soloncouncil : Pages:1–17

Original Research

Combining Knowledge Representation and Machine Learning for Improved Healthcare Claims Fraud Detection

Thuy Linh Pham[1]

[1]Da Nang Institute of Digital Technology, Department of Artificial Intelligence, Tran Cao Van Street, Da Nang, Vietnam.

Abstract

Healthcare claims fraud detection has become increasingly vital as healthcare systems grow in scale and complexity. The sheer volume of claims, alongside the heterogeneous nature of clinical, administrative, and billing records, creates a challenging environment for conventional rule-based or purely data-driven methods. This paper introduces an integrative framework that leverages knowledge representation structures together with advanced machine learning techniques to enhance fraud detection outcomes. By encoding essential domain knowledge in a structured format, our approach captures the semantic relationships and logical constraints inherent in healthcare claims. These representations guide data preprocessing, feature engineering, and explainability to complement powerful predictive models that target various fraud patterns, from upcoding and unbundling to fictitious billing. We discuss how domain-specific ontologies, rule-based inference engines, and first-order logic statements can interact with supervised and unsupervised learning approaches to capture both explicit and implicit indicators of fraudulent behavior. We detail an empirical evaluation that tests the proposed integrated system on a large-scale claims dataset, comparing performance metrics against baseline methods. The results underscore the value of a knowledge-infused pipeline, indicating superior detection accuracy, reduced false positives, and heightened interpretability of fraudulent cases. By bridging symbolic knowledge representation with robust machine learning algorithms, the proposed method promises a more reliable and comprehensible means of addressing healthcare claims fraud detection.

1. Introduction

The healthcare industry manages vast repositories of patient data and financial transactions, generating voluminous claims that must be processed quickly and accurately [1]. Modern healthcare settings include multiple layers of data acquisition, coding practices, and reimbursement procedures, each of which can be exploited by bad actors who aim to inflate or fabricate claims for financial benefit. The heterogeneous nature of healthcare data, which might involve diagnostic codes, treatment details, provider profiles, billing patterns, and patient demographics, complicates the development of robust fraud detection mechanisms [2]. The existence of hidden correlations among these diverse data types complicates the modeling task for purely traditional or rule-based systems. Consequently, researchers and practitioners have sought to capitalize on advances in machine learning, deep learning, and knowledge representation to build sophisticated frameworks capable of detecting subtle patterns of fraudulent behavior.

A core difficulty arises from the fact that many healthcare providers submit claims under legitimate but nuanced diagnostic classifications [3]. There might be inflation of treatment volumes or strategic unbundling of services intended to maximize reimbursement. Simple, static rules fail to capture the continuous changes in coding standards or newly emerging fraudulent schemes, leading to significant gaps in detection. Automated systems that rely exclusively on labeled examples and data-driven modeling may also experience challenges, especially if the fraudulent classes are underrepresented or if adversaries evolve their methods to exploit known detection vulnerabilities [4]. This complexity highlights the necessity of integrating domain-specific rules, ontologies, and knowledge graphs into the detection pipeline, ensuring that the system not only reacts to observed historical data but also incorporates robust encodings of medical and billing knowledge.

Various factors underscore the importance of bridging knowledge representation and machine learning for detecting claims fraud. First, healthcare data often include structured fields such as ICD codes, CPT codes, and standardized terminologies alongside unstructured clinical notes [5]. A knowledge-based framework provides a unifying representation for these disparate data sources, mapping coded elements to relevant medical concepts and capturing semantically meaningful relationships. Second, domain knowledge in healthcare has evolved over decades, manifested in complex regulatory guidelines, standard definitions of legitimate care pathways, and well-known fraud scenarios. Encoding these aspects explicitly can reduce the system's reliance on purely empirical signals, leading to better generalization when patterns shift [6]. Additionally, knowledge-based methods enable transparent reasoning about flagged claims. For example, a system might highlight a contradiction in a claim indicating a highly invasive procedure performed shortly after discharge from a routine outpatient visit [7]. This transparency is pivotal in domains such as healthcare, where interpretability can expedite audits and legal proceedings.

While domain knowledge can be incorporated through rule-based systems, these have traditionally struggled with inflexibility and high maintenance costs. Machine learning, especially when augmented by large-scale data, can detect complex interactions that exceed human capability [8]. Yet purely data-driven approaches may lack interpretability and may inadvertently learn spurious correlations driven by sampling biases. By melding knowledge representation and machine learning, one can exploit the strengths of both. Symbolic approaches and logic-based formalisms can ensure interpretability and compliance with known clinical and billing guidelines [9]. In parallel, machine learning can detect emerging and unforeseen behaviors that rule sets or ontologies alone might miss. This synergy results in an intelligent system that can robustly handle the dynamic nature of healthcare claims and the ever-shifting tactics of fraudulent actors.

Detecting healthcare fraud is not only a matter of financial concern but also influences patient care quality, provider reputations, and regulatory compliance [10]. Excessive or fraudulent claims might prompt unnecessary procedures, putting patients at risk and distorting the allocation of healthcare resources. Therefore, timely and accurate detection of anomalies can improve patient outcomes and reduce financial strain on insurers and governments. The stakes are high, further underscoring the need for highly reliable, interpretable, and adaptive detection systems [11]. In many real-world scenarios, detection must happen in near-real time or with minimal latency to prevent payouts to potentially fraudulent claims. This imposes additional challenges in model design and inference speed, highlighting the need for optimized data processing pipelines that can handle large-scale claims with minimal computational overhead. [12]

In the evolving landscape of healthcare analytics, explainability and trust have emerged as critical features. Traditional black-box machine learning methods, while effective in certain tasks, may fail to address the regulatory requirements and legal complexities surrounding fraud investigations. By incorporating explicit symbolic structures, such as ontologies and rule-based constraints, fraud examiners and auditors can better interpret the system's decisions and trace the rationale behind suspicious-claim alerts [13]. This interpretability not only aids in compliance but also fosters the trust of stakeholders, including healthcare organizations and providers who might otherwise be skeptical of opaque algorithms.

This paper explores the intersection of knowledge representation and machine learning for healthcare claims fraud detection, proposing a comprehensive framework for encoding domain knowledge into a data-driven pipeline. We discuss methods for constructing healthcare-specific knowledge graphs, formulating inference rules derived from clinical guidelines, and integrating those with advanced machine learning architectures that balance flexibility and interpretability [14]. Furthermore, we present detailed mathematical formulations for classification and anomaly detection tasks, highlighting how domain knowledge can structure model features and direct attention to potentially fraudulent patterns. Through an empirical evaluation on large-scale real-world claims data, we benchmark our integrated approach against classic rule-based and purely data-driven methods. We highlight improvements in detection rates and reductions in false positives, while demonstrating enhanced interpretability through well-structured audit logs and knowledge-based explanations [15]. In the subsequent sections, we delve into the technical machinery of knowledge representation paradigms, logic and symbolic methods, hybrid modeling strategies, data integrity considerations, and empirical validation metrics. Ultimately, we seek to illustrate how bridging these domains can improve both the robustness and the reliability of healthcare fraud detection strategies, positioning them for broader adoption in an era defined by exponential data growth and adversarial sophistication.

2. Knowledge Representation Paradigms in Healthcare Fraud Detection

Knowledge representation plays a pivotal role in capturing domain expertise and clinical context for use in complex tasks such as fraud detection [16]. By employing symbolic structures, one can systematically embed notions of legitimate treatment pathways, relationships among medical procedures, and known red flags in billing behavior. A primary representation strategy involves ontologies, which define concepts, attributes, and relationships in a hierarchical fashion [17]. For instance, one might have an ontology that situates a specific billing code within a broader class of procedures, thereby enabling reasoning about procedural equivalences or redundancies. Such structured representations formalize domain knowledge that traditional feature extraction approaches often overlook.

An important concept in knowledge representation is the notion of a knowledge graph, which may encode data as nodes representing entities (such as patients, providers, or procedures) and edges capturing semantic relationships (such as "performs," "receives," or "codes-for") [18]. In a healthcare claims scenario, nodes might include unique identifiers for medical providers, specific claim entries, or clinical diagnoses. Edges can represent claim-submission events, referral patterns, or conceptual relationships that group billing codes together. By constructing such a graph, one can detect anomalies by identifying substructures that deviate from typical or permissible configurations [19]. For example, repeated edges of certain relationships might signal suspiciously high volumes of specific, high-reimbursement procedures by a single provider, especially when contextualized against peer-group norms.

