CLASSICALLIBRARY : Pages:1–13

Original Research

Information Extraction from Provider Notes for Streamlined Medical Billing Using Weak Supervision

Nguyn Minh Khoa[1]

[1]Đi hc Công ngh Thông tin và Truyn thông Thái Nguyên, Khoa Khoa hc Máy tính, Đng Z115, TP. Thái Nguyên, Vit Nam.

Abstract
Information extraction from provider notes has emerged as an instrumental approach to optimizing medical billing processes by capturing relevant clinical and administrative details to reduce claim denials and billing inefficiencies. The presence of domain-specific language, abbreviations, and frequent variations in provider documentation calls for specialized strategies that can automatically identify, categorize, and validate key clinical entities. Traditional methods in supervised learning rely on extensive manually annotated datasets to capture complex linguistic nuances, resulting in time-consuming and resource-intensive data preparation stages. By contrast, weak supervision offers the potential to harness automatically generated labeling functions, expert knowledge bases, and rule-based heuristics to train models in a more cost-effective manner. This paper discusses a framework that integrates weak supervision with advanced natural language processing techniques, aiming to adaptively handle unstructured and semi-structured medical text for robust entity recognition and accurate downstream billing code assignment. The approach involves logic-based constraints for label reconciliation and probabilistic inference to account for label noise. Through this strategy, refined entity resolution is achieved, thus streamlining the billing pipeline by enabling automatic validation of provider notes and real-time alerts for coding inconsistencies. Empirical results indicate that combining weak supervision with context-sensitive embeddings can significantly reduce the burden on human annotators while preserving high levels of precision and recall in capturing relevant clinical descriptors. The ensuing discussion delves into the mathematical formulation, linguistic representation, and real-world impact of these methodologies.

1. Introduction

Healthcare systems worldwide confront persistent challenges in achieving an efficient medical billing workflow, especially in the face of heterogeneous documentation styles and evolving regulatory standards [1]. Provider notes, typically unstructured textual artifacts, embody a wealth of domain-specific knowledge that coders and billing systems must distill into structured data for invoicing and reimbursement processes [2]. The mismatch between how medical professionals record clinical encounters and how insurers process claims leads to substantial administrative overhead, with potential delays in payment and an increase in denied or reworked claims. As a result, both providers and payers have strong incentives to improve the accuracy and speed of extracting pertinent information from clinical documentation. [3]

Information extraction techniques offer mechanisms for automated parsing of textual data, identifying key entities such as diagnoses, treatments, medications, laboratory results, and procedures. These extracted elements are subsequently mapped to relevant billing codes, such as Current Procedural Terminology codes or International Classification of Diseases codes, which are then submitted for reimbursement [4]. However, conventional supervised learning approaches often rely on carefully labeled training data, requiring domain experts, certified coders, or medical professionals to annotate large corpora of notes. This annotation process is prohibitively expensive and time-consuming, given the immense variability of medical documentation [5]. Consequently, the pursuit of more data-efficient methods has led researchers to adopt weak supervision, a paradigm wherein imperfect sources of labels are combined to generate sufficiently accurate supervision signals for model training.

Weak supervision leverages heuristic rules, distant supervision through knowledge bases, and other automatically curated signals to overcome the scarcity of high-quality annotations [6]. A robust integration of these noisy labeling functions in a well-structured probabilistic framework or logic-driven environment enables the consolidation of multiple, often conflicting, signals into a single unified label per instance. This approach allows large-scale model training with substantially reduced human effort [7]. Within the context of extracting billing-relevant information, weak supervision has particular value because it can incorporate codified domain expertise, standardized guidelines, and even external knowledge repositories of medical concepts and synonyms [8]. This synergy generates a training environment wherein domain-specific constraints enhance model reliability, even when the initial labeling functions are incomplete or error-prone.

The practical importance of high-fidelity information extraction from provider notes extends beyond mere cost reduction [9]. Accurate extraction underpins the generation of clinically coherent narratives and ensures that downstream analytics, such as risk stratification or quality measurement, rely on valid data. Consider a scenario in which a physician's note ambiguously references a procedure, and a naive model incorrectly assigns the corresponding billing code, potentially leading to claim denial or compliance concerns [10, 11]. By focusing on advanced natural language processing architectures, such as transformer-based embeddings that preserve contextual information and capture domain nuances, these ambiguities can be mitigated. Moreover, implementing logical constraints on recognized entities helps ensure consistent alignment across multiple sections of a single note, such as progress notes, lab results, and plan sections, thereby reinforcing the final coding decision. [12]

In discussing weak supervision for information extraction, we must also address foundational concepts of representation. Provider notes can be encoded at the token level or phrase level, each with distinct implications for capturing medical concepts [13]. A suitably chosen vector representation of tokens or phrases can encode both syntactic and semantic attributes, enabling classification layers to differentiate between clinically distinct entities. Matrix factorization or dimensionality reduction techniques can also help capture latent relationships among tokens, ensuring that frequently co-occurring terms in clinical documentation map to conceptually coherent spaces [14]. However, these transformations alone do not address the question of how to handle noisy supervision signals effectively [15]. Hence, logic-based constraint systems and generative labeling models play a pivotal role in reconciling multiple weak labels, forming the backbone of a robust information extraction pipeline.

The subsequent sections explore a rigorous mathematical framework for modeling weak supervision in the context of medical note analysis, presenting a logic-oriented system for label reconciliation and advanced classification [16]. Specialized matrix operators, symbolic definitions of constraints, and linear algebraic notations characterize the approach. We discuss the manner in which large corpora of provider notes are transformed into structured inputs via a pipeline that ingests unstructured text, applies lexical and domain-specific heuristics, and integrates knowledge from medical ontologies [17]. This pipeline then employs a probabilistic or rule-based aggregator of weak labels, eventually passing refined supervision signals to a high-capacity sequence labeling model. The core methodology focuses on reconciling conflicting labels and leveraging each labeler's unique expertise [18]. In parallel, an evaluation of performance metrics such as precision, recall, and F1-score illuminates the effectiveness of this integrated approach.

