



Detecting Healthcare Billing Fraud with Graph Neural Networks

Emma Richardson
Asst Professor

Abstract

Healthcare payment accuracy directly influences both cost and quality of care, yet high rates of fraud, waste, and abuse continue to persist. These challenges are exacerbated by the inherent complexity and lack of transparency surrounding payment analytics. Addressing these issues, this study proposes an effective, data-driven solution for detecting potentially fraudulent or erroneous billing practices, using unsupervised machine learning to analyze payer-provider relationships. The primary motivation is threefold: first, such unsupervised approaches eliminate the need for a labeled training set; second, they take advantage of the richness and heterogeneity of the data, including subrogation and encounter information; and third, they have no bias toward any specific type of anomaly.

Core methodological contributions include the novel application of unsupervised and graph-based methods for billing analysis and anomaly detection, and the integration of external data into traditional billing tables to uncover new patterns. Testing demonstrates that model-enriched Normalized Payment Amount can help identify fraudulent encounters, offering long-term potential for both real-time monitoring and assisting state-level agencies responsible for managing Medicaid-funded healthcare.

Keywords: Healthcare Payment Integrity, Fraud Waste and Abuse Detection, Unsupervised Machine Learning, Graph-Based Anomaly Detection, Payer–Provider Network Analytics, Billing Pattern Analysis, Medicaid Payment Oversight, Healthcare Claims Analytics, Encounter-Level Fraud Detection, Normalized Payment Amount Modeling, Data-Driven Healthcare Compliance, Subrogation Data Integration, Network-Based Risk Scoring, Real-Time Payment Monitoring, Public Healthcare Program Integrity.

1. Introduction

Payment integrity—the validation of payers’ adherence to contractual and regulatory stipulations governing health plan benefits—is paramount in healthcare. Unquantified issues in payment accuracy erode efficiencies, inflate expenditures for both payers and providers, disrupt provider-payer relationships, and diminish trust in healthcare funding. Billing amendments based



on overpayment subrogation investigations detect only a sliver of misallocated claims, and fraud-control efforts remain subscale relative to the economic stakes. The critical mass of diverse payer-provider interactions embodied in healthcare billing data presents fertile ground for advanced anomaly-detection analyses capable of addressing the accuracy challenge.

Research focuses on unsupervised techniques that mitigate the dependence on expert labelers and on graph neural networks that leverage the triple complexity of billing data centered on the relationships between patients, payers, providers, and facilities. The objectives are twofold: to deploy a composite set of unsupervised models to uncover unexpected patient-billing characteristics and relationships and to provide a pertinent set of fraud- and risk-control flags from within the research community that can be fed back for validation into business decision support systems in near real time. Accepting that numerous anomalies exist and that neither domain expertise nor sufficient labeled training data can be marshaled, the analyses embrace an exploratory approach that acknowledges the data in aggregate and seeks to invoke multiple lenses for anomaly detection as a basis for discourse.

1.1. Overview of the Study

Accurate healthcare payments are critical for the sustainable functioning of the U.S. healthcare system, but payment accuracy relies on many manual and automated checks looking for anomalies—higher probability of inaccuracy, fraud, or wastage. Anomaly detection in billing data is an active area of research, and given the sensitivity of the healthcare data, labeled anomaly samples are not always available. Most research works there also employ supervised methods. Meanwhile, the recent attention on using graphs for representation learning has led to the foundation for using Graph Neural Networks to learn relationships between entities with varied degree of mappings. Healthcare billing data exhibit rich relationships: between payers and providers, providers and facilities, facilities and services, procedures and diagnosis codes, and

others.




Fig 1: Unsupervised Graph Representation Learning for Healthcare Payment Integrity: Detecting Billing Anomalies through Multi-Entity Embeddings

Deep learning Graph Neural Networks require large amounts of labeled data for effective learning, with important security and privacy considerations. Payer-provider representation sets created automatically from billing data are useful to payer groups seeking to understand anomalies in claims. Unsupervised deep learning processes based on representation sets do not require labeled samples, however. Using clustering or density-distribution based anomaly detection on payer, provider, facility, and claim embeddings has a similar effect on accuracy improvement compared to building an expert correctness inference model without labeled data.



	claim_id	payer	provider
1183	C01184	P05	PR036
909	C00910	P08	PR011
733	C00734	P01	PR059
802	C00803	P08	PR040
1021	C01022	P01	PR043
860	C00861	P08	PR059

2. Background and Motivation

The effectiveness of advanced analytics for healthcare payment integrity hinges on the accuracy of models targeting billing anomalies. This capacity grows with the detail and heterogeneity of billing data domains—especially claims and encounter data. These domains capture billing patterns defined by health insurers that impact financial performance through subrogation recoveries. However, a dearth of unsupervised approaches applies to anomaly detection; related work rests on labeled examples where the true underlying distributions and relationships remain hidden. Detecting anomalies without supervision curbs design bias, leveraging data richness rather than testing on engineered outliers. Anomalies arise from real-life deviations rather than being devised for study.