The integration of ontologies or knowledge graphs with statistical methods can improve the interpretability of fraud detection systems. When an anomaly is flagged, the system can identify the exact node or edge relationships that contributed to the score, linking them back to known domain concepts [20]. This bridging capability is crucial for explaining results to auditors or experts in the field, who can then verify whether the suspicion is valid. The synergy between domain knowledge and data analysis is further underscored by the fact that healthcare claims often exhibit cyclical patterns, correlated with seasonality, provider specialties, or regional coding practices. A knowledge-centric approach allows for explicit modeling of these patterns, reducing the confusion that arises when data-driven models alone encounter seasonal or regional outliers. [21]

One might formalize essential aspects of domain knowledge using first-order logic expressions. Suppose we define a domain of procedures P and diagnoses D, with a relation R(p, d) indicating that procedure p is appropriate for diagnosis d. A simplified logic statement might read: [22]

$$\forall d \in D \, \forall p \in P, \; R(p, d) \rightarrow \text{ApprovedClaim}(p, d),$$

meaning that if a procedure is recognized as valid for a given diagnosis, the resulting claim is at least conditionally approved subject to other constraints. Conversely, we might encode a constraint that says certain procedures are never permissible under certain diagnoses, a statement that can be logically formulated to filter claims that violate standard medical practice [23]. Such representations enable explicit reasoning about the legitimacy of claims, guiding both unsupervised anomaly detection modules and supervised classifiers that exploit domain-specific feature engineering.

While symbolic representation offers clarity, it also raises the challenge of integration with large-scale data systems. Healthcare organizations often store millions of claims records, each potentially linked to multiple diagnoses, procedures, or reimbursement codes [24]. Efficient indexing and query processing over knowledge graphs is nontrivial, particularly when real-time or near-real-time detection is desired. An additional consideration is that the knowledge encoded in ontologies may lag behind rapidly evolving clinical practices or new fraud techniques. Consequently, the approach must be designed to accommodate updates to the knowledge base, including the introduction of new billing codes, revised medical guidelines, or emergent red flags flagged by regulators. [25]

Knowledge representation can also assist in bridging the semantic gap that arises when data in structured databases coexists with unstructured information in clinical notes. Although the primary data for fraud detection might be structured claim entries, textual descriptions of treatments can be mined for potential inconsistencies. Through entity extraction and mapping to standardized terminologies, textual data can be transformed into structured assertions that fit into the knowledge graph [26]. This transformation makes it possible to capture evidence hidden in free-form text, such as an unexpected mention of a procedure code that does not match the patient's recorded condition. The integration of textual knowledge with structured data can substantially improve the coverage and accuracy of fraud detection models.

From a methodological standpoint, it is crucial to strike a balance between over-constraining the system with static knowledge and allowing adaptive, data-driven updates [27]. Overly rigid ontologies or rule sets might fail to detect novel fraud patterns, while purely adaptive methods risk explainability and might generate excessive false positives. The middle ground lies in a hybrid architecture that leverages symbolic knowledge for interpretability, consistency checks, and direct enforcement of well-known constraints, while also harnessing machine learning for discovery of nuanced or emerging fraud signals [28]. This approach ensures that known and evolving aspects of the fraud problem are addressed in a unified framework.

Designing a robust knowledge-based system for healthcare fraud detection further involves considerations of system scalability, maintenance, and integration with legacy healthcare information systems. The data pipelines that feed into the knowledge graph must manage updates, merges, and potential conflicts in real time or near-real time [29]. Terminologies and coding systems must be aligned to avoid semantic drift. In practical implementations, various standardized vocabularies such as SNOMED CT, ICD-10, or CPT might need to be cross-referenced and harmonized. These steps require both technical solutions and governance processes to ensure that the knowledge layer remains accurate, up to date, and effectively utilized by downstream detection modules [30]. Despite these complexities, the benefits of harnessing domain knowledge—from improved interpretability to reduced reliance on purely empirical correlations—underscore the value of knowledge representation paradigms for healthcare fraud detection.

3. Logic and Symbolic Methods in Healthcare Claims Processing

Logic-based approaches have long been a cornerstone in expert systems, enabling explicit rule encoding and symbolic inference to guide decision-making. In the context of healthcare claims fraud detection, logic and symbolic methods can serve as a robust framework for incorporating high-level knowledge that might be difficult to capture using raw data alone [31]. This includes regulatory guidelines, clinical protocols, and standard billing practices. By formulating these guidelines in a logical framework, one can systematically verify claims against a known set of permissible or impermissible conditions.

A fundamental aspect of logic-based approaches lies in the use of first-order logic statements [32]. Consider a domain where each claim is composed of a set of procedures P, diagnoses D, and provider attributes A. One might define a predicate Fraudulent(c) which asserts that claim c is potentially fraudulent [33]. Additional predicates capture relationships like ProvidedBy(c, a) and JustifiedBy(c, d). A highly simplified logic statement might read:

$$\forall c \left(\exists d \in D,\ \neg\text{JustifiedBy}(c,d)\right) \implies \text{Fraudulent}(c),$$

meaning that if a claim has no justifying diagnosis, it raises the possibility of fraud [34]. More sophisticated logic statements account for partial or probabilistic information, leading to frameworks such as probabilistic logic or fuzzy rule-based systems.

The transition from purely Boolean logic to fuzzy logic can be valuable in healthcare claims analysis, as certain billing scenarios may be partially justified. For example, a procedure might be borderline appropriate for a given diagnosis, and domain experts may assign a certain confidence level to its legitimacy [35]. Fuzzy membership functions, such as $\mu_{\text{Legitimate}}(x)$, define the degree to which a claim x meets known guidelines. By combining multiple fuzzy rules, the system can generate a composite fraud suspicion score. Symbolic inference engines that support these fuzzy concepts can then integrate partially violated constraints, capturing nuances that purely discrete approaches might miss.

Symbolic reasoning extends to the concept of rewriting systems and unification, which can be relevant when handling complex or nested billing codes [36]. For example, certain procedures might be considered special cases of broader categories, or multiple codes might unify to represent a single type of clinical activity. Such rewriting rules can help normalize claims data into canonical forms, ensuring that the knowledge base applies consistently across diverse representations. Once normalized, claims can be systematically compared against known patterns of fraud, such as unusual service combinations or contradictory codes [37]. This approach ensures that logically equivalent, but syntactically varied, claims are treated uniformly.

Logic engines integrated into a healthcare fraud detection platform can also exploit backward chaining. Instead of forward-chaining through all possible inferences, the system begins with a suspicious claim hypothesis and attempts to find supporting evidence from the knowledge base [38]. This technique streamlines the investigative process, allowing automated or human auditors to trace the reasons behind each flag. For instance, if a claim is suspected of fraudulent upcoding, backward chaining can help identify that a certain procedure code is only permissible if the patient's medical record includes a higher-level diagnosis [39]. In the absence of that, the logic chain fails, supporting the fraud hypothesis.

Despite their strengths, purely symbolic systems might suffer from fragility in the face of incomplete data or unmodeled phenomena. Healthcare data can be messy, with missing or erroneous information, as well as rapidly evolving coding standards that may not be fully reflected in the knowledge base [40]. Symbolic systems also require careful curation and regular updates to their rule sets. If not maintained, the system risks becoming obsolete or generating an excessive number of false positives. This challenge can be addressed, in part, by coupling logic-based components with machine learning models that adapt to changing data distributions. [41]

Symbolic logic can also guide the feature engineering stage of data-driven models. Suppose we define a set of logic constraints that characterize suspicious provider behavior, such as performing an unusually high count of procedures per patient visit or claiming contradictory procedures. These constraints can be encoded as numeric or categorical features, reflecting the degree of violation of each logic rule [42]. A machine learning model, whether it is a neural network or a gradient boosting system, can incorporate these features, thus leveraging domain-driven signals. This hybrid strategy respects the interpretability of symbolic approaches while benefiting from the powerful pattern-recognition capabilities of modern machine learning architectures.

Once logic is integrated into the detection pipeline, there arises the question of conflict resolution [43]. In certain claims, multiple rules may produce conflicting inferences. For instance, one rule might declare a procedure set permissible for a certain diagnosis, while another rule might disallow it under certain comorbidities [44]. More advanced logic frameworks incorporate default reasoning or exception handling, where rules have priorities or exceptions to manage such conflicts systematically. The logic might specify that certain guidelines override others in special cases, reflecting real-world healthcare complexities. Such structured conflict resolution is vital for maintaining a coherent and accurate knowledge representation system. [45]

Another perspective on logic use in healthcare fraud detection is model checking. By modeling the claims processing workflow as a state machine and specifying properties that represent normal or anomalous transitions, one can apply techniques from temporal or modal logic. For instance, if a patient's claim history states that a certain procedure must be preceded by a qualifying test, one could specify a temporal logic property that flags any claim in which the procedure is billed but no matching test is found in the previous time window [46]. These advanced logic-based techniques provide a rigorous mechanism for verifying system states, further contributing to the robustness of fraud detection strategies.