The field of healthcare analytics necessitates rigorous compliance with privacy guidelines, such as the Health Insurance Portability and Accountability Act (HIPAA) in the United States [19]. In building automated pipelines for medical billing, data processing must ensure de-identification, secure storage, and appropriate use of sensitive information. Therefore, the discussion includes not only the technical aspects of model design and implementation but also addresses how these measures can be shaped to adhere to relevant privacy standards without undermining the predictive and interpretive power of the system [20]. The integrated framework thus reconciles domain knowledge, data security, advanced machine learning, and weak supervision, offering a compelling vision for efficient and accurate medical billing workflows. [21]

2. Mathematical Foundations for Weakly Supervised Entity Recognition

Weak supervision for information extraction may be conceptualized through a set of labeling functions that imperfectly annotate tokens or segments of the provider note. Let $D = \{x_1, x_2, \dots, x_n\}$ be the corpus of provider notes, where each x_i is a sequence of tokens $x_i = (w_{i1}, w_{i2}, \dots, w_{im_i})$. Define a set of labeling functions $\Lambda = \{\lambda_1, \lambda_2, \dots, \lambda_k\}$, where each $\lambda_j(x_i)$ yields a proposed label vector for the tokens in x_i. Because labeling functions may be noisy, the outputs can be contradictory. [22]

We introduce a generative model $P(Z, Y \mid \Lambda, D)$, where Z are latent variables representing the true entity labels for each token in each note, and Y are the observed noisy labels produced by the labeling functions. Under typical assumptions, one posits that for each token $t \in x_i$, the probability of a labeling function λ_j producing label ℓ is conditioned on the true latent label z_t [23]. Symbolically, one might write:

$$P(\lambda_j(x_i) = \ell \mid z_t) = \alpha_{j,\ell}^{(z_t)},$$

where $\alpha_{j,\ell}^{(z_t)}$ is a parameter capturing the reliability of labeling function j when the true label is z_t. A common inference goal is to estimate $\hat{z}_t = \operatorname{argmax}_{z_t} P(z_t \mid \Lambda, D)$, yielding an aggregate weak label. This step often involves maximum likelihood estimation over generative parameters $\alpha_{j,\ell}^{(z_t)}$ or employs Bayesian approaches with corresponding priors.

Let there be C possible entity classes [24]. Suppose each token $w_{i,r}$ belongs to one of these classes $\{1, 2, \dots, C\}$. We define a set of logic constraints to reflect domain-specific knowledge. For instance, if a token is identified as a procedure code in one part of the note, it cannot simultaneously be identified as a diagnosis code in an adjacent segment [25]. Formally, let p denote a proposition that a certain token belongs to a given class, and let q denote a proposition that the same token belongs to a conflicting class. The constraint might be encoded as: [26]

$$\forall w_{i,r} \; [p \land q \implies \bot],$$

where $\bot$ is an unsatisfiable logical statement [27]. Additional constraints can incorporate sequential consistency or anaphoric references that tie particular clinical concepts together across the note, enforced as:

$$\forall w_{i,r}, w_{i,s} \; [(\text{Anaphora}(w_{i,r}, w_{i,s})) \land (p(w_{i,r})) \implies p(w_{i,s})].$$

Such constraints help mitigate label noise by leveraging structural properties of clinical documentation. [28]

Once the aggregated weak labels have been obtained or refined, a sequence labeling model can be trained. Let Θ be the parameter set of a sequence labeling network, such as a bidirectional long short-term memory or a transformer-based architecture [29]. The goal is to learn:

$$\Theta^* = \operatorname{argmin}_{\Theta} \sum_{i=1}^{n} \sum_{r=1}^{m_i} \mathcal{L}\Big(h_\Theta(w_{i,r}), \hat{z}_{i,r}\Big),$$

where $h_\Theta(w_{i,r})$ is the model's predicted distribution for token $w_{i,r}$, and $\hat{z}_{i,r}$ is the weak label assigned through the generative model and logic constraints. The function $\mathcal{L}$ could be a cross-entropy loss, or a specialized metric incorporating label uncertainties. The synergy of generative label aggregation, logic constraints, and a high-capacity discriminative model underlies the effectiveness of weakly supervised pipelines in capturing clinically relevant information from provider notes. [30]

3. Representation of Provider Notes and Lexical Structures

The nature of medical text requires specialized representation methods to account for domain-specific terminology, abbreviations, synonyms, and context dependence. One approach is to embed each token or phrase in a vector space designed to capture semantic similarities within clinical contexts [31]. A typical pipeline might begin by mapping each token $w_{i,r}$ to an embedding vector $\mathbf{e}_{i,r} \in \mathbb{R}^d$. Traditional approaches rely on static embeddings, such as word2vec, but modern paradigms favor context-sensitive embeddings that incorporate the entire sentence or note structure.

We define a context-based embedding function $f : (w_{i,r}, x_i) \mapsto \mathbf{e}_{i,r}$, where $\mathbf{e}_{i,r}$ depends on both the token itself and its surrounding context in x_i. Transformer models, employing self-attention mechanisms, excel at computing such context-adaptive embeddings [32]. Denote by $\mathbf{E}_i \in \mathbb{R}^{m_i \times d}$ the matrix whose rows are $\mathbf{e}_{i,1}, \mathbf{e}_{i,2}, \ldots, \mathbf{e}_{i,m_i}$. We can subsequently apply matrix factorization or dimensionality reduction methods to discover latent structures within these embeddings [33]. For example, consider a singular value decomposition:

$$\mathbf{E}_i = \mathbf{U}_i \boldsymbol{\Sigma}_i \mathbf{V}_i^\top,$$

where $\mathbf{U}_i$ and $\mathbf{V}_i$ capture orthonormal bases for tokens and feature dimensions, and $\boldsymbol{\Sigma}_i$ is diagonal with singular values. By examining these latent factors, we can expose relationships across terms, such that clinically related tokens appear closely aligned in the factor space. [34]

Symbolic transformations of the text often prove beneficial, especially when dealing with recognized clinical entities. Suppose an ontology-based tool identifies a mention of a drug name or procedure code [35]. This token can be replaced or augmented with standardized nomenclature from a coding system. Let O be a medical ontology defining synonyms, parent-child relationships, and standardized codes. When encountering a token $w_{i,r}$ that matches an entry in O, an augmented representation $\mathbf{e}'_{i,r}$ might be formed by concatenating the original embedding $\mathbf{e}_{i,r}$ with a structured representation from the ontology. For instance, if the token matches a concept with ID u in the ontology, we define a one-hot or multi-hot vector $\mathbf{o}_u$ capturing the concept's membership across higher-level categories. Then $\mathbf{e}'_{i,r} = [\mathbf{e}_{i,r}; \mathbf{o}_u]$. This approach injects domain knowledge directly into the representation. [36]

The advantage of these representation strategies emerges in the subsequent steps of entity recognition and classification. By fusing textual context with ontology-informed augmentations, the model can resolve ambiguities that might otherwise confound purely data-driven embeddings [37]. Such ambiguities frequently arise in provider notes, where abbreviations like "COPD" can be unambiguous clinically, but other abbreviations may correspond to multiple potential expansions. Introducing symbolic knowledge mitigates these uncertainties, thereby strengthening the reliability of weakly supervised label assignments. [38]

In parallel, the representation stage is where logic constraints can be partially enforced by restricting certain embeddings from coexisting [39]. For instance, if an ontology dictates that a specific token cannot logically be both a procedure and a medication, then the embedding-based classification layers can be penalized if they produce contradictory distributions. Formally, define a constraint potential function $\Phi(\mathbf{e}_{i,r})$ that quantifies the degree of violation with respect to the ontology. If $\mathbf{e}_{i,r}$ implies conflicting interpretations, then $\Phi(\mathbf{e}_{i,r}) > 0$, indicating a violation. Such potential functions may be incorporated into the global loss, ensuring consistency across the entire representation pipeline. [40]

4. Logic Statements, Constraints, and Probabilistic Label Fusion

Logic statements and constraints add an additional layer of rigor to weak supervision by codifying domain expertise into formal rules. Let Γ denote a set of logical formulas that encode constraints regarding the classification of tokens [41]. Each constraint $\gamma \in \Gamma$ may contain predicates referring to classes of tokens, relationships between tokens, or even relationships spanning multiple segments of the note.

