

Federated Learning Framework for Privacy-Preserving Clinical Named Entity Recognition

Farah Nadira Binti Salleh¹

¹Melaka Digital Sciences University, Department of Computer Engineering, Jalan Teknologi 3, Ayer Keroh, Melaka, Malaysia.

Abstract

Federated learning has emerged as a promising paradigm for collaborative model training without centralized data aggregation. This approach offers a privacy-preserving framework capable of accommodating stringent requirements associated with medical data. Clinical Named Entity Recognition relies on identifying and extracting pertinent medical concepts from unstructured text. However, the sharing of sensitive clinical information raises data ownership and privacy concerns, hindering collaborative progress. Leveraging federated learning circumvents these challenges by enabling multiple clinical sites to train shared models without exchanging patient data. This paper examines an advanced federated learning framework designed to address the privacy constraints of clinical text, focusing on sophisticated embedding techniques for named entities as well as specialized aggregation protocols to ensure secure model updates. Beyond classical encryption, the proposed approach includes theoretical and practical considerations that balance performance and confidentiality. Through the integration of encryption schemes and noise perturbations, the architecture supports real-time collaboration among institutions to broaden the scope and scale of data-driven medical research. Extensive theoretical analysis and experimentation demonstrate the feasibility of privacy-preserving implementations for tasks that require domain-specific accuracy. This work offers robust insights, including how encryption, aggregation, and distributed machine learning can be unified to tackle the unique challenges of clinical named entity recognition, thereby facilitating both improved patient outcomes and research discoveries.

1. Introduction

Research on intelligent health systems has pivoted toward techniques that can harness large volumes of data while complying with rigorous privacy mandates [1]. Health records contain a wealth of information such as diagnoses, laboratory results, medications, and treatment progress notes. Identifying key clinical entities and their contextual relevance remains a cornerstone of medical text analysis [2]. Numerous computational strategies have been proposed to extract structured medical facts from unstructured narratives and to automate the indexing of patient-specific risk factors. The objective is to combine scalability with confidentiality, ensuring that regulatory constraints concerning patient data remain intact. [3]

Emerging techniques in distributed machine learning offer powerful mechanisms for combining knowledge from disparate sources. Such approaches typically require either a centralized repository or a mechanism to share potentially sensitive information. Centralized methods face legal and ethical bottlenecks because data hosting in a single location is susceptible to security risks and difficulties in obtaining patient consent [4]. By contrast, federated learning circumvents centralized storage, orchestrating model training through repeated synchronous or asynchronous communications among participating sites. Each site trains locally on its proprietary data, transmitting only model gradients or weight updates [5]. This avoids direct exposure of medical records and thus offers a measure of confidentiality. Yet challenges remain, including unbalanced data distributions, network latencies, heterogeneity in local computing resources, and the inherent complexities of integrating advanced natural language processing

models with encryption or differential privacy [6]. The nuanced domain of clinical texts, with specialized terminologies, abbreviations, and context dependencies, adds further intricacy to establishing an effective federated training ecosystem.

Such an approach is well-aligned with the task of named entity recognition for clinical texts, in which domain-specific terminologies, drug names, disease mentions, and other key medical entities must be meticulously identified. Distinguishing abbreviations from fully spelled-out terms and disambiguating repeated mentions of the same entity create additional complexities [7]. The significance of extracting these elements accurately extends far beyond academic interest, influencing patient care, biomedical research, and public health initiatives. For instance, accurate entity recognition can ensure more targeted clinical decision support and facilitate the extraction of large-scale evidence from textual corpora [8]. At the same time, privacy remains paramount, because clinical narratives often contain personally identifiable information. Direct data sharing across multiple sites could inadvertently compromise anonymity, highlighting the value of a privacy-preserving framework. [9, 10]

Federated learning frameworks typically incorporate security measures for gradient updates. This often involves encryption or noise addition that safeguards partial model information at intermediary training steps. Substantial work has been done to integrate mechanisms such as homomorphic encryption, secure multiparty computation, and differential privacy [11]. Each strategy has different trade-offs in terms of computational overhead, communication complexity, and model accuracy. Homomorphic encryption allows certain arithmetic operations to be performed directly on encrypted data, guaranteeing that only permissible computations occur before decryption [12]. Differential privacy introduces mathematical noise to ensure that individual data samples cannot be easily reconstructed from model outputs, establishing rigorous bounds on data leakage.