Ultimately, logic and symbolic methods offer a structured way to incorporate domain knowledge directly into the detection process, enabling consistent, interpretable, and traceable analysis of healthcare claims. The method complements machine learning, which excels at finding hidden patterns in high-dimensional data but may fail to encode domain constraints naturally [47]. By combining these approaches, one can achieve a balance that addresses both known rules and emergent behaviors, forming a cornerstone for a hybrid framework that weaves together knowledge representation and machine learning for advanced healthcare claims fraud detection.

4. Hybrid Machine Learning Strategies

The interplay between knowledge representation and machine learning creates a rich environment for hybrid strategies that capitalize on symbolic logic while retaining the adaptability and predictive power of data-driven algorithms. In healthcare claims fraud detection, such hybrid strategies might begin by leveraging symbolic or ontology-based processes to transform or annotate raw claims data [48]. These annotations can highlight key medical codes, known contradictory service pairings, or patterns extracted from patient histories. The transformed data thus carry domain structure explicitly encoded, which can significantly improve learning algorithms' ability to detect subtle anomalies. [49]

A common architecture for hybrid approaches begins with a knowledge-based feature engineering layer. Consider a set of claims $\{c_1, c_2, ..., c_n\}$. Each claim c_i is represented originally as a set of codes $\{x_1, x_2, \ldots, x_k\}$ and contextual information such as provider ID, patient demographics, or date of service. By applying a knowledge-based transformation, one might derive additional features such as the average permissible frequency of each code over a defined period, or a measure of how often these codes co-occur in legitimate circumstances. Formally, define a feature extraction function ϕ_k derived from the knowledge base K [50]. Thus, each claim is mapped to a high-dimensional feature vector:

$$\phi_k(c_i) = (f_1(c_i, K), f_2(c_i, K), \ldots, f_m(c_i, K)),$$

where each f_j is a feature derived using domain-specific logic or ontology-based constraints [51]. These specialized features might signal, for example, potential unbundling if the combination of codes violates known billing guidelines.

Following this annotation, the data can be fed into a range of machine learning models. Classifiers such as random forests or gradient boosting machines can naturally handle high-dimensional, sparse feature vectors and incorporate the new, semantically rich features [52]. Deep learning models, particularly those employing attention mechanisms, can incorporate domain features into embedding layers, allowing the model to weigh the importance of each knowledge-derived dimension adaptively. Alternatively, autoencoder-based anomaly detection methods can use these domain-enhanced features to learn normal claim distributions, flagging claims that deviate significantly as potential fraud.

In some cases, knowledge-based constraints can be embedded directly into the learning algorithm [53]. For instance, certain neural network architectures allow for the inclusion of constraints in their loss function. Consider a labeling function $L : C \to \{0, 1\}$, where 1 indicates a fraudulent claim. The classification objective might be expressed as:

$$\min_{\theta} \left[\sum_{i=1}^{n} \ell\big(f_\theta(\phi_k(c_i)), L(c_i)\big) + \lambda \cdot \Omega(\theta) \right],$$

where ℓ is a standard cross-entropy loss, and $\Omega(\theta)$ is a regularization term [54]. One can incorporate logic-based constraints into $\Omega(\theta)$, for example by penalizing parameter configurations that violate known domain rules. This approach enforces consistency with established medical or billing guidelines, guiding the model away from spurious patterns that might yield high empirical accuracy but low interpretability or domain validity. [55]

A complementary perspective involves structured prediction, where the model must infer multiple interrelated labels simultaneously. For example, each element of a claim might be assigned a "legitimate" or "questionable" status, and these decisions must cohere with domain constraints. Let $\mathbf{y} = (y_1, y_2, ..., y_m)$ denote the predicted labels for the codes in claim c. A knowledge-based constraint might forbid certain combinations of $(y_1, y_2, \ldots, y_m)$ if they violate recognized medical logic [56]. The model's structured loss can be augmented with a constraint function $\Gamma(\mathbf{y}, K)$ that penalizes or outright disallows illegal configurations. These structured prediction methods can be particularly effective in capturing complex relationships within a claim, moving beyond the realm of single-label classification.

Hybrid approaches also open the door to advanced techniques such as multi-instance learning or relational learning, particularly valuable in capturing the interplay between multiple claims from the same provider or patient. For instance, a provider might submit hundreds of claims, and while any single claim could appear legitimate in isolation, the aggregate pattern might reveal suspicious billing frequencies [57]. A knowledge-based approach could group claims by provider, extracting features about cumulative billing, while a machine learning model operating on these groupings can more effectively detect anomalies. This relational perspective can be expressed with logic statements specifying normal or expected frequencies of certain procedures. The machine learning model can then incorporate aggregated features that measure deviations from these expectations. [58]

Another branch of hybrid systems employs active learning, where the knowledge base helps select the most informative claims for manual review. If the system is uncertain about a claim's legitimacy, or if a claim pattern is unprecedented with respect to the existing knowledge representation, it can be flagged for expert verification. Feedback from the expert can then update both the model (providing a labeled example) and the knowledge base (introducing new rules or ontology expansions if the claim is indeed an instance of a novel fraud method) [59]. This iterative cycle of model training and knowledge refinement can keep the detection system aligned with emerging trends in healthcare fraud, while continuously improving domain coverage and interpretability.

Some hybrid frameworks also employ generative modeling as part of the detection pipeline [60]. A knowledge base might specify constraints on how claims are validly generated. By building a generative model that samples from legitimate distributions under these constraints, one can approximate normal claim patterns. A separate discriminator model then evaluates whether real claims align with these synthetically generated patterns [61]. If a real claim diverges significantly from the synthetic normal set, it is flagged for additional scrutiny. The knowledge-driven constraints embedded in the generative model ensure that the sampling space captures clinically grounded claim structures rather than arbitrary data-driven distributions.

Designing and deploying hybrid strategies in production requires careful attention to computational overhead [62]. The knowledge-based preprocessing steps, such as graph queries or logic inference, can be expensive at scale. Therefore, indexing, caching, or approximate inference techniques might be necessary to handle real-time fraud detection over millions of claims. A layered architecture may be adopted, where a fast screening model quickly identifies obviously legitimate or obviously suspicious claims, and a more computationally intensive knowledge-based model is only invoked for borderline cases [63]. This architecture balances computational efficiency with the higher accuracy and explainability offered by deep knowledge integration.

Overall, hybrid machine learning strategies lie at the intersection of symbolic representation and data-driven modeling, harnessing the advantages of each domain. They incorporate domain knowledge to guide feature construction, shape loss functions, and enforce structured constraints, while retaining the ability to learn from large-scale data and detect novel fraud schemes [64]. This synergy provides a blueprint for robust healthcare claims fraud detection systems that are both effective at capturing emerging patterns and transparent in their reasoning, thereby meeting the dual challenges of adaptability and interpretability in an increasingly complex healthcare landscape.

5. Data Integrity and Clinical Semantics

Data integrity is paramount for any fraud detection system, but it assumes special significance in healthcare due to the sensitive, multifaceted nature of clinical and billing data [65]. Errors or intentional manipulations in diagnoses, procedure codes, or patient information can either hide fraud or generate a high rate of false positives, crippling the system's reliability. Ensuring data integrity demands a comprehensive strategy that includes cleansing, validation, consistency checks, and semantic alignment with recognized healthcare terminologies.

Clinical semantics refers to the layer of meaning that bridges raw codes and actual medical procedures, diagnoses, or patient states [66]. For instance, a certain procedure code might be conceptually subsumed by a broader procedure category, and a specific diagnosis code may share relationships with comorbidities or derivative conditions. These relationships constitute a semantic network that helps contextualize healthcare data. In the context of fraud detection, semantic knowledge can reveal inconsistencies [67]. If a claim references a complex surgical procedure performed on the same day as a conflicting condition, the semantic relationships among the procedure, diagnosis, and temporal context can alert the system to potential misconduct.