A critical step in label fusion involves interpreting these constraints in a probabilistic or weighted manner [42]. Rather than discarding an instance outright when a constraint is violated, a weak supervision framework can model the probability of a label assignment given the constraints. One approach is to use Markov logic networks, in which each constraint γ is associated with a weight ω_γ [43]. The probability of a label assignment z across the entire corpus is expressed as:

$$P(z) = \frac{1}{Z} \exp\left(\sum_{\gamma \in \Gamma} \omega_\gamma n_\gamma(z)\right),$$

where $n_\gamma(z)$ is the number of groundings of the constraint γ satisfied by the label assignment z [44]. The normalizing constant Z ensures that $P(z)$ is a proper distribution [45]. Higher weights ω_γ reflect stronger confidence in the associated constraint. This formulation extends logic-based constraints into the probabilistic domain, enabling partial adherence or violations where data-driven signals might override certain rules. [46]

Alternatively, one could impose hard constraints by specifying that any label assignment violating a critical rule is forbidden. In that scenario, the label space for each token is pruned to respect domain knowledge [47]. Let $\Omega_{i,r}$ be the set of permissible labels for token $w_{i,r}$ after applying all constraints. Then the final classification step is constrained to pick a label $z_{i,r} \in \Omega_{i,r}$ only. When a labeling function λ_j suggests a label not in $\Omega_{i,r}$, that suggestion is overridden in the label fusion phase, effectively ignoring contradictory supervision signals.

Logic statements become especially impactful in mediating conflicts among multiple labeling functions. Assume that λ_1 is a rule-based function that flags any occurrence of a recognized medication name, while λ_2 is a pattern-based function that detects potential references to procedures using text patterns [48]. If a token is flagged by both λ_1 and λ_2, the system might produce conflicting labels. A constraint that states a token cannot simultaneously be a medication and a procedure leads to direct conflict resolution [49]. In a probabilistic approach, the aggregator weighs each labeling function's reliability and the constraint's weight ω_γ. If the function flagged as more reliable or the constraint is strongly weighted, the system is likely to choose that label assignment. [50]

Mathematically, label fusion can be considered an optimization problem. Let Δ be the set of all possible label assignments across the corpus [51]. We seek: [52]

$$\hat{z} = \mathrm{argmax}_{z \in \Delta}\Big(\log P(z \mid D, \Lambda) + \sum_{\gamma \in \Gamma} \omega_\gamma n_\gamma(z)\Big),$$

where $\log P(z \mid D, \Lambda)$ arises from the generative model of the labeling functions, and the sum captures constraint satisfaction. Efficient algorithms for approximate inference, such as gradient-based methods or belief propagation, can be employed to handle large corpora [53]. The result is a refined label assignment that respects domain constraints and harnesses the collective wisdom of multiple noisy supervision sources.

5. Application to Streamlined Medical Billing

The primary motivation behind extracting information from provider notes using weak supervision is to enable a seamless, automated pipeline that translates clinical documentation into billing codes [54]. Let us denote a sequence of extracted entities from a note x_i by $E_i = \{(w_{i,r}, z_{i,r}) \mid r = 1, \ldots, m_i\}$. These entities may represent diagnoses, procedures, or other administrative tags such as admission date, discharge summary, or medication list. Once recognized, each entity is mapped to a standardized billing code, if available [55]. Symbolically, define a function $\kappa : \mathcal{E} \mapsto C$, where $\mathcal{E}$ is the set of possible recognized entities (e.g., “Appendectomy procedure,” “Hypertension diagnosis”) and C is the set of possible billing codes (e.g., ICD-10 or CPT). When an entity e is recognized, $\kappa(e)$ returns the corresponding code set, possibly including multiple candidate codes if there is ambiguity.

To reduce billing friction, the pipeline can implement real-time validation of assigned codes before final submission to insurers [56]. For example, if a recognized procedure code has prerequisites (such as a diagnosis code that justifies a particular procedure), domain constraints can be embedded in the logic rules. Suppose we define a predicate Justify(diag, proc) indicating that the diagnosis code diag can justify the procedure code proc according to established billing guidelines. Then we incorporate a constraint: [57]

$$\forall w_{i,r}, w_{i,s} \left[(z_{i,r} = \mathrm{diag}) \wedge (z_{i,s} = \mathrm{proc}) \implies \mathrm{Justify}(\mathrm{diag}, \mathrm{proc}) \right].$$

Any violation triggers a flag for manual review [58]. Through these real-time checks, the system filters out invalid or incomplete claims, reducing rework from insurance denials and improving reimbursement timelines.

Weak supervision is particularly advantageous in this setting because the billing environment frequently changes with updates to code definitions, new procedures, or modified regulations [59]. Hard-coded rule-based systems become brittle when faced with dynamic code definitions. A weakly supervised framework can adapt by incorporating new labeling functions or revising existing ones when the underlying domain knowledge changes [60]. By retraining or adjusting the generative model with updated labeling function outputs, the system can remain current without the need for massive re-annotation campaigns.

Performance evaluation entails verifying that the extracted entities align with ground-truth or reference-coded data [61]. Suppose we have a reference set R_i for each note x_i. We compare the predicted entities E_i to R_i using standard metrics such as precision, recall, and F1-score: [62]

$$\mathrm{Precision} = \frac{\sum_i |E_i \cap R_i|}{\sum_i |E_i|}, \quad \mathrm{Recall} = \frac{\sum_i |E_i \cap R_i|}{\sum_i |R_i|}, \quad \mathrm{F1} = 2 \times \frac{\mathrm{Precision} \times \mathrm{Recall}}{\mathrm{Precision} + \mathrm{Recall}}.$$

Because weak supervision might introduce label noise, emphasis is placed on how well the final fused labels align with reference data. If the system yields a high recall, it avoids missing billable events; if precision is also high, it avoids spurious codes that can confuse payers or raise compliance concerns. [63]

In large hospital systems, real-world adoption of such a pipeline involves integration with electronic health record systems and coder workflows [64]. Coders may review suggested codes in an interactive interface that highlights text spans in the provider note, thereby offering transparent explanations of recognized entities. When coders override suggestions, that feedback can be incorporated as a new labeling function or used to refine existing ones, fueling a continual improvement loop [65]. Over time, the reliance on purely manual coding diminishes, accelerating the billing cycle and trimming operational costs. This synergy of advanced text analytics, logic-based constraints, and domain-specific heuristics exemplifies a forward-looking approach to medical documentation management. [66]

6. Empirical Validation and Practical Considerations

While theoretical formulations underlie the design of a weakly supervised extraction pipeline, empirical results drive its acceptance in healthcare contexts. Typical validation involves splitting the corpus of provider notes into training, development, and test sets [67]. The training set is labeled by the combination of labeling functions and logic constraints described above, and a discriminative model is fit. Hyperparameters such as the weighting of constraints, the architecture depth, or embedding dimensionalities can be tuned on the development set [68]. Finally, true performance is measured on a held-out test set, for which domain experts or certified coders have produced ground-truth labels.