Federated strategies also must consider potential threats arising from malicious participants or inference attacks [13]. A dishonest party might attempt to analyze the received model parameters to uncover sensitive attributes of local data at other sites. Sophisticated adversaries might gather repeated model updates in an attempt to recreate subsets of training data, necessitating robust protocols that strictly limit information revealed across communication rounds. Additionally, the variety of data distributions among different hospitals and clinics, especially with text-based inputs, can skew the learned model if not adequately addressed [14]. Imbalanced corpora and domain mismatch can degrade performance, particularly with specialized medical jargon.

This paper seeks to integrate advanced concepts in privacy-preserving machine learning with the unique linguistic structures found in clinical text, delivering a federated framework for named entity recognition that can be extended to large consortia of data providers [15]. The focus lies in the synergy between cryptographic constructs, federated optimization, and nuanced natural language processing components that capture specialized clinical semantics. The following sections explore the technical and theoretical underpinnings of the model architecture, the interplay between secure aggregation and distributed language model fine-tuning, and the methodological choices that anchor this system [16]. Validation and theoretical bounds are also provided, revealing how well the framework scales while satisfying non-negotiable privacy guarantees. The ultimate intent is to demonstrate that such a federated approach can meet the accuracy demands of clinical named entity recognition, accommodate large datasets distributed across multiple medical institutions, and simultaneously preserve confidentiality.

The structure of the discussion unfolds with a detailed examination of federated model construction, including the logic and notation that formalize parameter updating across multiple participants [17]. Further exploration highlights how domain-specific embeddings can be fused with distributed optimization for medical text. Comprehensive mathematical formulations delineate the trade-offs between encryption schemes and differential privacy [18]. Results from computational experiments and theoretical proofs illustrate the capacity of the framework to deliver high entity extraction accuracy while offering robust protection against unauthorized data inference. This holistic perspective on federated learning for clinical named entity recognition offers a concrete example of how cutting-edge distributed computing can be harnessed to improve the safety and efficacy of healthcare analytics [19]. The culmination of this research underscores a path forward for privacy-preserving text analytics and lays the groundwork for

expanded use of advanced federated algorithms in other sensitive domains such as genomics, imaging, and multi-omics data integration.

2. Federated Aggregation and Secure Communication

Federated learning involves a collaboration among multiple institutions, each of which maintains a local dataset. Let there be N participating sites, denoted by the set $\{1, 2, \dots, N\}$. Each site i holds a local dataset D_i of cardinality $|D_i|$, which contains clinical text samples [20]. Denote by w_t^i the local model parameters at site i after t communication rounds. The global model parameters at round t are given by W_t [21]. The federated procedure typically orchestrates local training followed by aggregation according to

$$W_{t+1} = \sum_{i=1}^N \frac{|D_i|}{\sum_{k=1}^N |D_k|} w_t^i.$$

This aggregation mechanism is known as a weighted average, in which larger local datasets contribute more heavily [22]. The local models are updated by minimizing local loss functions. If $\mathcal{L}_i(w_t^i)$ denotes the loss on dataset D_i , the update step is often represented as

$$w_{t+1}^i = w_t^i - \eta \nabla \mathcal{L}_i(w_t^i),$$

where η is the learning rate and $\nabla \mathcal{L}_i$ is the local gradient operator. The aggregated parameters are then broadcast to all sites for the next round of local training.

Security arises from ensuring that model parameters, gradients, or derived values do not reveal private information [23]. The function

$$F_i(w, D_i) = \sum_{(x,y) \in D_i} f(w, x, y)$$

may embody the local objective for each site, where f is a sample-level loss term defined by the structure of the neural architecture or other machine learning model [24]. Malicious entities may attempt partial reconstruction of (x, y) from updates. The presence of named entities—diagnoses, medications, or personal details—magnifies the risk if partial gradients are intercepted. Consequently, encryption or perturbation of the update vectors may be introduced, represented by a function $\text{Enc}(\cdot)$ or $\text{Perturb}(\cdot)$, yielding

$$\text{Enc}(w_t^i) = \text{Enc}(w_t^i) \quad \text{or} \quad \text{Perturb}(w_t^i).$$

The secure aggregator then computes [25]

$$\widehat{W}_t = \sum_{i=1}^N \frac{|D_i|}{\sum_{k=1}^N |D_k|} \text{Dec}(\text{Enc}(w_t^i)),$$

or in the noise-based scenario,

$$\widehat{W}_t = \sum_{i=1}^N \frac{|D_i|}{\sum_{k=1}^N |D_k|} \text{Aggregate}(\text{Perturb}(w_t^i)).$$

An important requirement is that $\text{Dec}(\cdot)$ or $\text{Aggregate}(\cdot)$ does not expose the underlying data. The aggregator must act as a partially trusted or cryptographically constrained entity [26]. A typical approach is to use threshold encryption schemes or secure multiparty computation protocols. In threshold encryption, the public key is shared, but the private key is split among multiple parties, ensuring that no single

entity can decrypt [27]. Secure multiparty computation frameworks distribute the computations among multiple colluding-resistant servers.