Data ingestion pipelines in healthcare fraud detection often rely on standardized terminologies like ICD for diagnoses, CPT for procedures, and LOINC for laboratory tests. While these terminologies are widely adopted, local customizations or older versions can create mismatches between systems [68]. One approach to ensuring semantic consistency involves mapping all codes to a canonical reference ontology maintained within the knowledge representation layer. This ontology might define hierarchical relationships, such as a partial order relation $\leq$ over diagnostic codes, indicating that one diagnosis is a specialized form of another. More formally, given a set of diagnoses D and a partial order $\leq$, one can ensure claims referencing a diagnosis d_j also inherit the properties of its parent d_i if $d_i \leq d_j$ [69]. Such an inheritance structure is essential for advanced reasoning.

Clinical semantics also informs the extraction of additional features for machine learning. For instance, a neural network trained on standard code embeddings might benefit significantly from a preprocessing step that aligns codes with their semantic neighbors in an ontology [70]. This alignment can be implemented using concept embedding techniques, where each code is mapped to a vector space in which semantic proximity correlates with vector similarity. By referencing an ontology, one can ensure that codes referencing similar medical concepts are also closer in the embedding space, thus guiding the model to identify suspicious patterns spanning multiple related codes [71]. For instance, if a provider consistently bills for procedures whose embeddings lie far from typical diagnoses for a given specialty, that pattern might warrant further scrutiny.

A crucial yet often overlooked aspect of data integrity is temporal consistency. Many fraud schemes exploit timing ambiguities, such as billing for a follow-up procedure before the primary procedure is recorded or claiming services for a patient who had already been discharged [72]. Advanced knowledge-based systems can incorporate constraints stating that certain procedures must follow or precede other events. Such temporal logic can be expressed formally. If $T(c)$ denotes the time of claim c, and $\text{Precedes}(c_1, c_2)$ states that c_1 must occur before c_2, one can write:

$$\forall c_1, c_2,\ \text{Precedes}(c_1, c_2) \implies \big(T(c_1) < T(c_2)\big).$$

Violations of these temporal constraints can highlight suspicious or impossible sequences, revealing potential data tampering or fraudulent claims. [73]

Patient-level data integrity checks further involve verifying personal information, coverage periods, and potential identity theft scenarios. Fraud rings sometimes use stolen patient identities to submit claims for services never rendered. By cross-referencing claims data against membership or insurance coverage databases, the system can flag patients who appear multiple times under different providers in improbable intervals or who receive contradictory treatments at overlapping times [74]. These checks are facilitated by the knowledge representation layer, which might store connections between patient entities, coverage rules, and claim events, thereby creating an integrated environment for anomaly detection.

Maintaining data integrity also involves addressing errors or inconsistencies introduced during data entry or coding. The healthcare ecosystem relies on a variety of personnel—from coders to administrators—who might introduce mistakes inadvertently [75]. While not fraudulent, these errors can degrade detection performance. Automated cleaning steps, such as verifying code validity against the active version of ICD, can catch basic inconsistencies [76]. More advanced logic-based consistency checks can identify claims that are unlikely or impossible based on known medical standards. This requires that the system remain flexible enough to mark suspicious anomalies without automatically classifying them as fraud, allowing for subsequent human review or automated correction.

In many cases, data integrity improvements rely on feedback loops from the fraud detection process itself [77]. If the detection system frequently flags certain types of claims as suspicious only to discover that they result from systematic data entry errors, domain experts can refine the knowledge base to account for these recurrent patterns. This iterative approach addresses the dynamic nature of healthcare data, where new billing codes or updated guidelines can modify the landscape of what is considered normal. By capturing this domain evolution in the knowledge base, the system remains robust against data drifts that might otherwise confound machine learning models [78].

Ethical and legal considerations are also integral to managing data integrity in healthcare fraud detection. Patient confidentiality and data protection regulations necessitate secure data handling, anonymization where appropriate, and strict access controls. Logic-based approaches can help enforce access policies, such as restricting the visibility of certain codes to authorized entities, while machine learning components benefit from the additional structural clarity about what data subsets can be legitimately processed together [79]. The synergy between semantic alignment, integrity checks, and compliance forms the cornerstone of a credible detection system that can effectively operate in the highly regulated healthcare domain.

In summary, data integrity and clinical semantics go hand in hand to form the foundation of a reliable fraud detection framework. By aligning data with standard ontologies, ensuring temporal and contextual consistency, and capturing domain knowledge in logical constraints, one can create a robust environment that supports advanced machine learning methods [80]. The resulting pipeline does more than simply detect anomalies: it orchestrates a medically coherent representation of claims, thereby reducing noise, supporting interpretability, and fostering trust in automated systems. These aspects collectively make data integrity and clinical semantics indispensable elements in the quest to unearth fraudulent practices while preserving the veracity of legitimate healthcare transactions. [81]

6. Empirical Evaluation and Performance Metrics

A compelling demonstration of the proposed integration of knowledge representation and machine learning for healthcare claims fraud detection lies in a rigorous empirical evaluation. Such an evaluation typically involves large-scale real-world claims datasets, enriched with domain knowledge structures like ontologies, rule sets, or knowledge graphs. The overarching goal is to measure improvements in detection rates, false positive reductions, and interpretability relative to purely data-driven or purely rule-based baselines. [82]

Dataset preparation begins with the gathering of healthcare claims records, spanning diverse providers, diagnoses, and procedures. Each record is assigned a known or suspected ground truth label indicating whether it is legitimate or fraudulent. In practice, assembling a definitive labeled set can be challenging, as suspected fraud often undergoes extended investigations before final confirmation [83]. Nonetheless, approximate labels might be drawn from historical audits or legal verdicts, while unlabeled data can be used in semi-supervised or unsupervised contexts. Once the data are compiled, the knowledge representation layer is established. This involves linking relevant codes to standardized terminologies, creating or updating ontologies that define relationships among medical concepts, and formalizing logic rules that specify permissible claim structures or known red flags. [84]

The experimental pipeline is structured to test different detection methodologies. One baseline might be a purely data-driven approach, such as a gradient boosting classifier trained on raw claim features. Another baseline might be a purely rule-based system that flags claims violating certain hard-coded constraints [85]. In contrast, the proposed hybrid approach applies knowledge-driven feature engineering or direct logic-based constraints integrated into a machine learning model. By systematically comparing the performance of these distinct strategies, one can isolate the added value of knowledge representation.

Common performance metrics include precision, recall, F1 score, and the area under the Receiver Operating Characteristic curve (ROC-AUC) [86]. Since fraud detection is often treated as a binary classification task, these metrics provide insight into a system's capacity to correctly identify fraudulent cases (true positives) while minimizing the burden of false positives on auditors. However, the evaluation extends beyond traditional classification measures [87]. In real healthcare settings, an excessive false positive rate can overwhelm human reviewers, so the precision at high recall or the precision at a given detection threshold is also critical. If the system identifies only a subset of suspicious claims for immediate audit, it is essential that these flagged cases be highly indicative of potential fraud to avoid resource wastage.

Another relevant set of metrics encompasses interpretability and time-to-audit [88]. One might measure the proportion of flagged claims for which the system can provide a knowledge-based rationale, such as citing a violated logic rule or pointing to anomalous connections in the knowledge graph. This rationale can be assigned a complexity measure, reflecting how easily a domain expert can comprehend the system's explanation. Moreover, the time needed for an auditor to validate a flagged claim could be tracked to gauge whether knowledge-driven explanations indeed streamline the investigative process. [89]

Empirical results might reveal, for example, that a knowledge-enriched model outperforms a purely data-driven baseline by a significant margin in terms of F1 score, particularly in low-sample scenarios. This is because symbolic features and constraints help guide the model even when fraudulent samples are sparse. Conversely, one might observe that the purely data-driven approach does well when large volumes of labeled data are available, but fails to adapt when new fraud patterns or codes appear [90]. The knowledge-based method, continuously updated with domain information, proves more robust in the face of changing data distributions.

A quantitative illustration might note that the precision for the top 5 percent of flagged claims rises from 40 percent in a baseline to 60 percent in the hybrid approach, reducing the strain on human auditors and potentially saving millions in misallocated reimbursements. Additionally, a case-by-case analysis might highlight how certain suspicious claim patterns, such as repeated submission of unbundled codes in a short time window, are more readily detected by the knowledge-based system that explicitly encodes constraints on consecutive procedures. [91]

Incorporating advanced logic statements into the detection pipeline can also change how system updates are deployed. Rather than retraining models with every minor shift in the data, new logic rules might be introduced to capture emergent fraud trends [92]. An empirical study might measure the adaptation time for these updates: a symbolic approach could incorporate new knowledge in hours or days, while purely data-driven models require comprehensive retraining on newly labeled data. By quantifying how quickly the system incorporates new rules and how these rules improve detection performance, one can illustrate the operational advantages of knowledge-based integration.