Billing data captures numerous rich relationships among patients, providers, insurance plans, procedures, and other entities in payer-provider interactions. These relationships span billed services, actual reimbursements, and coverage. Unsupervised methods can harness the value embedded in these relationships, with recently proposed graph neural networks especially promising for modeling linked structures across domains. Deep neural networks offer powerful representation-learning mechanisms, and a message-passing mechanism supports various types of relationships. Utilizing these strengths enhances billing models focused on integrity and fraud deterrence and enables billing anomalies and potential early fraud detection. This direction is salient in light of the large financial stakes involved in billing inaccuracies, both for payers and for fraudulent providers who deploy the same tactics against multiple payers.



Equation 1) Core quantities the paper relies on

1.1 Entity–relationship billing graph

Let the billing graph be:

$$G = (V, E)$$

- V : nodes (patient/provider/facility/payer/procedure/encounter, etc.)
- E : edges (relationships), possibly typed: $r \in \mathcal{R}$

A **node feature vector** for node v is:

$$\mathbf{x}_v \in \mathbb{R}^{d_x}$$

An **edge** can have attributes too:

$$e = (u, v, r, t), \quad \mathbf{a}_e \in \mathbb{R}^{d_a}$$

where t is timestamp (paper stresses temporality).

2.1. Scope and Relevance of Healthcare Billing Data

Healthcare billing data are rich and heterogeneous; corresponding dimensions in labels, feature sets, data quality, privacy, and governance offer strong foundations for unsupervised learning at scale. Claims data contain billing statements detailing claims submission, payment, and denial; these tables reflect key relationships among stakeholders, description and reimbursement of services, and structural constraints of medical procedure pricing regimes. Encounter data document services rendered; they describe patients, providers, facilities, and procedures used; and they denote service timing. Subrogation data contain information on liability identification for claim payment. Payment data account for the flow of capital and generate costs. Together, these four domains provide a near-complete picture of paid claims considered for automatic anomaly detection.

The multi-dimensional nature, complemented with respect to time, enhances modeling capabilities. Binary labels for specific types of fraud allow semi-supervised learning, while




completeness and representational power enable a two-stage paradigm, where general anomaly detection is followed by detection specialization. Privacy and governance concerns can be addressed through sampling, anonymization, and appropriate data exploitation frameworks.

3. Data Foundations for Healthcare Billing

Comprehensive modeling of billing anomalies relies on solid data prerequisites. Key stages that constitute core elements of a healthcare billing data foundation—data sources and integration, quality and preprocessing, and lifecycle management—are briefly presented. Healthcare organizations employ a multitude of systems to manage business operations in a digital form. To cover the entire range of business operations, claims data originate from billing systems; payments data, from payment systems; encounters, from the population health or care management systems; and subrogation data, from the subrogation management system. Following integration and centralization in a data lake, multiple dimensions of the data, each to an extent, are used. A representation model reflecting at least six types of entities—patients, service providers, service facilities, procedures, payer plans, and encounter events—and their pairwise relationships forms the basis for model development.

Healthcare billing data collected by payers consist of claims, encounters, subrogation, and payments information. Claims data detail service requests or authorizations submitted to payers by service providers or members; payments information specifies payments settled by payers; encounters data reports contacts between patients and the healthcare system; and subrogation information indicates patients' involvement in accidents. Claims, payments, and subrogation datasets are provided by the payments division; population health or care management systems supply encounter details. The breadth and heterogeneity of payment-related data, engaging diverse constituents and use cases, render it particularly suitable for completely unsupervised modeling.



3.1. Data Sources and Integration

Healthcare billing data scope encompasses claims, encounter, payment, and subrogation records. Major sources comprise payer information systems (primary source for all domains), clinical management databases (encounters), electronic billing systems (claims), and subrogation services (subrogation). Data integration follows a zero-party approach: source systems supply extracted data snapshots in a common schema (e.g., Healthcare Claims Data Structure). Schema alignment and entity resolution employ DBpedia-Mapping-based similarity metrics and supervised learning, with integration orchestrated by Apache NiFi and supported by Talend.

Data integration across distinct domains, governed by different siloed systems, stands out as the core data-centric challenge. The analysis utilizes integrated snapshot-style data from disparate systems, with a dedicated ingestion pipeline transforming updates to a single cross-domain schema. Claims, subrogation, and payment data undergo routine ingestion and integration, while data for encounters is supplemented on-demand with wider temporal joins.

Payer systems typically expose structured claims data as processed files, continuously updated to reflect the latest status of claims (e.g., additional payments, subrogation claims).