Communication complexity is influenced by how frequently updates are exchanged and the dimension of model parameters. Modern deep learning architectures for text-based tasks can have millions of parameters, making naive encryption or frequent synchronization computationally impractical [28]. Therefore, the system design often includes strategies for partial parameter sharing or compression. Sparse gradient approaches transmit only substantial gradient components above a certain threshold, reducing overhead [29]. Additionally, quantization of parameter values can mitigate bandwidth usage while preserving accuracy within acceptable bounds. Each modification, though, must be examined for its influence on privacy [30]. More frequent gradient sharing may expose more pathways for inference attacks, while heavy compression can degrade model performance.

The interplay between these concerns and domain-specific text representations is crucial. Clinical texts often rely on specialized embeddings capturing medical lexicons, abbreviations, and domain ontologies such as ICD codes or SNOMED CT concepts [31]. The dimension of embedding vectors might be larger than typical word embeddings, compounding the bandwidth challenges for secure updates. This underscores the need for both cryptographic and algorithmic innovations, combining partial encryption with efficient gradient compression and domain-relevant feature extraction [32]. Without these optimizations, federated learning for entity recognition tasks becomes prohibitively expensive, either in computation or in data exposure risk.

Latency is also a pivotal factor [33]. Healthcare institutions might span geographical regions, linking data silos that operate under varying network conditions. Delays in communication may stall the global update, especially if a synchronous approach is employed. Asynchronous methods allow partial updates from available sites, though they introduce greater complexity in reconciling out-of-sync models [34]. Let τ_i be the delay associated with site i . One asynchronous update rule could be [35]

$$W_{t+1} = (1 - \alpha)W_t + \alpha \sum_{i \in S_t} \frac{|D_i|}{\sum_{k \in S_t} |D_k|} (w_t^i - W_t),$$

where S_t is the subset of sites that reported updates at round t , and α is an aggregation weight controlling stability. This distribution-based approach attempts to remain robust in the face of missing or delayed participants while preserving some notion of fairness and proportional representation. The tension between synchronous and asynchronous methods highlights that design choices in federated learning must consider the realities of large-scale deployment across diverse clinical contexts. [36]

In summary, the federated aggregation paradigm, coupled with secure communication protocols, enables distributed learning while protecting sensitive textual data. The intricacies of model size, specialized embeddings, and dynamic network environments necessitate careful consideration of cryptographic techniques, update frequencies, and asynchronous strategies. Clinical named entity recognition stands to benefit substantially from these structural and methodological refinements, as the diversity and volume of medical text are critical for achieving robust entity detection that generalizes across institutions [37]. Nevertheless, effectively managing the information flow within this distributed environment requires additional layers of privacy-preserving techniques, as discussed next.

3. Privacy Preservation through Differential and Homomorphic Methods

Clinical text data are especially sensitive, necessitating formal privacy assurances [38]. Methods based on differential privacy (DP) provide quantifiable measures of data exposure. The guarantee can be expressed logically as follows: for any two neighboring datasets D_i and D'_i differing in exactly one sample, and for any possible output set Ω , a randomized mechanism M satisfies (ϵ, δ) -differential privacy if [39]

$$\Pr[M(D_i) \in \Omega] \leq e^\epsilon \Pr[M(D'_i) \in \Omega] + \delta.$$

In this setting, each site can incorporate a local randomized mechanism to obfuscate gradients or model updates. Let Δ represent the global sensitivity of the function on which the randomization is applied. The addition of noise sampled from a suitable distribution, for instance a Gaussian with variance tied to Δ^2 , ensures minimal data leakage [40]. Concretely, if Δ is the maximum change to the gradient norm caused by altering a single sample, then

$$\tilde{g}_t^i = g_t^i + \mathcal{N}(0, \sigma^2 \Delta^2),$$

where g_t^i is the raw gradient vector for site i at round t , σ is a noise scaling parameter, and $\mathcal{N}$ denotes the Gaussian distribution. This yields the differentially private gradient $\tilde{g}_t^i$. The aggregator operates on $\tilde{g}_t^i$, thus ensuring that attempts to identify specific samples from model updates are bounded by ϵ . The trade-off involves calibrating σ to achieve acceptable privacy while retaining sufficient accuracy for named entity recognition. [41]