Unsupervised methods offer another dimension for evaluation, often using anomaly scores or outlier detection metrics [93]. A knowledge-enhanced autoencoder might reconstruct normal claim patterns more accurately than a purely data-driven autoencoder, leading to sharper distinctions when fraudulent claims deviate from legitimate structures. Metrics such as the reconstruction error distribution help visualize the difference in normal vs. anomalous patterns [94]. The presence of domain knowledge might significantly compress the variance of normal claim reconstructions, making anomalies stand out more clearly.

Scalability and runtime form another critical area of empirical evaluation. Large healthcare systems process millions of claims per day, so the computational overhead introduced by knowledge-based inferences can be nontrivial [95]. One might measure throughput (claims processed per second) and latency (time per claim) under different indexing strategies or inference algorithms. A well-architected system could maintain near-real-time throughput by caching frequently accessed ontology segments or precomputing certain logic-based features offline. The empirical study would demonstrate that the overhead remains manageable while delivering improved detection outcomes. [96]

In sum, a robust empirical evaluation for an integrated knowledge representation and machine learning approach encompasses multiple baselines, uses diverse metrics capturing both predictive performance and interpretability, and addresses practical considerations of scalability and adaptability. The results typically validate the core hypothesis: that the synergy of explicit domain knowledge with data-driven pattern recognition yields superior fraud detection capabilities, providing not just higher accuracy but also a deeper, more transparent understanding of potentially fraudulent claims [97]. This comprehensive evaluation closes the loop between theoretical design and real-world efficacy, offering compelling evidence for the viability of hybrid systems in modern healthcare environments.

7. Conclusion

This paper has examined the confluence of knowledge representation and machine learning to enhance healthcare claims fraud detection. We have discussed how ontologies, logic-based constraints, and symbolic reasoning offer a structured means of capturing medical domain expertise, billing guidelines, and clinical semantics [98]. These knowledge-based elements complement advanced machine learning models, which excel in pattern recognition and adaptation to emerging fraudulent behaviors. By combining symbolic and data-driven methods, the resultant frameworks gain not only improved accuracy and recall but also heightened transparency and interpretability—essential factors in a regulated domain such as healthcare, where trust and verifiability are paramount.

The presented perspective underscores how structured knowledge can guide feature extraction, shape loss functions, and define valid claim configurations [99]. It allows for the explicit encoding of rules governing permissible procedure-diagnosis pairs, expected treatment pathways, and provider-level norms. Simultaneously, machine learning approaches harness statistical patterns from large-scale claims data, capturing nuanced signals of fraudulent conduct that might elude purely symbolic systems. The synergy lies in bridging these dual channels of information, thereby reducing false positives and better highlighting genuine anomalies [100]. Empirical evaluations further reinforce that knowledge-rich approaches can be more robust when confronted with shifting codes, novel schemes, or limited labeled data.

Moving forward, future work might focus on designing architectures that tightly integrate ontological reasoning with deep learning layers, enabling near-real-time inference without compromising interpretability. Another promising area involves leveraging federated or distributed learning protocols to preserve patient privacy while still aggregating insights from multiple healthcare institutions [101]. The potential also exists for automated knowledge base enrichment, where discovered fraud patterns can be reverse-engineered into new logic rules or ontology expansions, thus ensuring a self-improving detection pipeline. By continuing to refine these ideas, researchers and practitioners can build comprehensive solutions that both reflect the sophisticated nature of modern healthcare data and anticipate the evolving methods of fraudulent actors. Through this advanced fusion of knowledge representation and machine learning, the healthcare industry gains a powerful tool for safeguarding financial integrity, regulatory compliance, and, ultimately, patient well-being. [102]

References

[1] L. J. J, K. Singh, and B. Chakravarthi, "Digital forensic framework for smart contract vulnerabilities using ensemble models," *Multimedia Tools and Applications*, vol. 83, pp. 51469–51512, 11 2023.

[2] A. Maghsoudi, Y. H. Sada, H. Zhu, S. Yarlagadda, D. Guffey, S. Nowakowski, A. Li, and J. Razjouyan, "Application of natural language processing to assess the performance status documentation quality metric in patients with non–small-cell lung cancer.," *Journal of Clinical Oncology*, vol. 41, pp. e13582–e13582, 6 2023.

[3] M. Ismail, T. Lalani, A. Kielar, C. Hong, J. Yacoub, C. Lim, V. Surabhi, K. Shanbhogue, S. Nandwana, X. Liu, C. Santillan, M. R. Bashir, and J. Lee, "Lessons learned: strategies for implementing and the ongoing use of li-rads in your practice.," *Abdominal radiology (New York)*, 10 2024.

[4] Y. Gao, D. Dligach, T. Miller, M. M. Churpek, O. Uzuner, and M. Afshar, "Progress note understanding - assessment and plan reasoning: Overview of the 2022 n2c2 track 3 shared task.," *Journal of biomedical informatics*, vol. 142, pp. 104346–104346, 4 2023.

[5] K. D. Forbus, C. Riesbeck, L. Birnbaum, K. Livingston, A. Sharma, and L. Ureel, "Integrating natural language, knowledge representation and reasoning, and analogical processing to learn by reading," in *Proceedings of the National Conference on Artificial Intelligence*, vol. 22, p. 1542, Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 2007.

[6] M. I. Miller, L. C. Shih, and V. B. Kolachalama, "Machine learning in clinical trials: A primer with applications to neurology.," *Neurotherapeutics : the journal of the American Society for Experimental NeuroTherapeutics*, vol. 20, pp. 1066–1080, 5 2023.

[7] C. Micale, S. Golder, K. O'Connor, D. Weissenbacher, R. Gross, S. Hennessy, and G. Gonzalez-Hernandez, "Patient-reported reasons for antihypertensive medication change: A quantitative study using social media.," *Drug safety*, vol. 47, pp. 81–91, 11 2023.

[8] J. Valinejad, Z. Guo, J.-H. Cho, and I.-R. Chen, "Social media-based social–psychological community resilience analysis of five countries on covid-19," *Journal of Computational Social Science*, vol. 6, pp. 1001–1032, 8 2023.

[9] A. B. Chapman, D. O. Scharfstein, A. E. Montgomery, T. Byrne, Y. Suo, A. Effiong, T. Velasquez, W. Pettey, and R. E. Nelson, "Using natural language processing to study homelessness longitudinally with electronic health record data subject to irregular observations.," *AMIA ... Annual Symposium proceedings. AMIA Symposium*, vol. 2023, pp. 894–, 1 2024.

[10] C. Xiang and M. Abouelyazid, "The impact of generational cohorts and visit environment on telemedicine satisfaction: A novel investigation," 2020.

[11] Y. Cong, A. N. LaCroix, and J. Lee, "Clinical efficacy of pre-trained large language models through the lens of aphasia.," *Scientific reports*, vol. 14, pp. 15573–, 7 2024.

[12] A. Sharma, M. Witbrock, and K. Goolsbey, "Controlling search in very large commonsense knowledge bases: a machine learning approach," *arXiv preprint arXiv:1603.04402*, 2016.

[13] F. Dipaola, M. A. Gebska, M. Gatti, A. G. Levra, W. H. Parker, R. Menè, S. Lee, G. Costantino, E. J. Barsotti, D. Shiffer, S. L. Johnston, R. Sutton, B. Olshansky, and R. Furlan, "Will artificial intelligence be "better" than humans in the management of syncope?," *JACC. Advances*, vol. 3, pp. 101072–101072, 7 2024.

[14] V. V. Kumar, A. Sahoo, R. Kumar, and N. Loyd, "Public healthcare informatics for covid-19 from social media data.," *Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual International Conference*, vol. 2024, pp. 1–4, 7 2024.

[15] S. Kondra, W. Xu, and V. V. Raghavan, "Unlocking the power of clinical notes: Natural language processing in healthcare," *Acta Scientific Medical Sciences*, pp. 37–44, 6 2024.

[16] N. Powell, M. Blank, A. Luintel, S. Elkhalifa, R. Bhogal, M. Wilcock, M. Wakefield, and J. Sandoe, "Narrative review of recent developments and the future of penicillin allergy de-labelling by non-allergists.," *npj antimicrobials and resistance*, vol. 2, pp. 18–, 7 2024.