Encounter records released by clinical management systems capture clinic visits and hospital stays. Such records not only document actual occurrences of services, procedures, and treatments but also furnish critical information (diagnoses, procedures, supporting services) against which claims are validated.




Fig 2: Harmonizing the Billing Cycle: A Zero-Party Data Integration Framework for Cross-Domain Healthcare Revenue Integrity

3.2. Data Quality and Preprocessing

Data quality comprises various dimensions that, when properly addressed, enhance the value of the data for analytical purposes. Completeness indicates the extent of missing information due to data censoring, transferring non-covered patients to another line of business, or the reverse in the case of low-cost subrogation events involving a cover line within certain limit values; it is handled by applying simple censoring rules. Accuracy relates to potential misleading values; for example, claims billed later than a certain time window are removed to reduce spurious associations. Consistency measures conflicting information within the same dataset and between source data; for instance, when patients are part of multiple subrogation groups, the largest group for each month takes precedence, with others deleted for top-n processes. Lack of consistency across datasets is naturally expected due to the use of diverse data domains. Duplication refers to identical yet redundant entries in the same table, two or more of which having been retained in the addition process. Furthermore, in analytic contexts, normalization plays a key role by reducing the degree of sparsity associated with the set of features used in the model: for example, length-of-stay can be adjusted by a ratio with the same measure for patients within the same major diagnostic category (MDC). These considerations guide data preparation, which also encompasses duplicates deletion, textual data cleaning,



normalization of features central to the models being developed, handling of missing values, and censoring of specific entries when deemed necessary.

The segregation of data or some of its subsets into processing, training, calibration, and validation is another relevant decision that profoundly affects the resultant models. In the case of such unsupervised statistical models, an initial processing phase focuses on validation. Operating on the complete set of data across all available domains unveils a comprehensive overview of the historical patterns, associations, and direct influences present throughout the system. This combined model must be viewed as an exploratory environment exploring all possible combinations without constraining relationships, and run without measuring, scoring, or thresholding its output. Once the behavior of the entire set is fully understood, it is possible to create a working application dedicated to the specific scope or set of tasks, as is commonly done when exploiting supervised algorithms.

Equation 2) “Model-enriched Normalized Payment Amount” (and related normalization)

2.1 Generic feature normalization (example given: LOS)

Let:

- LOS_i = length of stay of patient i
- $MDC(i)$ = patient’s major diagnostic category group
- $\overline{LOS}_{MDC(i)}$ = average LOS in that MDC group

Step-by-step normalized LOS ratio:

1. Compute group mean:

$$\overline{LOS}_g = \frac{1}{|S_g|} \sum_{j \in S_g} LOS_j$$

2. Divide individual by its group mean:



$$LOS_i^{norm} = \frac{LOS_i}{\overline{LOS_{MDC(i)}}}$$

Interpretation: $LOS_i^{norm} > 1$ means longer-than-peer stays.

2.2 Normalized Payment Amount as a z-score within a relationship

A practical “model-enriched” normalization consistent with the paper’s payer–provider relational framing is:

Let a claim i belong to payer–provider pair (p, pr) . Define:

- A_i = paid amount
- $\mu_{p,pr}$ = expected mean payment for that payer–provider relation (learned from history/model)
- $\sigma_{p,pr}$ = expected std deviation for that relation

Step-by-step z-score:

1. Compute residual:

$$r_i = A_i - \mu_{p,pr}$$

2. Scale by expected variability:

$$Z_i = \frac{r_i}{\sigma_{p,pr}} = \frac{A_i - \mu_{p,pr}}{\sigma_{p,pr}}$$


- If $Z_i \gg 0$: unusually high paid amount for that payer–provider relationship
- If $Z_i \ll 0$: unusually low paid amount (may indicate underpayment, miscoding, reversals, etc.)

4. Methodological Framework



The conceptual design for detecting billing anomalies employs unsupervised machine learning algorithms in conjunction with a graph neural network model. Unusual billing events or patterns relative to the respective relationships (provider-payer, community service usage, etc.) are identified using unsupervised anomaly detection algorithms. Relevant features are extracted from the underlying entity-relationship framework, and models are parameterized, trained, and evaluated. The resulting anomaly score is then standardized and thresholded to create an alerting mechanism for these types of events.

The well-documented research interests of the machine learning community facilitate the unambiguous selection of applicable unsupervised techniques. A non-exhaustive sample of techniques utilized in anomaly detection includes clustering methods, density-based methods, temporal pattern monitors, and various types of autoencoders. The designs of these models generally consider the tasks of feature generation, anomaly scoring, calibration and thresholding, and explainability. These considerations are leveraged in the definition of the anomaly detection portion of the overall framework.