Homomorphic encryption enables computations on encrypted data. The function [42]

$$H(\text{Enc}(a), \text{Enc}(b)) = \text{Enc}(a \circ b),$$

where $\circ$ can be addition or multiplication, indicates that certain algebraic operations can be performed in ciphertext space. Federated learning benefits from this property because sites can encrypt their model updates, then the aggregator can sum them without decrypting intermediate values. Let $\text{Enc}_i(\cdot)$ and $\text{Dec}_i(\cdot)$ be the encryption and decryption procedures for site i . For partial homomorphic schemes that support addition, the aggregator computes [43]

$$\text{Enc}\left(\sum_{i=1}^N w_t^i\right) = \prod_{i=1}^N \text{Enc}(w_t^i),$$

in some homomorphic encryption frameworks where multiplication of ciphertext corresponds to the sum of plaintexts. The aggregator eventually provides the result back to the sites, which can decrypt collectively using threshold keys [44]. The aggregator does not learn the individual local updates. This guards against data leakage even if the aggregator is compromised, provided the cryptographic scheme remains unbroken.

Some federated frameworks combine differential privacy with homomorphic encryption [45]. The logic behind this dual strategy is:

$$\forall i \in \{1, \dots, N\}, \exists \text{Enc}_i : \text{Enc}_i \circ \text{DP}\left(w_t^i\right) \rightarrow \text{Enc}_i(\tilde{w}_t^i),$$

where $\tilde{w}_t^i$ is the perturbed or noised version of the model parameter. Even if an adversary obtains the ciphertext, the presence of differential privacy ensures that any attempt to invert the encryption to find the original data is hindered by statistical uncertainty [46]. The synergy between these methods offers a strong defense but introduces computational overhead that could be substantial for large models.

The complexities of implementing these privacy protocols in a clinical text environment are nontrivial [47]. Medical entity extraction often relies on large contextual embeddings that may produce high-dimensional parameter tensors. Encrypting or perturbing these tensors can magnify memory usage and slow computations. Nonetheless, the impetus for robust solutions is strong, given the potential for reidentification attacks [48]. In typical textual clinical data, identifiers can be a single token away from medical entities of interest, underscoring the importance of advanced privacy protections.

Both theoretical and empirical frameworks have been proposed to evaluate privacy-utility trade-offs [49]. One approach is to measure entity recognition F1 scores at various noise scales or encryption intensities. Another involves bounding the success rate of membership inference attacks, which aim to determine if a particular record is in the training set [50]. Mathematically, membership inference can be cast in terms of a hypothesis test on model outputs. Let r represent a record, and let θ be model

parameters. An adversary attempts to determine if r was in the training data by analyzing θ [51]. In scenarios with strong privacy guarantees, the advantage of the adversary over random guessing is minimized.

In broader theoretical terms, the combination of DP and encryption can be aligned with the structure: [52]

$$\phi = (\text{DPGuarantee}(\varepsilon, \delta)) \wedge (\text{HE_properties}),$$

where ϕ is the composite privacy requirement. This statement emphasizes that the system upholds differential privacy constraints for local updates while operating under homomorphic encryption for aggregated computations [53]. The local data remain inaccessible to other sites or external adversaries. The cost of these operations must be weighed against the practical demands of real-world healthcare workflows, which may be time-sensitive and resource-constrained.

Scaling these methods also demands strategies for partial or selective encryption of parameters [54]. For instance, in neural networks trained for entity recognition, lower-level embeddings might remain unencrypted if they contain less direct information about specific patients, while higher-level parameters closely tied to actual text patterns might undergo encryption. Logic-based policies can govern which segments of model updates to encrypt, ensuring comprehensive coverage for highly sensitive layers while reducing overhead for more general layers [55]. A possible formulation uses an indicator function $I(\ell \in S)$ that toggles encryption based on parameter layer ℓ belonging to a sensitive set S . Then the operation might be: [56]

$$\tilde{w}_{t,\ell}^i = \begin{cases} \text{Enc}(w_{t,\ell}^i), & \ell \in S, \\ w_{t,\ell}^i, & \ell \notin S. \end{cases}$$

Further refinements can incorporate approximate encryption methods or integer-based representations to reduce precision while maintaining adequate performance for named entity tasks.