[17] R. Mylvaganam, R. Lawrence, I. Goldberg, F. Rahaghi, S. Chiu, S. C. Malaisrie, D. Schimmel, R. Avery, K. Martin, and M. J. Cuttica, "Differences in referral to a chronic thromboembolic pulmonary hypertension center following acute pulmonary embolism: a locoregional experience.," *Journal of thrombosis and thrombolysis*, vol. 55, pp. 691–699, 2 2023.

[18] H. Chen, E. Cohen, D. Wilson, and M. Alfted, "Improving patient safety event report classification with machine learning and contextual text representation.," *Proceedings of the Human Factors and Ergonomics Society ... Annual Meeting. Human Factors and Ergonomics Society. Annual meeting*, vol. 67, pp. 1063–1069, 10 2023.

[19] L. L. Guo, E. Steinberg, S. L. Fleming, J. Posada, J. Lemmon, S. R. Pfohl, N. Shah, J. Fries, and L. Sung, "Ehr foundation models improve robustness in the presence of temporal distribution shift.," *Scientific reports*, vol. 13, pp. 3767–, 3 2023.

[20] D. Kumar and N. P. Bassill, "Analysing trends of computational urban science and data science approaches for sustainable development," *Computational Urban Science*, vol. 4, 11 2024.

[21] T. Wambsganss, I. Benke, A. Maedche, K. Koedinger, and T. Käser, “Evaluating the impact of learner control and interactivity in conversational tutoring systems for persuasive writing,” *International Journal of Artificial Intelligence in Education*, 5 2024.

[22] Y. Dang, F. Li, X. Hu, V. K. Keloth, M. Zhang, S. Fu, M. F. Amith, J. W. Fan, J. Du, E. Yu, H. Liu, X. Jiang, H. Xu, and C. Tao, “Systematic design and data-driven evaluation of social determinants of health ontology (sdoho).,” *Journal of the American Medical Informatics Association : JAMIA*, vol. 30, pp. 1465–1473, 6 2023.

[23] V. E. Nambudiri, “Decoding the abcs of natural language processing in dermatology.,” *Journal of the European Academy of Dermatology and Venereology : JEADV*, vol. 38, pp. 2201–2202, 11 2024.

[24] K. Ding, M. Chetty, A. N. Hoshyar, T. Bhattacharya, and B. Klein, “Speech based detection of alzheimer’s disease: a survey of ai techniques, datasets and challenges,” *Artificial Intelligence Review*, vol. 57, 10 2024.

[25] W. Liang, N. Rajani, X. Yang, E. Ozoani, E. Wu, Y. Chen, D. S. Smith, and J. Zou, “Systematic analysis of 32,111 ai model cards characterizes documentation practice in ai,” *Nature Machine Intelligence*, vol. 6, pp. 744–753, 6 2024.

[26] S. N. El-Tallawy, J. V. Pergolizzi, I. Vasiliu-Feltes, R. S. Ahmed, J. K. LeQuang, H. N. El-Tallawy, G. Varrassi, and M. S. Nagiub, “Incorporation of "artificial intelligence" for objective pain assessment: A comprehensive review.,” *Pain and therapy*, vol. 13, pp. 293–317, 3 2024.

[27] M. T. Mefford, A. P. Ambrosy, R. Wei, C. Zheng, R. V. Parikh, T. N. Harrison, M.-S. Lee, A. S. Go, and K. Reynolds, “Rule-based natural language processing to examine variation in worsening heart failure hospitalizations by age, sex, race and ethnicity, and left ventricular ejection fraction.,” *American heart journal*, vol. 278, pp. 117–126, 9 2024.

[28] T. T. Khoei and A. Singh, “A survey of emotional artificial intelligence and crimes: detection, prediction, challenges and future direction,” *Journal of Computational Social Science*, vol. 7, pp. 2359–2402, 7 2024.

[29] H. A. L. Rocha, E. Z. M. Solha, V. Furtado, F. L. Justino, L. A. L. Barreto, R. G. da Silva, Ítalo Martins de Oliveira, D. W. Bates, L. P. de Góes Cavalcanti, A. S. L. Neto, and E. A. de Oliveira, “Covid-19 outbreaks surveillance through text mining applied to electronic health records.,” *BMC infectious diseases*, vol. 24, pp. 359–, 3 2024.

[30] T. E. Workman, A. Ahmed, H. M. Sheriff, V. K. Raman, S. Zhang, Y. Shao, C. Faselis, G. C. Fonarow, and Q. Zeng-Treitler, “Chatgpt-4 extraction of heart failure symptoms and signs from electronic health records.,” *Progress in cardiovascular diseases*, vol. 87, pp. 44–49, 10 2024.

[31] S. Kayikci and T. M. Khoshgoftaar, “Blockchain meets machine learning: a survey,” *Journal of Big Data*, vol. 11, 1 2024.

[32] Y. Yang, X. Liu, Q. Jin, F. Huang, and Z. Lu, “Unmasking and quantifying racial bias of large language models in medical report generation.,” *Communications medicine*, vol. 4, pp. 176–, 9 2024.

[33] A. K. Saxena, R. R. Dixit, and A. Aman-Ullah, “An lstm neural network approach to resource allocation in hospital management systems,” *International Journal of Applied Health Care Analytics*, vol. 7, no. 2, pp. 1–12, 2022.

[34] C. Liman, J. Schein, A. Wu, X. Huang, S. Thadani, A. Childress, S. H. Kollins, and S. Bhattacharjee, “Real world analysis of treatment change and response in adults with attention-deficit/hyperactivity disorder (adhd) alone and with concomitant psychiatric comorbidities: results from an electronic health record database study is the united states.,” *BMC psychiatry*, vol. 24, pp. 618–, 9 2024.

[35] A. A. T. Monfared, Y. Stern, S. Doogan, M. Irizarry, and Q. Zhang, “Understanding barriers along the patient journey in alzheimer’s disease using social media data.,” *Neurology and therapy*, vol. 12, pp. 899–918, 4 2023.

[36] Y. Sun, L. Liu, Y. Xu, X. Zeng, Y. Shi, H. Hu, J. Jiang, and A. Abraham, “Alternative data in finance and business: emerging applications and theory analysis (review),” *Financial Innovation*, vol. 10, 12 2024.

[37] J. He, F. Li, J. Li, X. Hu, Y. Nian, Y. Xiang, J. Wang, Q. Wei, Y. Li, H. Xu, and C. Tao, “Prompt tuning in biomedical relation extraction.,” *Journal of healthcare informatics research*, vol. 8, pp. 206–224, 2 2024.

[38] J. G. Stenzel, N. R. Schultz, and M. J. Marino, “Automated classification of online reviews of otolaryngologists.,” *Laryngoscope investigative otolaryngology*, vol. 9, pp. e70036–, 11 2024.

[39] J. R. Vest, R. J. Hinrichs, and H. Hosler, “How legal problems are conceptualized and measured in healthcare settings: a systematic review.,” *Health & justice*, vol. 11, pp. 48–, 11 2023.

[40] N. Kedia, S. Sanjeev, J. Ong, and J. Chhablani, "Chatgpt and beyond: An overview of the growing field of large language models and their use in ophthalmology.," *Eye (London, England)*, vol. 38, pp. 1252–1261, 1 2024.

[41] S. J. Rao, A. Isath, P. Krishnan, J. A. Tangsrivimol, H. U. H. Virk, Z. Wang, B. S. Glicksberg, and C. Krittanawong, "Chatgpt: A conceptual review of applications and utility in the field of medicine.," *Journal of medical systems*, vol. 48, pp. 59–, 6 2024.

[42] I. A. Hashem, A. Siddiqa, F. A. Alaba, M. Bilal, and S. M. Alhashmi, "Distributed intelligence for iot-based smart cities: a survey," *Neural Computing and Applications*, vol. 36, pp. 16621–16656, 7 2024.

[43] S. Gan, C. Kim, J. Chang, D. Y. Lee, and R. W. Park, "Enhancing readmission prediction models by integrating insights from home healthcare notes: Retrospective cohort study.," *International journal of nursing studies*, vol. 158, pp. 104850–104850, 7 2024.

[44] A. Lois, R. Yates, M. Ivy, C. Inaba, R. Tatum, L. Cetrulo, Z. Parr, J. Chen, S. Khandelwal, and A. Wright, "Accuracy of natural language processors for patients seeking inguinal hernia information.," *Surgical endoscopy*, vol. 38, pp. 7409–7415, 10 2024.

[45] N. Obradovich, S. S. Khalsa, W. Khan, J. Suh, R. H. Perlis, O. Ajilore, and M. P. Paulus, "Opportunities and risks of large language models in psychiatry.," *NPP - digital psychiatry and neuroscience*, vol. 2, 5 2024.