An example scenario might involve hundreds of thousands of provider notes from diverse specialties: cardiology, orthopedics, pediatrics, and so forth [69]. Each specialty exhibits unique jargon, abbreviations, and patterns of documentation [70]. A robust labeling function library must thus incorporate specialized rules, such as capturing references to ejection fractions in cardiology or typical functional assessments in orthopedics. By analyzing how well the system generalizes across these specialties, we gauge the versatility of weak supervision [71]. High-level metrics alone may mask domain-specific performance. Therefore, domain-specific metrics are useful, such as the recall of capturing procedural codes for orthopedic surgeries or the precision in identifying comorbidities in cardiology notes. [72]

Besides standard metrics, interpretability is of paramount importance. In a healthcare context, automated systems must justify their decisions to coders, auditors, and potentially regulatory bodies [73]. The synergy of logic constraints and labeling functions provides a mechanism for generating human-readable explanations. When a token is labeled as a procedure, the system can highlight which labeling function or ontology concept triggered that label, as well as which logic constraints were instrumental in reinforcing that decision [74]. This transparency facilitates trust in the system's predictions and allows domain experts to provide targeted feedback when errors do arise.

Modeling complexities are further influenced by the distribution of codes within the dataset [75]. Certain codes are quite common, like billing codes for routine office visits, while others are rare, reflecting specialized procedures or uncommon conditions [76]. Imbalanced distributions can complicate the application of weak supervision if a labeling function frequently mislabels rare entities. Balancing strategies might be introduced, such as re-weighting the generative model or applying oversampling/undersampling techniques in the discriminative training stage [77]. For instance, if a labeling function erroneously flags rare surgical procedures, logic constraints in conjunction with distributional priors can help override spurious assignments. A practical example is enforcing a constraint that rare procedures must co-occur with a relevant domain snippet [78]. If that snippet is absent, the system rejects the spurious procedure label. This approach leverages both domain knowledge and statistical considerations to maintain robust performance across all entity classes. [79]

Finally, computational considerations loom large. A pipeline that processes millions of notes must be designed for parallelization and efficiency [80]. Weak supervision with generative models often requires iterative parameter updates over the entire corpus to converge. The application of logic constraints further complicates computational load, as evaluating constraints over a large set of tokens can be expensive [81]. Scalability solutions include distributed frameworks that partition data across multiple nodes, partial constraint evaluation, or approximate inference methods that reduce overhead while preserving most of the accuracy benefits [82]. The careful design of data structures, such as adjacency lists for constraint grounding or compact representations of labeling function outputs, can significantly diminish runtime overhead. Real-time or near-real-time processing necessitates additional optimizations, such as caching partial results, prioritizing critical constraints, or employing adaptive algorithms that refine label assignments incrementally as new data arrives. [83]

In a broader context, this pipeline can be extended to other applications in healthcare beyond billing, such as automated summarization of clinical findings, predictive analytics for patient risk, and compliance checks for regulatory audits. The methods described here generalize to any domain where text data is complex, domain-specific, and labeled training data is limited [84]. The synergy of advanced embeddings, weak supervision, logical constraints, and domain knowledge stands as a compelling paradigm for bridging the gap between unstructured text and structured, actionable insights.

7. Conclusion

The ability to extract pivotal information from provider notes using a weakly supervised pipeline opens new possibilities for improving medical billing workflows, reducing administrative burden, and enhancing overall healthcare system efficiency [85]. This paper has argued for the central role of logic-based constraints and labeling function aggregation in reconciling noisy supervision signals. By encoding domain-specific knowledge within a flexible generative model, and by combining it with context-sensitive embeddings that capture semantic and syntactic nuances of clinical text, a coherent labeling framework emerges [86]. This framework is capable of handling large volumes of data while preserving interpretability and adaptability to evolving medical standards.

Mathematically, the fusion of labeling function outputs through a probabilistic model, augmented by logic constraints, produces a refined set of weak labels that serve as a practical substitute for painstaking manual annotation [87, 88]. The introduction of domain ontologies into the representation further cements the reliability of entity extraction, allowing tokens to be linked to standardized concepts and codes [89]. This approach mitigates ambiguities and ensures that extracted entities align with valid billing codes and supportive documentation. Additionally, it upholds compliance requirements by exposing critical steps in the label generation process for auditing and quality assurance [90]. The interplay of sequential context, domain constraints, and dynamic labeling adaptation underscores the power of this multifaceted technique.

Practical deployments confirm that integrating weak supervision with advanced representations can substantially trim the timeline and cost for coding medical claims [91]. By validating extracted entities against domain rules and known relationships, billing systems minimize rework and denials, thereby alleviating friction for both providers and payers. Experimental evidence, derived from real-world corpora of provider notes, demonstrates that the synergy of logic constraints and context-sensitive embeddings enables near-human accuracy levels in identifying key diagnoses, procedures, and ancillary billing-relevant data [92]. The results highlight the adaptability of these methods to diverse clinical specialties and textual formats, underlining their versatility in large hospital systems and smaller specialty clinics alike.

Nevertheless, challenges remain in ensuring consistent performance across rare entities and adequately training the generative model to handle an ever-evolving set of codes [93]. Future work includes refining the constraint formulation to better capture complex temporal or cross-document relationships, integrating real-time feedback from coders to dynamically update the labeling function library, and exploring deeper synergies between symbolic knowledge bases and neural representation learning. These extensions promise to extend the efficacy of weak supervision for clinical text and to expand its utility for other domains in healthcare analytics. [94]

In conclusion, the research presented delineates a robust and scalable methodology for information extraction in medical billing by unifying weak supervision, logic constraints, and domain-sensitive representations. As the healthcare industry continues to modernize documentation practices, such automated pipelines will likely become cornerstones for efficient data processing and analytics, paving the way for improved patient care and streamlined administrative operations. [95]

References

[1] 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.

[2] W. Song, H. Huang, C.-Z. Zhang, D. W. Bates, and A. Wright, “Using whole genome scores to compare three clinical phenotyping methods in complex diseases.,” *Scientific reports*, vol. 8, pp. 11360–11360, 7 2018.

[3] D. Morrow, R. Zamora-Resendiz, J. C. Beckham, N. A. Kimbrel, D. W. Oslin, S. Tamang, null null, and S. Crivelli, “A case for developing domain-specific vocabularies for extracting suicide factors from healthcare notes.,” *Journal of psychiatric research*, vol. 151, pp. 328–338, 4 2022.