In sum, the privacy-preserving dimension for federated learning in clinical text is shaped by the synergy of differential privacy, homomorphic encryption, secure multiparty computation, and practical constraints such as dimensional complexity and resource limitations. Quantifiable privacy provides confidence to both data owners and regulatory bodies [57]. End-to-end encryption strategies preclude unauthorized data exposures. Balancing these safeguards with performance is a decisive engineering challenge [58]. The next discussions move deeper into how domain-specific semantics for clinical named entity recognition can be integrated into this framework, underscoring the interactions between the language modeling components and the privacy-preserving mechanisms.

4. Semantic Representation for Clinical Named Entity Recognition

Contextual embedding models have revolutionized named entity recognition across a range of domains, and clinical text is no exception [59, 60]. Large-scale pretraining on general corpora is often insufficient, given the distinctive vocabulary and structure in medical narratives. Domain adaptation of contextual models, such as fine-tuning BERT-like architectures on clinical corpora, has become a standard practice to capture specialized entities relating to diagnoses, symptoms, procedures, and personal information. Let $E(\cdot)$ denote the embedding function that maps textual tokens to continuous vector representations [61]. Denote by $\{x_1, x_2, \dots, x_m\}$ a sentence from a clinical note. Then

$$\{v_1, v_2, \dots, v_m\} = E(\{x_1, x_2, \dots, x_m\}),$$

where $v_i \in \mathbb{R}^d$. The dimension d may be very large in specialized medical models. To identify named entities, a typical neural architecture processes the sequence of embeddings through contextual layers, often transformers, to capture dependencies [62]. A classification head then assigns labels to each token, distinguishing entity boundaries. Let y_i represent the label for token x_i , where y_i can be B_Entity, I_Entity, O or other tagging conventions.

In federated learning contexts, each clinical site might train a version of E augmented by local domain knowledge [63]. Let E_i be the locally adapted embedding function at site i , reflecting local data distributions, annotation styles, or subdomain specializations. The global model merges these representations in a privacy-preserving way. If the aggregator blindly averages the embedding layers, it might dilute specialized knowledge [64]. Alternatively, each site can train only higher layers, leaving a core embedding shared across participants. Denote the final classification layer by CL . Then the local model can be described as [65]

$$CL_i(E_i(x_1), E_i(x_2), \dots).$$

When aggregated, certain parts of CL_i may be combined across sites, while E_i remains partially local. This hierarchical strategy can reduce communication overhead, since only a subset of parameters is shared. Let θ_E be the parameters for the embedding function, and θ_{CL} for the classification layer. Federated updates can be partitioned:

$$\begin{aligned}\theta_{E,t+1} &= \theta_{E,t} - \eta \sum_{i=1}^N \nabla \mathcal{L}_{E,i}(\theta_{E,t}, \theta_{CL,t}), \\ \theta_{CL,t+1} &= \theta_{CL,t} - \eta \sum_{i=1}^N \nabla \mathcal{L}_{CL,i}(\theta_{E,t}, \theta_{CL,t}),\end{aligned}$$

where $\mathcal{L}_{E,i}$ and $\mathcal{L}_{CL,i}$ isolate the partial gradients relevant to each component. In practice, the aggregator might only request updates for θ_{CL} , and sites keep θ_E local or update it less frequently. This modular approach aligns with privacy objectives, as the low-level embeddings that might inadvertently encode personal details remain sequestered. [66]

Clinical text often includes additional metadata such as admission date, discharge date, or lab results embedded in note structures. These features can function as auxiliary signals for entity recognition, albeit with the risk of reidentification [67]. Incorporating them into the modeling pipeline demands a robust anonymization protocol. It is sometimes feasible to represent date or time intervals by relative offsets rather than absolute timestamps, thereby reducing the identifiability [68]. Let $\text{Offset}(t)$ be a function that normalizes a timestamp t by referencing it to an event or to a generic timeline. This lowers the resolution of potential personal data while preserving enough temporal context to assist in entity boundary detection. Such data transformation complements cryptographic approaches by reducing the direct presence of unique identifiers within the training process.

The classification architecture for entity detection in clinical text might be realized by conditional random fields or other structured prediction layers on top of the contextual embeddings [69]. Let $\text{CRF}(\{v_1, \dots, v_m\})$ be a function that models label dependencies. Federated training of a CRF layer can be expressed as merging the transition matrices from multiple sites. Let Ψ_i be the transition matrix for site i , capturing label-to-label transitions [70]. An aggregated transition matrix might be

$$\Psi = \sum_{i=1}^N \alpha_i \Psi_i,$$

where α_i depends on local dataset size or reliability weighting [71]. This parameter merging is performed in the encrypted or differentially private space, ensuring that local label statistics remain hidden. The logical premise

$$\forall i \in \{1, \dots, N\} : \Psi_i \in \mathcal{E} \longrightarrow \text{Enc}(\Psi_i),$$

forces each site to encrypt or perturb Ψ_i before it is shared [72]. The aggregator computes Ψ in ciphertext space or under noise constraints, mitigating the risk of reconstructing local label distributions.

