequently precede fragmented episodes. Patient-sharing network research suggests that provider network density, community structure, and shared-care relationships can reveal organizational coordination patterns that are not visible from patient-level summaries alone [19-22]. Health systems could use these insights to refine referral pathways, strengthen closed-loop communication, and target network-level care coordination improvements.

Evaluation Strategy

Predictive performance metrics

The model should be evaluated against clinically meaningful fragmentation definitions, including dispersed provider involvement, incomplete referral closure, duplicate care processes, unresolved care gaps, and cross-setting utilization patterns. Predictive evaluation could examine discrimination, calibration, precision-recall behavior, and decision-curve usefulness, but these metrics should be interpreted as conceptual evaluation domains rather than as evidence of actual model performance. Comparators should include non-graph baselines that use only patient-level features, allowing investigators to assess whether graph structure adds value beyond standard utilization summaries [8, 9]. Evaluation should also examine whether GNN explanations identify actionable causes of predicted fragmentation rather than merely reproducing general utilization risk.

Temporal validation and generalizability

Temporal validation should use forward-time testing so that the model is evaluated on future care periods rather than randomly mixed past and future observations. This is important because provider networks, referral patterns, and care delivery structures may change over time, particularly when health systems modify access pathways or care coordination programs. External validation should assess whether a model developed in one care network can generalize to another region, payer mix, or organizational structure, since fragmentation patterns can differ across populations and delivery systems [6, 20, 21]. The evaluation should therefore test not only whether the GNN can predict fragmentation conceptually, but also whether its relational representations remain meaningful across health system contexts.

Operational impact on care coordination

Operational evaluation should focus on whether the model improves care coordination workflow, not only whether it predicts an abstract fragmentation label. A pilot implementation could assess whether care coordinators find the explanations useful, whether flagged patients have clearer coordination plans, and whether the model helps prioritize patients with unresolved referrals, duplicate services, or cross-setting care instability. Fragmentation studies linking dispersed care to acute utilization, hospitalizations, medication burden, and adverse outcomes provide the rationale for evaluating downstream operational consequences [2, 3, 5, 12]. However, any operational impact should be tested prospectively and cautiously, because a conceptual model alone cannot establish reductions in utilization or harm.

Table 2 presents the evaluation, interpretability, deployment safeguard, and practical action requirements needed to translate the proposed graph neural network from conceptual prediction into responsible care coordination support.

Table 2. Evaluation, Interpretability, Deployment Safeguards, and Practical Action Framework for the Proposed Fragmentation Prediction Model

Domain	What Should Be Evaluated	Manuscript-Specific Implementation Question	Operational Risk if Neglected	Practical Action for Health Systems
Outcome definition	Fragmented care label based on dispersed providers, incomplete referral closure, duplicate processes, unresolved gaps, or cross-setting instability	Does the prediction target reflect a clinically meaningful and actionable fragmentation episode?	The model may predict utilization intensity rather than true care fragmentation	Define fragmentation with input from care coordination, primary care, referral management, and quality teams
Baseline comparison	Non-graph models using patient-level utilization and comorbidity summaries	Does the graph structure add value beyond standard patient-level risk factors?	The GNN may add complexity without meaningful operational benefit	Compare against transparent tabular and sequence-based baselines before deployment
Temporal validation	Forward-time testing across future patient-periods	Does the model work when provider networks and referral patterns change over time?	Random validation may overestimate readiness for real-world use	Use prospective or forward-time validation windows
Generalizability	Testing across regions, payer groups, care settings, and organizational structures	Does the model transfer across different delivery networks and referral ecosystems?	A model trained in one system may fail in another with different care pathways	Conduct external validation or local recalibration before broader use
Calibration	Agreement between predicted fragmentation risk and observed future fragmentation	Are high-risk and moderate-risk scores interpretable for prioritizing care coordination workload?	Poor calibration can overload teams with false alerts or miss patients needing intervention	Calibrate thresholds to care coordinator capacity and local intervention pathways
Explanation quality	Important providers, referral edges, care settings, and unresolved gaps contributing to risk	Can users see why a patient was flagged?	Care teams may distrust opaque scores or be unable to act on them	Provide concise subgraph explanations tied to modifiable coordination issues
Fairness and equity	Performance across age groups, chronic disease burden, insurance types, language groups, geography, and access patterns	Does the model systematically under-detect fragmentation in patients with incomplete or out-of-network data?	Existing access inequities may be amplified or hidden	Stratify evaluation and include uncertainty indicators for incomplete data