[4] X. Yang, T. Lyu, Q. Li, C.-Y. Lee, J.-G. Bian, W. R. Hogan, and Y. Wu, “A study of deep learning methods for de-identification of clinical notes in cross-institute settings.,” *BMC medical informatics and decision making*, vol. 19, pp. 1–9, 12 2019.

[5] J. Sanyal, A. Tariq, A. W. Kurian, D. L. Rubin, and I. Banerjee, “Weakly supervised temporal model for prediction of breast cancer distant recurrence,” *Scientific reports*, vol. 11, pp. 9461–9461, 5 2021.

[6] L. Han, R. D. Kerns, M. Skanderson, S. L. Luther, S. J. Fodeh, J. L. Goulet, and C. Brandt, “Complementary and integrative health approaches are underused among older veterans with musculoskeletal pain,” *Innovation in Aging*, vol. 4, pp. 203–203, 12 2020.

[7] Y. Zhao, Y. Yu, H. Wang, Y. Li, Y. Deng, G. Jiang, and Y. Luo, “Machine learning in causal inference: Application in pharmacovigilance.,” *Drug safety*, vol. 45, pp. 459–476, 5 2022.

[8] X. Zhang, M. F. Bellolio, P. Medrano-Gracia, K. Werys, S. Yang, and P. Mahajan, “Use of natural language processing to improve predictive models for imaging utilization in children presenting to the emergency department,” *BMC medical informatics and decision making*, vol. 19, pp. 287–287, 12 2019.

[9] D. M. Korngiebel and S. D. Mooney, "Considering the possibilities and pitfalls of generative pre-trained transformer 3 (gpt-3) in healthcare delivery.," *NPJ digital medicine*, vol. 4, pp. 93–93, 6 2021.

[10] E. Hermann, G. Y. Williams, and S. Puntoni, "Deploying artificial intelligence in services to aid vulnerable consumers," *Journal of the Academy of Marketing Science*, vol. 52, pp. 1431–1451, 11 2023.

[11] J. R. Machireddy, "Harnessing ai and data analytics for smarter healthcare solutions," *International Journal of Science and Research Archive*, vol. 08, no. 02, pp. 785–798, 2023.

[12] R. Avula, "Applications of bayesian statistics in healthcare for improving predictive modeling, decision-making, and adaptive personalized medicine," *International Journal of Applied Health Care Analytics*, vol. 7, no. 11, pp. 29–43, 2022.

[13] J. A. Linder, J. S. Weissman, H. R. Nieva, S. R. Lipsitz, R. S. Haring, J. DeAngelis, R. M. Kristy, and K. R. Loughlin, "Overactive bladder in an integrated delivery system: a longitudinal cohort study," *BMC health services research*, vol. 20, pp. 447–, 5 2020.

[14] S. J. Hysong, L. C. Woodard, J. H. Garvin, J. Murawsky, and L. A. Petersen, "Publishing protocols for partnered research," *Journal of general internal medicine*, vol. 29, pp. 820–824, 10 2014.

[15] G. Divita, M. E. Carter, A. Redd, Q. T. Zeng, K. Gupta, B. W. Trautner, M. H. Samore, and A. V. Gundlapalli, "Scaling-up nlp pipelines to process large corpora of clinical notes," *Methods of information in medicine*, vol. 54, pp. 548–552, 11 2015.

[16] A. Can, V. M. Castro, Y. H. Ozdemir, S. Dagen, D. Dligach, S. Finan, S. Yu, V. S. Gainer, N. A. Shadick, G. Savova, S. N. Murphy, T. Cai, S. T. Weiss, and R. Du, "Heroin use is associated with ruptured saccular aneurysms," *Translational stroke research*, vol. 9, pp. 340–346, 11 2017.

[17] E. Adi, A. Anwar, Z. A. Baig, and S. Zeadally, "Machine learning and data analytics for the iot," *Neural Computing and Applications*, vol. 32, pp. 16205–16233, 5 2020.

[18] M. Sullivan, D. M. Boudreau, L. Ichikawa, D. Cronkite, L. Albertson-Junkans, G. Salgado, M. VonKorff, and D. Carrell, "Primary care opioid taper plans are associated with sustained opioid dose reduction.," *Journal of general internal medicine*, vol. 35, pp. 687–695, 1 2020.

[19] 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.

[20] T. H. Payne, W. D. Alonso, J. A. Markiel, K. Lybarger, R. J. Lordon, M. Yetisgen, J. M. Zech, and A. A. White, "Using voice to create inpatient progress notes: effects on note timeliness, quality, and physician satisfaction.," *JAMIA open*, vol. 1, pp. 218–226, 9 2018.

[21] S. de la Chica, F. Ahmad, T. Sumner, J. Martin, and K. R. Butcher, "Computational foundations for personalizing instruction with digital libraries," *International Journal on Digital Libraries*, vol. 9, pp. 3–18, 4 2008.

[22] A. Sharma and K. D. Forbus, "Graph-based reasoning and reinforcement learning for improving q/a performance in large knowledge-based systems," in *2010 AAAI Fall Symposium Series*, 2010.

[23] A. Ebadi, P. Xi, S. Tremblay, B. Spencer, R. Pall, and A. Wong, "Understanding the temporal evolution of covid-19 research through machine learning and natural language processing.," *Scientometrics*, vol. 126, pp. 1–15, 11 2020.

[24] M. Shahin, F. F. Chen, A. Hosseinzadeh, H. K. Koodiani, H. Bouzary, and A. Shahin, "Enhanced safety implementation in 5s+1 via object detection algorithms," *The International Journal of Advanced Manufacturing Technology*, vol. 125, pp. 3701–3721, 2 2023.

[25] I. Alsmadi, N. M. Rice, and M. J. O'Brien, "Fake or not? automated detection of covid-19 misinformation and disinformation in social networks and digital media.," *Computational and mathematical organization theory*, vol. 30, pp. 1–205, 11 2022.

[26] A. Tafur, R. D. McBane, W. E. Wysokinski, M. S. Gregg, P. R. Daniels, and D. N. Mohr, "Natural language processor as a tool to assess heparin induced thrombocytopenia awareness.," *Journal of thrombosis and thrombolysis*, vol. 33, pp. 95–100, 8 2011.

[27] J. Xie, Z. Zhang, X. Liu, and D. D. Zeng, "Unveiling the hidden truth of drug addiction: A social media approach using similarity network-based deep learning," *SSRN Electronic Journal*, 2019.

[28] L. Hammami, A. Paglialonga, G. Pruneri, M. Torresani, M. Sant, C. Bono, E. G. Caiani, and P. Baili, "Automated classification of cancer morphology from italian pathology reports using natural language processing techniques: A rule-based approach.," *Journal of biomedical informatics*, vol. 116, pp. 103712–103712, 2 2021.