[46] R. Ibrahim, H. N. Pham, S. Ganatra, Z. Javed, K. Nasir, and S. Al-Kindi, "Social phenotyping for cardiovascular risk stratification in electronic health registries.," *Current atherosclerosis reports*, vol. 26, pp. 485–497, 7 2024.

[47] J. R. Machireddy, "Automation in healthcare claims processing: Enhancing efficiency and accuracy," *International Journal of Science and Research Archive*, vol. 09, no. 01, pp. 825–834, 2023.

[48] B. Kösoglu-Kind, R. Loredo, M. Grossi, C. Bernecker, J. M. Burks, and R. Buchkremer, "A biological sequence comparison algorithm using quantum computers.," *Scientific reports*, vol. 13, pp. 14552–, 9 2023.

[49] M. Wornow, Y. Xu, R. Thapa, B. Patel, E. Steinberg, S. Fleming, M. A. Pfeffer, J. Fries, and N. H. Shah, "The shaky foundations of large language models and foundation models for electronic health records.," *NPJ digital medicine*, vol. 6, pp. 135–, 7 2023.

[50] N. Tavabi, M. Singh, J. Pruneski, and A. M. Kiapour, "Systematic evaluation of common natural language processing techniques to codify clinical notes.," *PloS one*, vol. 19, pp. e0298892–e0298892, 3 2024.

[51] M. Muniswamaiah, T. Agerwala, and C. C. Tappert, "Context-aware query performance optimization for big data analytics in healthcare," in *2019 IEEE High Performance Extreme Computing Conference (HPEC-2019)*, pp. 1–7, 2019.

[52] J. M. Plasek, M. G. Amato, A. Salem, D. Foer, S. Lipsitz, G. P. Jackson, D. W. Bates, and L. Zhou, "Adverse drug events in ambulatory care: A cross-sectional study.," *Drug safety*, vol. 48, pp. 363–374, 12 2024.

[53] M. Muniswamaiah, T. Agerwala, and C. C. Tappert, "Federated query processing for big data in data science," in *2019 IEEE International Conference on Big Data (Big Data)*, pp. 6145–6147, IEEE, 2019.

[54] M. C. Sabaner, R. Anguita, F. Antaki, M. Balas, L. C. Boberg-Ans, L. F. Desideri, J. Grauslund, M. S. Hansen, O. N. Klefter, I. Potapenko, M. L. R. Rasmussen, and Y. Subhi, "Opportunities and challenges of chatbots in ophthalmology: A narrative review.," *Journal of personalized medicine*, vol. 14, pp. 1165–1165, 12 2024.

[55] S.-F. Chen, S. Loguercio, K.-Y. Chen, S. E. Lee, J.-B. Park, S. Liu, H. J. Sadaei, and A. Torkamani, "Artificial intelligence for risk assessment on primary prevention of coronary artery disease," *Current Cardiovascular Risk Reports*, vol. 17, pp. 215–231, 11 2023.

[56] G. Echefu, L. Batalik, A. Lukan, R. Shah, P. Nain, A. Guha, and S.-A. Brown, "The digital revolution in medicine: Applications in cardio-oncology.," *Current treatment options in cardiovascular medicine*, vol. 27, 11 2024.

[57] K. L. Kehl, J. Jee, K. Pichotta, M. A. Paul, P. Trukhanov, C. Fong, M. Waters, Z. Bakouny, W. Xu, T. K. Choueiri, C. Nichols, D. Schrag, and N. Schultz, "Shareable artificial intelligence to extract cancer outcomes from electronic health records for precision oncology research.," *Nature communications*, vol. 15, pp. 9787–, 11 2024.

[58] O. J. B. D. Walk, A. Pichon, H. R. Nieva, T. Sun, J. Li, J. Joseph, S. Kinberg, L. R. Richter, S. Crusco, K. Kulas, S. A. Ahmed, D. Snyder, A. Rahbari, B. L. Ranard, P. Juneja, D. Demner-Fushman, and N. Elhadad, "Contextualized race and ethnicity annotations for clinical text from mimic-iii.," *Scientific data*, vol. 11, pp. 1332–, 12 2024.

[59] M. Abouelyazid and C. Xiang, "Machine learning-assisted approach for fetal health status prediction using cardiotocogram data," *International Journal of Applied Health Care Analytics*, vol. 6, no. 4, pp. 1–22, 2021.

[60] Y. Liu, H. Wang, H. Zhou, M. Li, Y. Hou, S. Zhou, F. Wang, R. Hoetzlein, and R. Zhang, "A review of reinforcement learning for natural language processing and applications in healthcare.," *Journal of the American Medical Informatics Association : JAMIA*, vol. 31, pp. 2379–2393, 8 2024.

[61] A. Sharma, K. M. Goolsbey, and D. Schneider, "Disambiguation for semi-supervised extraction of complex relations in large commonsense knowledge bases," in *7th Annual Conference on Advances in Cognitive Systems*, 2019.

[62] M. P. Girouard, A. J. Chang, Y. Liang, S. A. Hamilton, A. S. Bhatt, J. Svetlichnaya, J. K. Fitzpatrick, E. C. B. Carey, H. R. Avula, S. Adatya, K. K. Lee, M. D. Solomon, R. V. Parikh, A. S. Go, and A. P. Ambrosy, "Clinical and research applications of natural language processing for heart failure.," *Heart failure reviews*, vol. 30, pp. 407–415, 12 2024.

[63] H. Xu, S. M. Sudheendra, J. Huh, A. Salecha, and J. Srivastava, "Influence of emotions on coping behaviors in crisis: a computational analysis of the covid-19 outbreak," *Journal of Computational Social Science*, vol. 7, pp. 1599–1623, 5 2024.

[64] B. Jeon and S. Chae, "0906 free-text documentation of sleep disturbance in acute myeloid leukemia: A natural language processing study," *SLEEP*, vol. 47, pp. A389–A389, 4 2024.

[65] C. Yuan and S. S. Agaian, "A comprehensive review of binary neural network," *Artificial Intelligence Review*, vol. 56, pp. 12949–13013, 3 2023.

[66] H.-T. Tsai, J. Wu, P. Gupta, E. R. Heinz, and A. Jafari, "Predicting blood transfusions for coronary artery bypass graft patients using deep neural networks and synthetic data," *Neural Computing and Applications*, vol. 36, pp. 21153–21162, 8 2024.

[67] S. A. Pahune, "How does ai help in rural development in healthcare domain: A short survey," *International Journal for Research in Applied Science and Engineering Technology*, vol. 11, pp. 4184–4191, 6 2023.

[68] W. Zhong, P. Y. Yao, S. H. Boppana, F. V. Pacheco, B. S. Alexander, S. Simpson, and R. A. Gabriel, "Improving case duration accuracy of orthopedic surgery using bidirectional encoder representations from transformers (bert) on radiology reports.," *Journal of clinical monitoring and computing*, vol. 38, pp. 221–228, 9 2023.

[69] J. D. Gaizo, C. Sherard, K. Shorbaji, B. Welch, R. Mathi, and A. Kilic, "Prediction of coronary artery bypass graft outcomes using a single surgical note: An artificial intelligence-based prediction model study.," *PloS one*, vol. 19, pp. e0300796–e0300796, 4 2024.

[70] M. L. Sbodio, V. López, T. L. Hoang, T. Brisimi, G. Picco, I. Vejsbjerg, V. Rho, P. M. Aonghusa, M. Kristiansen, and J. Segrave-Daly, "Collaborative artificial intelligence system for investigation of healthcare claims compliance.," *Scientific reports*, vol. 14, pp. 11884–, 5 2024.

[71] J. R. Machireddy, "Machine learning and automation in healthcare claims processing," *Journal of Artificial Intelligence General science (JAIGS) ISSN: 3006-4023*, vol. 6, no. 1, pp. 686–701, 2024.

[72] M. A. L. Khaniki, M. Mirzaeibonehkhater, M. Manthouri, and E. Hasani, "Vision transformer with feature calibration and selective cross-attention for brain tumor classification," *Iran Journal of Computer Science*, 12 2024.

[73] T. Akan, S. Alp, M. S. Bhuiyan, T. Helmy, A. W. Orr, M. M. R. Bhuiyan, S. A. Conrad, J. A. Vanchiere, C. G. Kevil, and M. A. N. Bhuiyan, "Viviechoformer: Deep video regressor predicting ejection fraction.," *Journal of imaging informatics in medicine*, 11 2024.