4.1. Unsupervised Learning for Anomaly Detection

Unsupervised learning techniques are fundamental to the anomaly detection approach. Clustering algorithms such as K-means or DBSCAN reveal unusual group memberships, density-based methods like Local Outlier Factor expose data points residing in low-density regions, models identifying abnormal temporal patterns or trajectories preprocess input via K-means clustering, and autoencoders detect anomalies by measuring reconstruction errors. Anomaly scores computed using these techniques are typically combined with a threshold to confirm or reject anomalous statuses.

A critical but often overlooked design dimension is the interpretability of model responses. Anomalies materialize in varied ways, from dormant behavior suddenly becoming active to the opposite pattern. Distinct explanations for distinct classes of anomalies enhance the ability of model outputs to inform domain experts.


4.2. Graph Neural Networks for Billing Relationship Modeling

In this research, Graph Neural Networks (GNNs) model interconnections in billing. Given the necessity of representing relationships between multiple billing parties and the inherent relational nature of billing data, GNNs were deemed suitable for the task. Several aspects characterize the selected GNN design.

First, both entity-types and edge-type attributes are important to describe the relationships in billing. Entities to represent in the billing relationship graph encompass patients, providers, facilities, payers, and service providers (or procedure codes). Edge attributes capture payer-provider pairs, facility-provider relationships, type of medical services rendered (procedures associated with claims), and temporal edges linking events across contributions. The model architecture is a message-passing-based GNN, although using relational graphs enabling multi-type messages may yield improved predictive performance. In particular, a distinct message for each payer-provider pair can enhance model expressiveness, enabling a richer representation of relationships with varying trust levels, fraud prevalence, and interdependence behind some payer-provider pairs than others. Apart from expressiveness, generalization may also be a concern. The number of patients and providers involved in billing is typically limited, further limiting the surfaces of relations to learn. Special attention is thus needed when developing a training model for GNN-based anomaly detection.



Second, the relative reliability of distinct sources of knowledge about the billing graph may differ. In particular, patient-provider relations have low probability to indicate fraudulent activity due to patients' limited choices of physicians in prevalent healthcare systems, while facility-provider links can be anticipated to have a larger proportion of suspicious events. Third, the temporal aspect is crucial for GNN modeling of billing-related events. The validity of claims is intrinsically subject to time: as time progresses, previously approved claims may be overturned when revised claims receive authorizations or fraud detection systems unveil incorrect billings. The temporal aspect affects also causality in the events: some procedures are prerequisites of certain treatments, external health events like accidents are prerequisites for subrogation claims, etc.



Equation 3) Unsupervised anomaly scores

3.1 K-means (distance-to-cluster anomaly)

Assume we embed/feature each claim as $\mathbf{x}_i \in \mathbb{R}^d$.



K-means objective:

$$\min_{\{\mathbf{c}_k\}, \{z_i\}} \sum_{i=1}^n \|\mathbf{x}_i - \mathbf{c}_{z_i}\|^2$$

- $\mathbf{c}_k$ = centroid of cluster k
- $z_i \in \{1, \dots, K\}$ = assignment

Step-by-step update rules:

2. Assignment step (choose closest centroid):

$$z_i = \operatorname{argmin}_k \|\mathbf{x}_i - \mathbf{c}_k\|^2$$

3. Centroid update (mean of assigned points):

$$\mathbf{c}_k = \frac{1}{|C_k|} \sum_{i: z_i=k} \mathbf{x}_i$$

Anomaly score from K-means:

$$s_i = \|\mathbf{x}_i - \mathbf{c}_{z_i}\|$$

3.2 Local Outlier Factor (LOF)

LOF compares local density of a point to densities of its neighbors.

Let k be number of neighbors.

Step-by-step:

3. k -distance of i : distance to its k -th nearest neighbor:

$$kdist(i)$$

4. Reachability distance between i and neighbor j :

$$reachdist_k(i, j) = \max\{kdist(j), \|\mathbf{x}_i - \mathbf{x}_j\|\}$$



3. Local reachability density (LRD):

$$lrd_k(i) = \left(\frac{1}{|N_k(i)|} \sum_{j \in N_k(i)} reachdist_k(i, j) \right)^{-1}$$

4. LOF score:

$$LOF_k(i) = \frac{1}{|N_k(i)|} \sum_{j \in N_k(i)} \frac{lrd_k(j)}{lrd_k(i)}$$

- $LOF_k(i) \approx 1$: similar density to neighbors (normal)
- $LOF_k(i) \gg 1$: much lower density than neighbors (outlier)

5. Model Architecture and Training Paradigms

Delineate concrete modeling components and training strategies.