Performance metrics for clinical named entity recognition typically include precision, recall, and F1 score for each class of entity [73]. These metrics can be computed locally by each site, or aggregated in a privacy-preserving manner. Global performance evaluation might rely on a hold-out set of anonymized data or synthetic data designed to approximate real clinical text [74]. Alternatively, partial model outputs can be encrypted and sent to a trusted evaluation server. Let Eval_i be the local evaluation function that returns metrics M for site i . A global aggregator can compute

$$M_{\text{global}} = \frac{1}{N} \sum_{i=1}^N \text{Eval}_i(\text{Enc}(\theta_{CL,i}), \text{Enc}(E_i)),$$

though in practice, localized evaluation is more straightforward [75]. Maintaining standardized annotation protocols across sites is essential for consistent named entity labeling, as subtle differences in labeling guidelines can degrade aggregated performance. This can be framed as a logical requirement: [76]

$$\exists \mathcal{G} : \forall i \in \{1, \dots, N\}, \text{AnnotationGuideline}(i) = \mathcal{G}.$$

Adherence to the same guidelines ensures that the aggregated model's notion of boundaries and entity classes remains coherent.

Through modular embedding strategies, data transformation policies, specialized structured prediction layers, and a federated training process, semantic representations tailored for clinical text can be learned while preserving patient confidentiality [77]. The success of these methods hinges on balanced data distributions, well-calibrated encryption or noise levels, and careful alignment of local adaptations with global model parameters. When properly orchestrated, federated approaches can unify diverse clinical corpora, capturing the domain intricacies required to identify medical entities accurately. This synergy of distributed optimization and specialized language modeling is the crux of advanced federated learning systems for privacy-preserving clinical named entity recognition. [78]

5. Implementation Details, Evaluation, and Theoretical Implications

The design of a privacy-preserving federated system for clinical named entity recognition combines cryptographic schemes, differential privacy methods, and advanced language modeling in a computationally feasible manner. One feasible path is to adopt a hybrid strategy, where partial encryption is employed for highly sensitive layers, and differential privacy noise is added to gradient updates in other layers [79]. This approach can reduce the overhead associated with fully homomorphic encryption while still achieving robust privacy assurances. Let Ω denote the set of model parameters that must be fully encrypted, and let Λ be the set of parameters subjected to noise-based DP [80]. Then the local update rule for site i might be expressed as

$$\forall \omega \in \Omega : \tilde{\omega}_t^i = \text{Enc}(\omega_t^i),$$

$$\forall \lambda \in \Lambda : \tilde{\lambda}_t^i = \lambda_t^i + \mathcal{N}(0, \sigma^2 \Delta^2).$$

The aggregator receives $\{\tilde{\omega}_t^i\}_{i=1}^N$ and $\{\tilde{\lambda}_t^i\}_{i=1}^N$, merging them into a new global parameter set.

Implementation typically relies on an underlying secure multiparty computation library or a specific homomorphic encryption backend that supports large-dimensional vectors. These systems rely on polynomial arithmetic in finite fields or ring structures for encryption [81]. Symbolically, one might denote the ring of polynomials modulo an irreducible polynomial by $\mathbb{Z}_p[x]/(x^m + 1)$. The encryption function $\text{Enc}(\cdot)$ maps real-valued vectors into polynomial rings via quantization, so each gradient or parameter is discretized. The aggregator then carries out addition or multiplication in the polynomial ring. Noise management is vital to prevent decryption failure or overflow [82]. A standard approach is to

keep gradient magnitudes bounded using gradient clipping, ensuring that the numerical range fits within the chosen integer representation. Let C denote the clipping threshold [83]. Each local site enforces:

$$g_t^i \leftarrow \frac{g_t^i}{\max(1, \frac{\|g_t^i\|_2}{C})}.$$

This operation normalizes outliers in the gradient space, ensuring stable encryption.