Data completeness	Availability of external encounters, referral status, provider affiliations, and care gap documentation	Is the observed graph a reliable representation of the patient's actual care network?	Missing edges may be misread as coordinated care	Integrate claims, health information exchange, and care management feeds where possible
Sparse-network uncertainty	New patients, low utilizers, out-of-network users, and patients with limited local records	Does the model communicate uncertainty when relational data are thin?	The model may produce misleadingly confident scores	Display uncertainty flags alongside risk scores
Human oversight	Care coordinator review before outreach or intervention	Is the model positioned as decision support rather than an automated care determination?	Automated labeling may misclassify patients or create inappropriate outreach	Require human review of risk explanations before action
Workflow integration	Placement in population health platforms, referral work queues, or care management dashboards	Does the alert connect to a specific next step?	Risk scores may increase administrative burden without improving coordination	Link each alert to referral closure, follow-up assignment, communication consolidation, or care plan review
Monitoring after deployment	Alert volume, usefulness, drift, false positives, care coordinator feedback, and unintended burden	Does the model remain useful as care pathways and provider networks evolve?	Model performance and user trust may degrade over time	Establish periodic monitoring, governance review, and retraining triggers

Limitations

Graph sparsity and data availability

A major limitation is that the graph may be incomplete when patients receive care outside the available EHR, claims feed, or health information exchange. Missing out-of-network encounters, undocumented referrals, and unavailable care gap data could make the graph appear more coordinated or less fragmented than the patient's actual experience. Fragmentation measurement studies have shown that observed patterns can depend on payer population, data source, and care setting, making data completeness central to valid interpretation [6, 7]. The model should therefore treat sparse graph structure as a source of uncertainty and should avoid assuming that absence of an observed edge means absence of care.

Implementation complexity and interpretability trade-offs

GNNs introduce implementation complexity because they require graph data pipelines, typed edge construction, temporal updating, computational infrastructure, and governance processes for clinical deployment. Their explanations may also be approximate, especially when message passing distributes information across many providers, settings, and care gap indicators. Surveys of GNNs in healthcare and general graph learning emphasize that architecture selection, scalability, interpretability, and validation remain active challenges for real-world applications [15, 16, 26]. For care coordination, the model should therefore be positioned as a decision-support tool that requires human review rather than as an automated determination of fragmented care.

Conclusion

A graph neural network for predicting fragmented care episodes offers a relational way to model how patients move through provider networks, referral chains, care settings, and unresolved care tasks. Instead of treating fragmentation as a static count of providers or encounters, the proposed model represents fragmentation as an evolving network pattern. This makes it possible to generate an early-warning signal before dispersed care becomes clinically or operationally harmful. The model's purpose is to support timely care coordination rather than replace professional judgment.

The key strength of the proposed approach is its ability to combine provider relationships, encounter sequences, referral status, cross-setting utilization, and care gap indicators in one dynamic graph. A patient embedding learned from this graph could capture whether the patient's care is distributed across a coherent network or scattered across poorly connected providers and settings. Interpretable subgraph explanations could then identify the specific relational drivers of the risk signal. This combination of prediction and explanation is especially important for care coordination teams that need actionable reasons, not only risk scores.

Several challenges remain before such a model could be implemented responsibly. Health systems would need sufficiently complete data on referrals, encounters, provider affiliations, care gaps, and out-of-network utilization. The model would also require validation across different care delivery structures, payer populations, and clinical domains. Operational integration would need careful design so that risk alerts support care coordinators without increasing alert fatigue or administrative burden.

Future work should prioritize pilot implementations within accountable care organizations, integrated delivery networks, and value-based care programs seeking to reduce fragmented care. These pilots should evaluate whether graph-based risk explanations help coordinators close referrals, consolidate communication, and identify patients whose care is becoming dispersed across settings. The most important contribution of this model is not the use of a complex algorithm by itself, but the reframing of care fragmentation as a dynamic network problem. That reframing could help health systems move from retrospective measurement toward earlier, more coordinated intervention.