Representation Learning for Billing Entities

To effectively calibrate detection capabilities, it becomes imperative to compute dedicated embeddings for the key billing entities, encompassing patients, billing providers, medical facilities, procedures, and participating payer plans. In the case of patients and records, dedicated numerical vectors are created through a multimodal approach that integrates traditional bag-of-words and temporal encoding techniques, thereby quantifying textual claims, the duration of preceding medical engagement, and the recency of medical activity. In the same vein, a hybrid representation learning strategy is employed to derive specialized temporal embeddings for the set of providers and facilities involved within the billing records throughout the designated period.

The remaining entities lack associated sequential information. For service codes utilized within the records, a Prototypical Network architecture is harnessed such that it guarantees similar embeddings for semantically equivalent codes and unique vectors for those that are specific while relying on a small set of labeled instances. Finally, an additional categorical embedding step incorporates the Payor Plan information. The complete training yield entity



vectors that engender expressiveness across all dedicated functions and modalities, along with a significant concurrence across time horizons and primary instances.

Temporal and Sequential Considerations

The event-based nature of the records implies the presence of time information. During training and testing, such time can be harnessed to enrich the dataset by creating sequences of events, limiting their dimension/numbers through monotonically sliding windows, and enforcing a causality constraint between antecedents and successors. This temporal information opens up the possibility of investigating the temporal aspect of billing, enabling the definition of Event-based Temporal Graphs and the subsequent implementation of time-aware Graph Neural Networks.

5.1. Representation Learning for Billing Entities

Anomaly detection models built upon data-billing-group and data-graph-group alternate in representation learning tasks on different entities, capturing complementary aspects. A small amount of specialized data may be sufficient for unsupervised representation learning, making it particularly compelling to integrate learned embeddings into domain-adapted models (e.g., autoencoders, classifiers). Popular implications include patients, providers, facilities, procedure codes, and payer plans. A multi-task approach is suggested with the objective of closely aligning the results of representation learning with the specific needs of all downstream tasks.

Multi-task representation learning of entities across sources and time supports accurate modeling of entities like patients and providers with diverse interactions. Patients exhibit temporal information through visits and billing events, which could be leveraged to improve the quality of representation learning. Hence, the resulting representation for patients might be further enhanced by capturing visit and billing sequences in addition to their other interactions. Another axis accidentally represented through the diverse interactions of patients is the value distribution, which can be exploited with domain-specific modeling techniques. The proposed methods also address an overlooked aspect for patients, providing cues to avoid misinformation that others do not pay attention to when predicting a behavior not intended to be leaked to the system.




Fig 3: Multi-Task Synergistic Embeddings: A Dual-Axis Representation Learning Framework for Healthcare Anomaly Detection

5.2. Temporal and Sequential Considerations

For time-stamped events, sequences are crucial in many settings, typically limiting the design of supervised algorithms (e.g., posing convolutions on time-stamped events). A sliding window approach is often used, requiring predefined sequences and disregarding causal dependencies. A hidden Markov model (HMM)-based direction offers maximum likelihood adaptation and dynamically adapts the number of sequences. Such approaches fit well with representation learning and unsupervised tasks.

When defined and understood in a broader context, temporality allows the use of more complex structures than simple sequences or the HMM perspective. More generally, when




processing, attending, or interacting with subjects having multiple events, it is natural to model the order (or temporal ordering) of these events. A natural construct is then a temporal graph, with marked edges containing the timestamps, but specific characteristics associated with time-stamped or sequential elements can also be applied independently to entities, followed by the synchronous construction of HTN and STN. The particular case of causal sequential information in multi-modal contexts can also be considered to obtain directional relationships that may assist training.

6. Evaluation and Validation Strategies

Rigorous assessment establishes the validity and utility of the approaches for detecting anomalies in healthcare payments. Evaluation and validation strategies are framed according to the intended use of the derived anomaly scores: supporting the capability to identify uncommon billing events and detecting relational inconsistencies during the temporal evolution of the billing data. Both aspects are critical to enabling any downstream forensic analysis.

Various metrics are defined for evaluating the performance of anomaly detection methods. Commonly used measures include precision, recall, F1-score, area under the receiver operating characteristic (ROC-AUC) curve, and area under the precision-recall (PR-AUC) curve. Metrics explicitly considering the cost of billing a patient are also valuable since minimizing the total billing cost is a natural objective for many healthcare payers. When the true workload-induced anomaly-suspected nature of the data is unknown, appropriate calibrators can be employed to determine optimal thresholds, allowing account for any general over- or under-estimation of the models.