[74] M. Wörsdörfer, "Biden's executive order on ai: strengths, weaknesses, and possible reform steps," *AI and Ethics*, 7 2024.

[75] L. Chen, Z. Fan, J. Chang, R. Yang, H. Hou, H. Guo, Y. Zhang, T. Yang, C. Zhou, Q. Sui, Z. Chen, C. Zheng, X. Hao, K. Zhang, R. Cui, Z. Zhang, H. Ma, Y. Ding, N. Zhang, X. Lu, X. Luo, H. Jiang, S. Zhang, and M. Zheng, "Sequence-based drug design as a concept in computational drug design.," *Nature communications*, vol. 14, pp. 4217–, 7 2023.

[76] K. D. Forbus, K. Lockwood, A. B. Sharma, and E. Tomai, "Steps towards a second generation learning by reading system.," in *AAAI Spring Symposium: Learning by Reading and Learning to Read*, pp. 36–43, 2009.

[77] P. Bota, G. Thambiraj, S. C. Bollepalli, and A. A. Armoundas, "Artificial intelligence algorithms in cardiovascular medicine: An attainable promise to improve patient outcomes or an inaccessible investment?," *Current cardiology reports*, vol. 26, pp. 1477–1485, 10 2024.

[78] G. Holste, M. Lin, R. Zhou, F. Wang, L. Liu, Q. Yan, S. H. V. Tassel, K. Kovacs, E. Y. Chew, Z. Lu, Z. Wang, and Y. Peng, "Harnessing the power of longitudinal medical imaging for eye disease prognosis using transformer-based sequence modeling.," *NPJ digital medicine*, vol. 7, pp. 216–, 8 2024.

[79] J. A. Rodger and J. Piper, "Assessing american presidential candidates using principles of ontological engineering, word sense disambiguation, data envelope analysis and qualitative comparative analysis," *International Journal of Speech Technology*, vol. 26, pp. 743–764, 10 2023.

[80] A. C. Brittain and J. M. Carrington, "Effective communication of system-level events for hospital system health and nurse well-being: A qualitative study," *Safety*, vol. 10, pp. 96–96, 11 2024.

[81] M. H. Murad, T. Nayfeh, S. Saadi, M. H. Farah, J. W. Njeru, K. E. Viola, C. Y. Y. Chen, L. J. Prokop, and Z. Wang, "Home and community-based services: A systematic review and evidence map.," *Journal of general internal medicine*, vol. 39, pp. 3315–3321, 8 2024.

[82] V. Berisha and J. M. Liss, "Responsible development of clinical speech ai: Bridging the gap between clinical research and technology.," *NPJ digital medicine*, vol. 7, pp. 208–, 8 2024.

[83] V. Bali, V. Turzhitsky, J. Schelfhout, M. Paudel, E. Hulbert, J. Peterson-Brandt, J. Hertzberg, N. R. Kelly, and R. H. Patel, "Machine learning to identify chronic cough from administrative claims data.," *Scientific reports*, vol. 14, pp. 2449–, 1 2024.

[84] A. Vaid, J. Jiang, A. Sawant, S. Lerakis, E. Argulian, Y. Ahuja, J. Lampert, A. Charney, H. Greenspan, J. Narula, B. Glicksberg, and G. N. Nadkarni, "A foundational vision transformer improves diagnostic performance for electrocardiograms.," *NPJ digital medicine*, vol. 6, pp. 108–, 6 2023.

[85] F. J. Griffith, G. I. Ash, M. Augustine, L. Latimer, N. Verne, N. S. Redeker, S. S. O'Malley, K. S. DeMartini, and L. M. Fucito, "Natural language processing in mixed-methods evaluation of a digital sleep-alcohol intervention for young adults.," *NPJ digital medicine*, vol. 7, pp. 342–, 11 2024.

[86] P. M. Heider and S. M. Meystre, "An extensible evaluation framework applied to clinical text deidentification natural language processing tools: Multisystem and multicorpus study.," *Journal of medical Internet research*, vol. 26, pp. e55676–e55676, 5 2024.

[87] A. Sharma, *Structural and network-based methods for knowledge-based systems*. PhD thesis, Northwestern University, 2011.

[88] S. F. Ahmed, M. S. B. Alam, M. Hassan, M. R. Rozbu, T. Ishtiak, N. Rafa, M. Mofijur, A. B. M. S. Ali, and A. H. Gandomi, "Deep learning modelling techniques: current progress, applications, advantages, and challenges," *Artificial Intelligence Review*, vol. 56, pp. 13521–13617, 4 2023.

[89] C. G. Arnold, B. Sonn, F. J. Meyers, A. Vest, R. Puls, E. Zirkler, M. Edelmann, I. M. Brooks, and A. A. Monte, "Accessing and utilizing clinical and genomic data from an electronic health record data warehouse.," *Translational medicine communications*, vol. 8, 3 2023.

[90] J. T.-Y. Wu, J. Corrigan, C. Su, C. Dumontier, J. La, A. Khan, S. Arya, A. H. S. Harris, L. Backhus, M. Das, N. V. Do, M. T. Brophy, S. S. Han, M. Kelley, and N. R. Fillmore, "The performance status gap in immunotherapy for frail patients with advanced non-small cell lung cancer.," *Cancer immunology, immunotherapy : CII*, vol. 73, pp. 172–, 7 2024.

[91] R. Wistrom, L. Khait, and G. Nelson, "Artificial intelligence in point-of-care ultrasound," *Current Emergency and Hospital Medicine Reports*, vol. 12, pp. 89–94, 5 2024.

[92] S. Pan, G. T. Goodnight, X. Zhao, Y. Wang, L. Xie, and J. Zhang, ""game changer": the ai advocacy discourse of 2023 in the us," *AI & SOCIETY*, 8 2024.

[93] V. Barcelona, D. Scharp, H. Moen, A. Davoudi, B. R. Idnay, K. Cato, and M. Topaz, "Using natural language processing to identify stigmatizing language in labor and birth clinical notes.," *Maternal and child health journal*, vol. 28, pp. 578–586, 12 2023.

[94] M. Muniswamaiah, T. Agerwala, and C. Tappert, "Big data in cloud computing review and opportunities," *arXiv preprint arXiv:1912.10821*, 2019.

[95] V. Sorin, B. S. Glicksberg, Y. Artsi, Y. Barash, E. Konen, G. N. Nadkarni, and E. Klang, "Utilizing large language models in breast cancer management: systematic review.," *Journal of cancer research and clinical oncology*, vol. 150, pp. 140–, 3 2024.

[96] M. L. Sung, C. León, J. I. Reisman, K. S. Gordon, R. D. Kerns, W. Li, W. Liu, A. Mitra, H. Yu, and W. C. Becker, "Disparities in receipt of medications for opioid use disorder before and during the covid-19 pandemic in the us veterans health administration.," *Substance use & addiction journal*, vol. 46, pp. 369–376, 11 2024.

[97] Z. Wang, Z. Chu, T. V. Doan, S. Ni, M. Yang, and W. Zhang, "History, development, and principles of large language models: an introductory survey," *AI and Ethics*, 10 2024.

[98] M. Muniswamaiah, T. Agerwala, and C. C. Tappert, "Federated query processing for big data in data science," in *2019 IEEE International Conference on Big Data (Big Data)*, pp. 6145–6147, IEEE, 2019.

[99] Y.-H. Sheu, C. Magdamo, M. Miller, S. Das, D. Blacker, and J. W. Smoller, "Ai-assisted prediction of differential response to antidepressant classes using electronic health records.," *NPJ digital medicine*, vol. 6, pp. 73–, 4 2023.

[100] E. Fröling, N. Rajaeean, K. S. Hinrichsmeyer, D. Domrös-Zoungrana, J. N. Urban, and C. Lenz, "Artificial intelligence in medical affairs: A new paradigm with novel opportunities.," *Pharmaceutical medicine*, vol. 38, pp. 331–342, 9 2024.

[101] M. Malgaroli, T. D. Hull, J. M. Zech, and T. Althoff, "Natural language processing for mental health interventions: a systematic review and research framework.," *Translational psychiatry*, vol. 13, pp. 309–, 10 2023.

[102] A. H. Zhao, D. I. Glazer, M. M. Hammer, K. S. Burk, P. J. DiPiro, and R. Khorasani, "Comparing thoracic and abdominal subspecialists' follow-up recommendations for abdominal findings identified on chest ct.," *Abdominal radiology (New York)*, vol. 48, pp. 1468–1478, 2 2023.