Experimental evaluation of such a system typically involves curated datasets or synthetic text that emulate real clinical notes [84]. One might utilize open clinical corpora like MIMIC-III, ensuring that data is split across multiple sites. The entity annotation can focus on critical categories, such as diagnoses, treatments, or personal identifiers [85]. To simulate realistic institutional distributions, each site might have a unique subset of the corpus reflecting local patient demographics. After a series of federated rounds, the final model is tested on a held-out subset. Common metrics include token-level F1 scores for entity detection [86]. Additionally, the privacy metrics revolve around quantifying membership inference attacks or calculating the (ϵ, δ) bounds under differential privacy. Let the final ϵ be computed as [87]

$$\epsilon_{\text{total}} = \epsilon_{\text{dp}} + \epsilon_{\text{composition}} + \epsilon_{\text{amplification}},$$

where ϵ_{dp} stems from the local DP noise, $\epsilon_{\text{composition}}$ arises from multiple training epochs, and $\epsilon_{\text{amplification}}$ may reflect the sampling ratio used in local batches.

The theoretical foundation of privacy in federated settings is grounded in bounding the amount of information each participant can glean about other participants' data. The cryptographic dimension ensures that aggregated computations do not leak local updates [88]. Differential privacy provides a bounding measure on the changes in outputs when individual records vary in the training set. Combining these approaches yields an overall system where an adversary controlling the aggregator or a subset of participants faces significant barriers to reconstructing sensitive data. Let Adv be an adversary, and let Π be the protocol controlling how encryption and noise addition are applied. A typical statement about security might read: [89]

$$\Pr(\text{Adv recovers } (x, y)) \leq \alpha(\epsilon, \delta, \text{cryptographic strength}),$$

where α is a function that decreases exponentially with the security parameter of the encryption scheme and the level of differential privacy. The aggregator is thus prevented from distinguishing whether a particular record is present or absent in the training corpus, and partial knowledge of intermediate parameters does not unravel local data. [90]

In practice, the system's performance is influenced by factors such as training hyperparameters, local data heterogeneity, and the level of encryption overhead. Heterogeneity in text styles and local annotation practices can slow convergence, requiring more communication rounds [91]. Domain-specific initialization, such as using a pretrained clinical language model, can mitigate these issues and accelerate the learning of robust entity representations. Through repeated testing, results often reveal a small drop in F1 scores compared to centralized training, reflecting the compromises needed to preserve privacy. Nevertheless, the domain coverage gained by federating distinct text corpora can lead to broader generalization capabilities. [92]

From a theoretical standpoint, certain bounds can be placed on the rate of convergence of federated optimization. In the simplest case, assume a strongly convex objective with Lipschitz continuous gradients [93]. Let $\nabla F(W)$ be the global gradient. Standard analyses show that under specific step-size rules, the deviation between federated updates and an optimal parameter set is bounded by terms that incorporate the variance of local gradients and the communication frequency [94]. When encryption or noise is introduced, additional variance arises, which can be incorporated into a generalized bound. Formally, if σ_{noise} is the magnitude of noise added per update, the asymptotic convergence rate might

degrade by a factor related to σ_{noise}^2 . Mathematically, one might encounter results of the form

$$\mathbb{E} \left[\|W_T - W^*\| \right] \leq O \left(\frac{\sigma_{\text{noise}}^2}{T} \right),$$

where W^* is the optimal set of parameters, and T is the number of rounds [95]. Similar logic applies to homomorphic encryption overhead, as polynomial arithmetic can introduce approximation errors. Despite these theoretical slowdowns, practical implementations often demonstrate acceptable trade-offs when the system is carefully tuned. [96]

This comprehensive view of implementation, evaluation, and theoretical implications illuminates the delicate interplay of advanced cryptographic methods, noise-based privacy, and domain-specific language modeling. Ultimately, the success of a federated approach for clinical named entity recognition hinges on balancing these elements to achieve meaningful performance gains while adhering to privacy mandates [97]. The final outcome is a scalable and secure pipeline that can be deployed across multiple clinical sites without risking the inadvertent disclosure of sensitive text segments. By encapsulating advanced neural architectures within a framework of robust privacy protections, one preserves both the integrity of patient data and the potential for collaborative medical breakthroughs.

6. Conclusion

Federated learning offers a potent solution for training sophisticated models on distributed clinical text data without centralizing patient records [98]. In this examination of privacy-preserving clinical named entity recognition, the architectural framework intertwines secure aggregation protocols, differential privacy methods, and domain-oriented embedding strategies to navigate the particular challenges of medical narratives. The logic-based descriptions and mathematical constructs introduced throughout the discussion serve as a foundation for understanding the constraints and trade-offs that arise when local gradients, encryption keys, and noise perturbations converge in a single system [99]. This paradigm not only respects legal and ethical imperatives but also potentially accelerates model development by tapping into a broad spectrum of clinical knowledge that would otherwise remain isolated within individual institutions.