Establishing ground truth for anomaly detection is generally difficult. Signature datasets of labeled examples or human-annotated resources are rare, particularly in healthcare. Two basic approaches can be employed to obtain labeled benchmarks: (i) expert-designed sets, in which a panel of healthcare specialists chooses a set of billing event anomaly signatures, and (ii) data synthesis, involving the artificial generation of labeled case data. Notably, semi-supervised techniques can be used to fill an example ground truth-loading gap by exploiting sets of expert-validated major or minor class instances. Human co-speaker adjudication permits supervised input into the learning phase while relaxing the need for full human annotation. Special precaution should be taken to ensure the representativeness and non-overlap of the ground-truth set, particularly when temporal and feature correlation may spur biases toward a detection model.



6.1. Evaluation Metrics for Anomaly Detection

Evaluation of billing anomaly detection requires assessment metrics that quantify model performance, indicate suitability for downstream objectives, and support informed decision-making by end users. In addition to standard measures such as precision, recall, F1 score, and area under the ROC curve, cost-based analysis can yield actionable insights.

Classification performance hinges on the intrinsic trade-off between false positives and false negatives, which the respective cost tolerances of analysts or business units largely dictate. A prominent use-case illustrates the potential for significant gains from the deliberate tuning of threshold values: Although a large controller organization seeks to minimize false positives so as to avoid wasting auditor resources, a small business unit incurs a sizeable cost for every missed error and constitutes the most significant source of return on investment for the audit program. At the other extreme, custom operating procedures of certain other business units promote rapid yet thorough claims-processing, leading to a relatively high volume of low-cost de-duplication requests. Consequently, crafting a distinct threshold, or calibrated set of thresholds, for each business unit produces large efficiency wins.

In anomaly detection, poor-class performance significantly undermines overall utility. For example, stakeholders invariably review the false positive candidates to minimize the chance



of missing real problems; hence, even a small number of false positives incurs fiscal loss. Such cardinal-false-positive-centric evaluation translates naturally into Precision@Specific-Threshold and, a fortiori, PR-AUC. In this manner, model formed aims and stakeholder needs directly shape evaluation. A comprehensive overview of performance metrics follows.

Equation 4) “Standardized and thresholder” alert mechanism

Assume a raw anomaly score s_i .

4.1 Standardization (z-score over scores)

4. Compute mean and std across all scores:

$$\mu_s = \frac{1}{n} \sum_{i=1}^n s_i, \quad \sigma_s = \sqrt{\frac{1}{n} \sum_{i=1}^n (s_i - \mu_s)^2}$$

5. Standardize:

$$\tilde{s}_i = \frac{s_i - \mu_s}{\sigma_s}$$

4.2 Thresholding (create an alert flag)

Pick a threshold τ (e.g., top 1–2% of scores, or calibrated rule):

$$Alert_i = \begin{cases} 1 & \tilde{s}_i \geq \tau \\ 0 & \text{otherwise} \end{cases}$$

6.2. Validation Data and Ground Truth Acquisition

Obtaining labeled data for anomaly-detection systems can prove exceptionally difficult as it typically requires a precoding classification by a subject-matter expert. This requirement creates biases in the validation set since very few anomalies exist in healthcare claim data. Several alternative approaches might prove beneficial: a class of well-defined fraud schemes might exist along with positive and negative examples that can provide the basis for a semi-supervised validation strategy; it is also possible to synthesize anomalies to perform task-specific



training and validation; finally, expert review of a small number of instances may help disentangle anomalies from normal observations. Any of these approaches can help obtain some level of reliable ground truth, although their inherent limitations will remain present.

Once a validation benchmark of some kind has been established, the appropriate metrics can be computed and used to guide model thresholds. Cost-based metrics integrating anomaly counts with the cost implications of errors can deepen the evaluation by showing how well the model supports real-life decisions even without a labeled validation set.

7. Conclusion

Sound evidence and scientific rationale support the growing recognition that billions of taxpayer dollars are drained through inaccuracy and fraud in the healthcare payment system. At the same time, methods for detecting billing anomalies with advanced analytics—especially those involving unsupervised machine learning techniques—remain relatively immature. This study proposes an effective strategy for detecting such anomalies using several complimentary unsupervised algorithms and graph-oriented neural networks, which represent the relationships and interactions among the various billing entities as a graph and embed the discrete nodes

across source systems.




Fig 4: Healthcare Payment Integrity Drivers

The results are expected to enrich the dialogue between payers and providers regarding payment accuracy and lay the groundwork for future directional work. Such work includes further improving the detection accuracy with appropriate validation data, expanding the detection to more service categories such as subrogation and consumer finance, doing so in near-real time, and supporting policy against healthcare fraud at the national level. Multiple directions for enhancing the model architecture and training paradigms have also been identified.

7.1. Final Thoughts and Future Directions

The value of accurate healthcare payment systems, rooted in appropriate billing, lies in their potential to ease financial burdens on payers, providers, and patients. However, healthcare payment systems harbor abundant inefficiencies that deeply affect their performance. The



viability of supervised machine learning for billing anomaly detection is limited by the high cost of accurately labeled examples and by the high dimensionality and complexity of the search space for anomalous billing behavior. An alternative approach tackles the problem using unsupervised learning to identify billing entities and their relationships, and to construct a set of models whose anomalies can expose spurious payment behavior. A more methodical exploration of these relationships is accomplished by the introduction of a graph neural network representation.