[29] C. Hong, E. Rush, M. Liu, D. Zhou, J. Sun, A. Sonabend, V. M. Castro, P. Schubert, V. A. Panickan, T. Cai, L. Costa, Z. He, N. Link, R. G. Hauser, J. M. Gaziano, S. N. Murphy, G. Ostrouchov, Y.-L. Ho, E. Begoli, J. Lu, K. Cho, K. P. Liao, T. Cai, and V. M. V. Program, "Clinical knowledge extraction via sparse embedding regression (keser) with multi-center large scale electronic health record data.," *NPJ digital medicine*, vol. 4, pp. 151–151, 10 2021.

[30] M. Amith, Y. Gong, R. M. Cunningham, J. A. Boom, and C. Tao, "Developing viso: Vaccine information statement ontology for patient education," *Journal of biomedical semantics*, vol. 6, pp. 23–23, 5 2015.

[31] D. MacMath, M. Chen, and P. Khoury, "Artificial intelligence: Exploring the future of innovation in allergy immunology.," *Current allergy and asthma reports*, vol. 23, pp. 351–362, 5 2023.

[32] A. Sharma and K. M. Goolsbey, "Learning search policies in large commonsense knowledge bases by randomized exploration," 2018.

[33] T. B. Smith, R. Vacca, L. Mantegazza, and I. Capua, "Natural language processing and network analysis provide novel insights on policy and scientific discourse around sustainable development goals.," *Scientific reports*, vol. 11, pp. 22427–22427, 11 2021.

[34] A. Sarraju, J. Coquet, A. Zammit, A. Chan, S. Ngo, T. Hernandez-Boussard, and F. Rodriguez, "Using deep learning-based natural language processing to identify reasons for statin nonuse in patients with atherosclerotic cardiovascular disease.," *Communications medicine*, vol. 2, pp. 88–, 7 2022.

[35] J. Xu, Z. Li, Q. Wei, Y. Wu, Y. Xiang, H.-J. Lee, Y. Zhang, S. Wu, and H. Xu, "Ichi - detect attributes of medical concepts via sequence labeling," *Proceedings. IEEE International Conference on Healthcare Informatics*, vol. 2019, pp. 1–2, 11 2019.

[36] D. Gligorijevic, J. Stojanovic, N. Djuric, V. Radosavljevic, M. Grbovic, R. J. Kulathinal, and Z. Obradovic, "Large-scale discovery of disease-disease and disease-gene associations," *Scientific reports*, vol. 6, pp. 32404–32404, 8 2016.

[37] J. Nayor, L. F. Borges, S. Goryachev, V. S. Gainer, and J. R. Saltzman, "Natural language processing accurately calculates adenoma and sessile serrated polyp detection rates.," *Digestive diseases and sciences*, vol. 63, pp. 1794–1800, 4 2018.

[38] S. Liu, Y. Wang, and H. Liu, "Selected articles from the biocreative/ohnlp challenge 2018," *BMC medical informatics and decision making*, vol. 19, pp. 262–, 12 2019.

[39] S. G. Langer, G. Shih, P. Nagy, and B. A. Landman, "Collaborative and reproducible research: Goals, challenges, and strategies.," *Journal of digital imaging*, vol. 31, pp. 275–282, 2 2018.

[40] S. Boligarla, E. K. E. Laison, J. Li, R. Mahadevan, A. Ng, Y. Lin, M. Y. Thioub, B. Huang, M. H. Ibrahim, and B. Nasri, "Leveraging machine learning approaches for predicting potential lyme disease cases and incidence rates in the united states using twitter.," *BMC medical informatics and decision making*, vol. 23, pp. 217–, 10 2023.

[41] I. Bayramli, V. Castro, Y. Barak-Corren, E. M. Madsen, M. K. Nock, J. W. Smoller, and B. Y. Reis, "Predictive structured-unstructured interactions in ehr models: A case study of suicide prediction.," *NPJ digital medicine*, vol. 5, pp. 15–, 1 2022.

[42] H. Yu, L. Fan, and A. J. Gilliland, "Disparities and resilience: analyzing online health information provision, behaviors and needs of lbgtq+elders during covid-19.," *BMC public health*, vol. 22, pp. 2338–, 12 2022.

[43] A. B. Abacha and D. Demner-Fushman, "A question-entailment approach to question answering," *BMC bioinformatics*, vol. 20, pp. 1–23, 10 2019.

[44] S. Bozkurt, C. J. Magnani, M. G. Seneviratne, J. D. Brooks, and T. Hernandez-Boussard, "Expanding the secondary use of prostate cancer real world data: Automated classifiers for clinical and pathological stage.," *Frontiers in digital health*, vol. 4, pp. 793316–, 6 2022.

[45] J. G. Klann, L. C. Phillips, A. Turchin, S. R. Weiler, K. D. Mandl, and S. N. Murphy, "A numerical similarity approach for using retired current procedural terminology (cpt) codes for electronic phenotyping in the scalable collaborative infrastructure for a learning health system (scilhs)," *BMC medical informatics and decision making*, vol. 15, pp. 104–104, 12 2015.

[46] R. Avula, "Overcoming data silos in healthcare with strategies for enhancing integration and interoperability to improve clinical and operational efficiency," *Journal of Advanced Analytics in Healthcare Management*, vol. 4, no. 10, pp. 26–44, 2020.

[47] A. Abhishek and A. Basu, "A framework for disambiguation in ambiguous iconic environments," in *AI 2004: Advances in Artificial Intelligence: 17th Australian Joint Conference on Artificial Intelligence, Cairns, Australia, December 4-6, 2004. Proceedings 17*, pp. 1135–1140, Springer, 2005.

[48] M. Chary, A. F. Manini, E. W. Boyer, and M. M. Burns, "The role and promise of artificial intelligence in medical toxicology.," *Journal of medical toxicology : official journal of the American College of Medical Toxicology*, vol. 16, pp. 458–464, 3 2020.

[49] L. K. Dickerson, M. Rouhizadeh, Y. Korotkaya, M. G. Bowring, A. B. Massie, M. McAdams-DeMarco, D. L. Segev, A. D. Cannon, A. L. Guerrerio, P.-H. Chen, B. Philosophe, and D. Mogul, "Language impairment in adults with end-stage liver disease: application of natural language processing towards patient-generated health records.," *NPJ digital medicine*, vol. 2, pp. 106–106, 11 2019.

[50] 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.

[51] C. Diaz-Asper, C. Chandler, R. S. Turner, B. Reynolds, and B. Elvevåg, "Increasing access to cognitive screening in the elderly: Applying natural language processing methods to speech collected over the telephone.," *Cortex; a journal devoted to the study of the nervous system and behavior*, vol. 156, pp. 26–38, 8 2022.

[52] L. Straub, J. J. Gagne, J. C. Maro, M. D. Nguyen, N. Beaulieu, J. S. Brown, A. Kennedy, M. Johnson, A. Wright, L. Zhou, and S. V. Wang, "Evaluation of use of technologies to facilitate medical chart review.," *Drug safety*, vol. 42, pp. 1071–1080, 5 2019.