Acknowledgements

None

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 20 Jan 2025

Revised: 11 Mar 2025

Accepted: 09 Apr 2025

Published online: 20 July 2025

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

- Kern LM, Bynum JP, Pincus HA. Care fragmentation, care continuity, and care coordination—how they differ and why it matters. *JAMA Intern Med.* 2024;184(3):236-7.
- Kern LM, Ringel JB, Rajan M, Colantonio LD, Casalino LP, Pinheiro LC, et al. Ambulatory care fragmentation and subsequent hospitalization: evidence from the REGARDS study. *Med Care.* 2021;59(4):334-40.
- 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.
- Joo JY. Fragmented care and chronic illness patient outcomes: a systematic review. *Nurs Open.* 2023;10(6):3460-73.
- Funk RJ, Pagani FD, Hou H, Zhang M, Yang G, Malani PN, et al. Care fragmentation predicts 90-day durable ventricular assist device outcomes. *Am J Manag Care.* 2022;28(12):e444-e450.
- Kern LM, Seirup JK, Rajan M, Jawahar R, Miranda Y, Stuard SS. Extent of health care fragmentation in different payer populations: evidence from the Hudson Valley of New York. *Popul Health Manag.* 2019;22(2):138-43.
- Kern LM, Ringel JB, Rajan M, Casalino LP, Colantonio LD, Pinheiro LC, et al. Ambulatory care fragmentation, emergency department visits, and race: a nationwide cohort study in the US. *J Gen Intern Med.* 2023;38(4):873-80.
- Solares JR, 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.
- Ashfaq A, Sant'Anna A, Lingman M, Nowaczyk S. Readmission prediction using deep learning on electronic health records. *J Biomed Inform.* 2019;97:103256.
- Munoz-Gama J, Martin N, Fernandez-Llatas C, Johnson OA, Sepúlveda M, Helm E, et al. Process mining for healthcare: characteristics and challenges. *J Biomed Inform.* 2022;127:103994.
- Choi E, Bahadori MT, Song L, Stewart WF, Sun J. GRAM: graph-based attention model for healthcare representation learning. In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*; 2017. p. 787-95.
- Shang J, Xiao C, Ma T, Li H, Sun J. GAMENet: graph augmented memory networks for recommending medication combination. *Proc AAAI Conf Artif Intell.* 2019;33(1):1126-33.
- Mao C, Yao L, Luo Y. MedGCN: medication recommendation and lab test imputation via graph convolutional networks. *J Biomed Inform.* 2022;127:104000.
- Lu H, Uddin S. A weighted patient network-based framework for predicting chronic diseases using graph neural networks. *Sci Rep.* 2021;11(1):22607.
- Wu Z, Pan S, Chen F, Long G, Zhang C, Yu PS. A comprehensive survey on graph neural networks. *IEEE Trans Neural Netw Learn Syst.* 2021;32(1):4-24.

Zhou J, Cui G, Hu S, Zhang Z, Yang C, Liu Z, et al. Graph neural networks: a review of methods and applications. *AI Open*. 2020;1:57-81.

Das IG, Ringel JB, Rajan M, Colantonio LD, Safford MM, Kern LM. Fragmented ambulatory care and medication count among older adults. *Am J Med Qual*. 2025;40(3):90-6.

Ong MS, Olson KL, Chadwick L, Liu C, Mandl KD. The impact of provider networks on the co-prescriptions of interacting drugs: a claims-based analysis. *Drug Saf*. 2017;40(3):263-72.

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.

Breslau J, Dana B, Pincus H, Horvitz-Lennon M, Matthews L. Empirically identified networks of healthcare providers for adults with mental illness. *BMC Health Serv Res*. 2021;21(1):777.

Vlaanderen FP, de Man Y, Tanke MA, Munneke M, Atsma F, Meinders MJ, et al. Density of patient-sharing networks: impact on the value of Parkinson care. *Int J Health Policy Manag*. 2022;11(7):1132-41.

Korsberg A, Cornelius SL, Awa F, O'Malley J, Moen EL. A scoping review of multilevel patient-sharing network measures in health services research. *Med Care Res Rev*. 2025;82(3):203-24.

Ramelson H, Nederlof A, Karmiy S, Neri P, Kiernan D, Krishnamurthy R, et al. Closing the loop with an enhanced referral management system. *J Am Med Inform Assoc*. 2018;25(6):715-21.

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.

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.

Boll HO, Amirahmadi A, Ghazani MM, de Morais WO, de Freitas EP, Soliman A, et al. Graph neural networks for clinical risk prediction based on electronic health records: a survey. *J Biomed Inform*. 2024;151:104616.

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.

Tariq A, Kaur G, Su L, Gichoya J, Patel B, Banerjee I. Adaptable graph neural networks design to support generalizability for clinical event prediction. *J Biomed Inform*. 2025;163:104794.