The detection of distinct forms of anomalous billing behavior has been previously conceived using simpler models, and the proposed framework further broadens the scope of detection by inferring richer sets of relationships from heterogeneous billing data spanning multiple years. The fully unsupervised framework is further applicable to other complex C5 domain problems with rich heterogeneous data, such as the detection of anomalous patterns in fraud, criminal, or cyberspace data. Supervised anomaly detection methods pursue the identification of anomalies in the format of clusters, sub-clusters, or denser regions of the feature space; however, for data theft and fraud cases, normal patterns frequently share little in terms of structure or data distributions. These models identify unexpected side effects of fraud incidents—such as telecommunication failures following telecommunication frauds—rather than the frauds or thefts themselves, and real-time monitoring of ongoing and new data streams appears unfeasible.

References

1. Ekin, T., Frigau, L., & Conversano, C. (2021). Health care fraud classifiers in practice. *Applied Stochastic Models in Business and Industry*, 37(5), 890–903.
2. Kanksha, Bhaskar, A., Pande, S., Malik, R., & Khamparia, A. (2021). An intelligent unsupervised technique for fraud detection in health care systems. *Intelligent Decision Technologies*, 15(1), 101–112.
3. Inala, R. (2023). AI-powered investment decision support systems: Building smart data products with embedded governance controls. *Journal for ReAttach Therapy and Developmental Diversities*, 6(10), 2251–2266.
4. Musal, R. M., & others. (2021). How to detect healthcare fraud? A systematic review. *Gaceta Sanitaria*, 35(Suppl. 2), S441–S449.
5. Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2021). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 32(1), 4–24.



6. Ying, C., Cai, T., Luo, S., Zheng, S., Ke, G., He, D., Shen, Y., & Liu, T.-Y. (2021). Do transformers really perform badly for graph representation? *Advances in Neural Information Processing Systems*, 34, 28877–28888.
7. Kolla, S. K., & Mangalampalli, B. M. (2024). Edge-Based Deep Learning Systems for Point-of-Care Diagnostic Intelligence. *Journal of Neonatal Surgery*, 13(1), 2387-2399.
8. Wang, J., Ma, A., Chang, Y., Gong, J., Jiang, Y., et al. (2021). scGNN is a novel graph neural network framework for single-cell RNA-seq analyses. *Nature Communications*, 12, 1882.
9. Li, M. M., Huang, K., & Zitnik, M. (2022). Graph representation learning in biomedicine and healthcare. *Nature Biomedical Engineering*, 6(12), 1353–1369.
10. Gu, Y., Yang, X., Tian, L., Yang, H., Lv, J., Yang, C., Wang, J., Xi, J., Kong, G., & Zhang, W. (2022). Structure-aware siamese graph neural networks for encounter-level patient similarity learning. *Journal of Biomedical Informatics*, 127, 104027.
11. Inala, R., & Somu, B. (2024). Agentic ai in retail banking: Redefining customer service and financial decision-making. *Journal of Artificial Intelligence and Big Data Disciplines*, 1(1), 1-19.
12. Cho, H. N., Ahn, I., Gwon, H., Kang, H. J., Kim, Y., Seo, H., Choi, H., Kim, M., Han, J., Kee, G., Jun, T. J., & Kim, Y.-H. (2022). Heterogeneous graph construction and HinSAGE learning from electronic medical records. *Scientific Reports*, 12, Article 20893.
13. Zou, Y., Pesaranghader, A., Song, Z., Verma, A., Buckeridge, D. L., & Li, Y. (2022). Modeling electronic health record data using an end-to-end knowledge-graph-informed topic model. *Scientific Reports*, 12, Article 17868.
14. Settupalli, L., & Gangadharan, G. R. (2022). WMTDBC: An unsupervised multivariate analysis model for fraud detection in health insurance claims. *Expert Systems with Applications*, 205, 119259.
15. Davuluri, P. S. L. (2023). AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems. *AI-Augmented Sanctions Screening: Enhancing Accuracy and Latency in Real Time Compliance Systems* (December 15, 2023).
16. Kumaraswamy, N., Markey, M. K., Barner, J. C., & Rascati, K. (2022). Feature engineering to detect fraud using healthcare claims data. *Expert Systems with Applications*, 210, 118433.
17. Hou, Z., Liu, X., Cen, Y., Dong, Y., Yang, H., Wang, C., & Tang, J. (2022). GraphMAE: Self-supervised masked graph autoencoders. *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 594–604.