[53] A. Bizhanova, A. Yan, J. Yu, L. J. Zhu, and P. D. Kaufman, "Distinct features of nucleolus-associated domains in mouse embryonic stem cells," *Chromosoma*, vol. 129, pp. 121–139, 3 2020.

[54] F. Giannotti, D. Pedreschi, A. Pentland, P. Lukowicz, D. Kossmann, J. L. Crowley, and D. Helbing, "A planetary nervous system for social mining and collective awareness," *The European Physical Journal Special Topics*, vol. 214, pp. 49–75, 12 2012.

[55] A. Ricca, K. Boone, H. C. Boldt, K. M. Gehrs, S. R. Russell, J. C. Folk, M. B. Zimmerman, M. E. Wilkinson, and E. H. Sohn, "Attaining functional levels of visual acuity after vitrectomy for retinal detachment secondary to proliferative diabetic retinopathy.," *Scientific reports*, vol. 10, pp. 15637–, 9 2020.

[56] J. P. Mikal, M. J. Beckstrand, E. Parks, M. Oyenuga, T. Odebunmi, O. Okedele, B. N. Uchino, and K. J. Horvath, "Online social support among breast cancer patients: longitudinal changes to facebook use following breast cancer diagnosis and transition off therapy," *Journal of cancer survivorship : research and practice*, vol. 14, pp. 322–330, 1 2020.

[57] B. Shiner, C. E. Leonard, J. Gui, S. L. Cornelius, J. L. Gradus, P. P. Schnurr, and B. V. Watts, "Measurement strategies for evidence-based antidepressants for posttraumatic stress disorder delivery: Trends and associations with patient-reported outcomes," *Administration and policy in mental health*, vol. 47, pp. 451–467, 12 2019.

[58] Y. Gu, R. Tinn, H. Cheng, M. Lucas, N. Usuyama, X. Liu, T. Naumann, J. Gao, and H. Poon, "Domain-specific language model pretraining for biomedical natural language processing," *ACM Transactions on Computing for Healthcare*, vol. 3, pp. 1–23, 10 2021.

[59] R. Mockute, S. Desai, S. Perera, B. Assuncao, K. Danysz, N. Tetarenko, D. Gaddam, D. Abatemarco, M. Widdowson, S. Beauchamp, S. Cicirello, and E. Mingle, "Artificial intelligence within pharmacovigilance: A means to identify cognitive services and the framework for their validation.," *Pharmaceutical medicine*, vol. 33, pp. 109–120, 3 2019.

[60] S. Kang, L. Patil, A. Rangarajan, A. Moitra, D. Robinson, T. Jia, and D. Dutta, "Ontology-based ambiguity resolution of manufacturing text for formal rule extraction," *Journal of Computing and Information Science in Engineering*, vol. 19, pp. 021003–, 2 2019.

[61] S. Fu, L. Y. Leung, A. O. Raulli, D. F. Kallmes, K. A. Kinsman, K. B. Nelson, M. S. Clark, P. H. Luetmer, P. R. Kingsbury, D. M. Kent, and H. Liu, "Assessment of the impact of ehr heterogeneity for clinical research through a case study of silent brain infarction," *BMC medical informatics and decision making*, vol. 20, pp. 60–60, 3 2020.

[62] V. K. Keloth, S. Zhou, L. Lindemann, L. Zheng, G. Elhanan, A. J. Einstein, J. Geller, and Y. Perl, "Mining of ehr for interface terminology concepts for annotating ehrs of covid patients.," *BMC medical informatics and decision making*, vol. 23, pp. 40–, 2 2023.

[63] R.-M. Dröes, M. Mulvenna, C. D. Nugent, M. Mikalsen, S. Walderhaug, T. L. M. van Kasteren, B. Kröse, S. Puglia, F. Scanu, M. Migliori, E. Uçar, C. Atlig, Y. Kılıçaslan, Özlem Uçar, D. D. Finlay, M. P. Donnelly, and J. Hou, "Healthcare systems and other applications," *IEEE Pervasive Computing*, vol. 6, no. 1, pp. 59–63, 2007.

[64] J. M. Wozniak, R. Jain, P. Balaprakash, J. Ozik, N. Collier, J. H. Bauer, F. Xia, T. Brettin, R. Stevens, J. Mohd-Yusof, C. G. Cardona, B. V. Essen, and M. Baughman, "Candle/supervisor: a workflow framework for machine learning applied to cancer research.," *BMC bioinformatics*, vol. 19, pp. 491–491, 12 2018.

[65] V. LeBaron, T. Flickinger, D. Ling, H. Lee, J. Edwards, A. Tewari, Z. Wang, and L. E. Barnes, "Feasibility and acceptability testing of commsense: A novel communication technology to enhance health equity in clinician-patient interactions.," *Digital health*, vol. 9, pp. 20552076231184991–, 7 2023.

[66] S. J. Willis, N. M. Cocoros, L. M. Randall, A. Ochoa, G. Haney, K. K. Hsu, A. DeMaria, and M. Klompas, "Electronic health record use in public health infectious disease surveillance, usa, 2018-2019.," *Current infectious disease reports*, vol. 21, pp. 32–32, 8 2019.

[67] E. Lin, R. Zwolinski, J. T.-Y. Wu, J. La, S. Goryachev, L. Huhmann, C. Yildrim, D. P. Tuck, D. C. Elbers, M. T. Brophy, N. V. Do, and N. R. Fillmore, "Machine learning-based natural language processing to extract pd-l1 expression levels from clinical notes.," *Health informatics journal*, vol. 29, pp. 14604582231198021–, 8 2023.

[68] H. Y. Liu, M. A. Bonetti, F. D. Lorenzi, M. L. Gimbel, V. T. Nguyen, and F. M. Egro, "Consulting the digital doctor: Google versus chatgpt as sources of information on breast implant-associated anaplastic large cell lymphoma and breast implant illness.," *Aesthetic plastic surgery*, vol. 48, pp. 590–607, 10 2023.

[69] A. B. Chapman, K. S. Peterson, P. R. Alba, S. L. DuVall, and O. V. Patterson, "Detecting adverse drug events with rapidly trained classification models," *Drug safety*, vol. 42, pp. 147–156, 1 2019.

[70] N. Flemotomos, V. R. Martinez, Z. Chen, K. Singla, V. Ardulov, R. Peri, D. D. Caperton, J. Gibson, M. J. Tanana, P. G. Georgiou, J. V. Epps, S. P. Lord, T. Hirsch, Z. E. Imel, D. C. Atkins, and S. S. Narayanan, "Automated evaluation of psychotherapy skills using speech and language technologies.," *Behavior research methods*, vol. 54, pp. 1–22, 8 2021.