18. Rampášek, L., Galkin, M., Dwivedi, V. P., Luu, A. T., Wolf, G., & Beaini, D. (2022). Recipe for a general, powerful, scalable graph transformer. *Advances in Neural Information Processing Systems*, 35, 14501–14515.
19. Gottimukkala, V. R. R. (2024). Federated Learning Approaches for Fraud Detection in International Payment Systems. https://www.jisem-journal.com/download/118_JISEM.pdf.
20. Adebayo, J., Mardanbeigi, D., & others. (2022). A novel fraud detection and prevention method for healthcare claim processing using machine learning and blockchain technology. *Digital Applications in Archaeology and Cultural Heritage*, 26, 100122.
21. Yun, S., Jeong, M., Yoo, S., Lee, S., Yoon, S., Kim, R., Kang, J., & Kim, H. J. (2022). Graph transformer networks: Learning meta-path graphs to improve GNNs. *Neural Networks*, 153, 104–119.
22. Qiu, Y., Lin, F., Chen, W., & Xu, M. (2023). Pre-training in medical data: A survey. *Machine Intelligence Research*, 20, 147–179.
23. Kolla, S. K., & Reddy, V. A. R. (2024). Evaluating Cloud-Native vs. Hybrid Architectures for Health Benefit Administration Systems. *International Journal of Medical Toxicology and Legal Medicine*, 27(5), 1042-1053.
24. Lu, J., Lin, K., Chen, R., & Lin, M. (2023). Health insurance fraud detection by using an attributed heterogeneous information network with a hierarchical attention mechanism. *BMC Medical Informatics and Decision Making*, 23, Article 62.
25. Xiao, S., Bai, T., Cui, X., Wu, B., Meng, X., & Wang, B. (2023). A graph-based contrastive learning framework for Medicare insurance fraud detection. *Frontiers of Computer Science*, 17(2), 172341.
26. Yoo, Y., Shin, J., & Kyeong, S. (2023). Medicare fraud detection using graph analysis: A comparative study of machine learning and graph neural networks. *IEEE Access*, 11, 88278–88294.
27. Hancock, J. T., Bauder, R. A., Wang, H., & Khoshgoftaar, T. M. (2023). Explainable machine learning models for Medicare fraud detection. *Journal of Big Data*, 10, Article 154.
28. Mao, Q., Liu, Z., Liu, C., & Sun, J. (2023). HINormer: Representation learning on heterogeneous information networks with graph transformer. *Proceedings of the ACM Web Conference 2023*, 599–608.
29. Nikolentzos, G., Vazirgiannis, M., Xypolopoulos, C., Lingman, M., & Brandt, E. G. (2023). Synthetic electronic health records generated with variational graph autoencoders. *npj Digital Medicine*, 6, Article 83.
30. Inala, R. Designing Scalable Technology Architectures for Customer Data in Group Insurance and Investment Platforms.



31. Meng, X., & Zou, T. (2023). Clinical applications of graph neural networks in computational histopathology: A review. *Computers in Biology and Medicine*, 164, 107201.
32. Paul, S. G., Saha, A., Hasan, M. Z., Noori, S. R. H., & Moustafa, A. (2024). A systematic review of graph neural network in healthcare-based applications: Recent advances, trends, and future directions. *IEEE Access*, 12, 15145–15170.
33. Boll, H. O., Amirahmadi, A., Ghazani, M. M., de Moraes, W. O., de Freitas, E. P., Soliman, A., Etmiani, F., Byttner, S., & Recamonde-Mendoza, M. (2024). Graph neural networks for clinical risk prediction based on electronic health records: A survey. *Journal of Biomedical Informatics*, 151, 104616.
34. Hong, B., Lu, P., Xu, H., Lu, J., Lin, K., & Yang, F. (2024). Health insurance fraud detection based on multi-channel heterogeneous graph structure learning. *Heliyon*, 10(9), e30045.
35. Kolla, T. (2024). Graph Neural Networks for HCC Risk Adjustment and Interoperability. *International Journal of Science, Research and Technology*, 7(6), 13244-13255.
36. Leevy, J. L., Hancock, J., & Khoshgoftaar, T. M. (2024). Medicare fraud detection: A comparative study on the effectiveness of one-class and binary classification models. *International Journal of Computers and Applications*, 47(3), 1–12.
37. Piya, F. L., Gupta, M., & Beheshti, R. (2024). HealthGAT: Node classifications in electronic health records using graph attention networks. *Proceedings of the IEEE International Conference on Connected Health: Applications, Systems and Engineering Technologies*, 132–141.
38. Johnson, R., Li, M. M., Noori, A., Queen, O., & Zitnik, M. (2024). Graph artificial intelligence in medicine. *Annual Review of Biomedical Data Science*, 7, 345–368.