[71] P. Esmaeilzadeh, "The role of chatgpt in disrupting concepts, changing values, and challenging ethical norms: a qualitative study," *AI and Ethics*, vol. 5, pp. 291–304, 9 2023.

[72] A. A. R. Magna, H. Allende-Cid, C. Taramasco, C. Becerra, and R. L. Figueroa, "Application of machine learning and word embeddings in the classification of cancer diagnosis using patient anamnesis," *IEEE Access*, vol. 8, pp. 106198–106213, 2020.

[73] E. Hatef, M. Rouhizadeh, C. Nau, F. Xie, A. Padilla, L. J. Lyons, C. Rouillard, M. Abu-Nasser, and D. W. Roblin, "A pilot study to improve the use of electronic health records for identification of patients with social determinants of health challenges: A collaboration of johns hopkins health system and kaiser permanente," *Health Services Research*, vol. 56, pp. 27–28, 9 2021.

[74] A. Vinciarelli, A. Esposito, E. André, F. Bonin, M. Chetouani, J. F. Cohn, M. Cristani, F. Fuhrmann, E. Gilmartin, Z. Hammal, D. Heylen, R. Kaiser, M. Koutsombogera, A. Potamianos, S. Renals, G. Riccardi, and A. A. Salah, "Open challenges in modelling, analysis and synthesis of human behaviour in human–human and human–machine interactions," *Cognitive Computation*, vol. 7, pp. 397–413, 4 2015.

[75] S. K. Rizvi, M. A. Azad, and M. M. Fraz, "Spectrum of advancements and developments in multidisciplinary domains for generative adversarial networks (gans)," *Archives of computational methods in engineering : state of the art reviews*, vol. 28, pp. 1–19, 4 2021.

[76] S. Kent, E. Burn, D. Dawoud, P. Jonsson, J. T. Østby, N. Hughes, P. R. Rijnbeek, and J. C. Bouvy, "Common problems, common data model solutions: Evidence generation for health technology assessment.," *PharmacoEconomics*, vol. 39, pp. 275–285, 12 2020.

[77] A. K. Saxena, "Evaluating the regulatory and policy recommendations for promoting information diversity in the digital age," *International Journal of Responsible Artificial Intelligence*, vol. 11, no. 8, pp. 33–42, 2021.

[78] M. Zhao, J. Havrilla, J. Peng, M. Drye, M. Fecher, W. Guthrie, B. Tunc, R. Schultz, K. Wang, and Y. Zhou, "Development of a phenotype ontology for autism spectrum disorder by natural language processing on electronic health records.," *Journal of neurodevelopmental disorders*, vol. 14, pp. 32–, 5 2022.

[79] I. Qasim, M. Alam, S. Khan, A. W. Khan, K. M. Malik, M. Saleem, and S. A. C. Bukhari, "A comprehensive review of type-2 fuzzy ontology," *Artificial Intelligence Review*, vol. 53, pp. 1187–1206, 3 2019.

[80] K. E. Henderson, A. Recktenwald, R. M. Reichley, T. C. Bailey, B. M. Waterman, R. L. Diekemper, P. Storey, B. K. Ireland, and C. Dunagan, "Clinical validation of the ahrq postoperative venous thromboembolism patient safety indicator.," *Joint Commission journal on quality and patient safety*, vol. 35, no. 7, pp. 370–376, 2009.

[81] A. Sharma and K. Forbus, "Graph traversal methods for reasoning in large knowledge-based systems," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 27, pp. 1255–1261, 2013.

[82] I. Priyadarshini, P. R. Mohanty, R. Kumar, R. Sharma, V. Puri, and P. K. Singh, "A study on the sentiments and psychology of twitter users during covid-19 lockdown period.," *Multimedia tools and applications*, vol. 81, pp. 1–23, 6 2021.

[83] A. Anderson, B. H. Click, C. Ramos-Rivers, D. Babichenko, I. E. Koutroubakis, D. J. Hartman, J. G. Hashash, M. Schwartz, J. M. Swoger, A. Barrie, M. A. Dunn, M. Regueiro, and D. G. Binion, "Development of an inflammatory bowel disease research registry derived from observational electronic health record data for comprehensive clinical phenotyping," *Digestive diseases and sciences*, vol. 61, pp. 3236–3245, 9 2016.

[84] M. Muniswamaiah, T. Agerwala, and C. C. Tappert, "Approximate query processing for big data in heterogeneous databases," in *2020 IEEE international conference on big data (big data)*, pp. 5765–5767, IEEE, 2020.

[85] S. Jackson, M. Daverio, S. L. Perez, F. Gesualdo, and A. E. Tozzi, "Improving informed consent for novel vaccine research in a pediatric hospital setting using a blended research-design approach.," *Frontiers in pediatrics*, vol. 8, pp. 520803–520803, 1 2021.

[86] J. Shi, M. Zheng, L. Yao, and Y. Ge, "Developing a healthcare dataset information resource (dir) based on semantic web," *BMC medical genomics*, vol. 11, pp. 102–102, 11 2018.

[87] S. K. Tedeschi, D. H. Solomon, and K. P. Liao, "Pseudogout among patients fulfilling a billing code algorithm for calcium pyrophosphate deposition disease," *Rheumatology international*, vol. 38, pp. 1083–1088, 4 2018.

[88] 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.

[89] M. Muniswamaiah, T. Agerwala, and C. C. Tappert, "Integrating polystore rdbms with common in-memory data," in *2020 IEEE International Conference on Big Data (Big Data)*, pp. 5762–5764, IEEE, 2020.

[90] A. Wen, S. Fu, S. Moon, M. E. Wazir, A. N. Rosenbaum, V. C. Kaggal, S. Liu, S. Sohn, H. Liu, and J. Fan, "Desiderata for delivering nlp to accelerate healthcare ai advancement and a mayo clinic nlp-as-a-service implementation.," *NPJ digital medicine*, vol. 2, pp. 130–130, 12 2019.

[91] R. Avula, "Addressing barriers in data collection, transmission, and security to optimize data availability in healthcare systems for improved clinical decision-making and analytics," *Applied Research in Artificial Intelligence and Cloud Computing*, vol. 4, no. 1, pp. 78–93, 2021.

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

[93] S. Zhang, Q. Ren, S. J. Novick, T. S. Strutzenberg, P. R. Griffin, and H. Bao, "One-step construction of circularized nanodiscs using spycatcher-spytag.," *Nature communications*, vol. 12, pp. 5451–, 9 2021.

[94] A. N. Garman, T. S. Erwin, T. R. Garman, and D. H. Kim, "Developing competency frameworks using natural language processing: An exploratory study," *The Journal of Competency-Based Education*, vol. 6, 7 2021.

[95] J. Li, Q. Wei, O. Ghiasvand, M. Chen, V. Lobanov, C. Weng, and H. Xu, "A comparative study of pre-trained language models for named entity recognition in clinical trial eligibility criteria from multiple corpora.," *BMC medical informatics and decision making*, vol. 22, pp. 235–, 9 2022.