The core mechanisms revolve around ensuring that neither individual tokens nor overall text patterns can be reverse-engineered from intermediate training signals [100]. Techniques such as homomorphic encryption, secure multiparty computation, and partial parameter sharing orchestrate the confidentiality of sensitive model layers. Simultaneously, differentially private updates provide statistical defenses against membership or property inference attacks. Considerations for efficient communication arise from the massive dimensionality of language model parameters, driving the need for compression, gradient clipping, and asynchronous aggregation methods [101]. Balancing these elements while maintaining semantic fidelity for specialized clinical entities is accomplished by modular architecture designs that partition encryption and noise-based modifications across different segments of the model.

Empirical evidence suggests that the incorporation of privacy mechanisms, while introducing certain overhead and minor performance degradation, can be carefully tuned to ensure that entity recognition metrics remain within clinically acceptable thresholds [102]. Moreover, theoretical constructs, ranging from convergence bounds to formal definitions of differential privacy, support the feasibility of scaling these methods to larger networks of hospitals or research institutions. The outcome is an infrastructure in which local adaptations to specialized medical jargon or site-specific data structures can be leveraged to construct a global model enriched by diverse textual corpora. [103]

The significance of these findings extends beyond the immediate scope of named entity recognition, offering a template that can be generalized to other tasks in healthcare analytics. The interplay of cryptographic primitives and advanced neural networks holds promise for radiology, pathology, and integrative multi-omics research, all of which share the need for confidentiality. Future directions may incorporate more refined domain-specific embedding strategies, automated data quality checks, and

dynamic protocols that adjust encryption or noise levels in real time based on the sensitivity of the processed content [104]. Such enhancements can further close the gap between the theoretical ideals of complete privacy and the practical demands of collaborative medical research.

In conclusion, federated learning for privacy-preserving clinical named entity recognition demonstrates a viable synergy of distributed optimization, cryptographic security, and domain-aware natural language processing [105]. This approach fosters the collective use of medical text across institutional boundaries while safeguarding individual patient identities. Through layered defenses, from homomorphic encryption to differential privacy, the risk of data leakage is mitigated, enabling the potential for large-scale, cooperative advancements in healthcare. The framework and analyses provided here illuminate a path forward in reconciling the needs for high-performing clinical text analytics and robust privacy protection, thereby contributing to an evolving landscape of secure, data-driven innovation in medicine. [106]

References

- [1] C. Tao, Y. Gong, H. Xu, and Z. Zhao, "Introduction: The international conference on intelligent biology and medicine (icibm) 2016: Special focus on medical informatics and big data," *BMC medical informatics and decision making*, vol. 17, pp. 77–77, 7 2017.
- [2] H. B. Fevrier, L. Liu, L. J. Herrinton, and D. Li, "A transparent and adaptable method to extract colonoscopy and pathology data using natural language processing.," *Journal of medical systems*, vol. 44, pp. 151–, 7 2020.
- [3] Y. Luo, W. K. Thompson, T. M. Herr, Z. Zeng, M. A. Berendsen, S. Jonnalagadda, M. B. Carson, and J. Starren, "Natural language processing for ehr-based pharmacovigilance: A structured review," *Drug safety*, vol. 40, pp. 1075–1089, 6 2017.
- [4] M. Waters, K. H. Kang, R. Brennenman, D. Caruthers, M. Spraker, and C. Abraham, "Convolutional neural networks naively trained on radiation oncology-specific data outperforms classical natural language processing approaches for automated identification of common toxicity terms.," *International Journal of Radiation Oncology*Biophysics*, vol. 111, no. 3, pp. S65–S66, 2021.
- [5] K. M. Zaki-Metias, J. J. MacLean, A. M. Satei, S. Medvedev, H. Wang, C. C. Zarour, and P. J. Arpasi, "The find program: Improving follow-up of incidental imaging findings.," *Journal of digital imaging*, vol. 36, pp. 804–811, 2 2023.
- [6] B. E. Jones, J. Ying, M. R. Nevers, P. R. Alba, O. V. Patterson, K. S. Peterson, E. Rutter, M. A. Christensen, S. Stern, M. M. Jones, A. Gundlapalli, N. C. Dean, M. C. Samore, and T. Greene, "Trends in illness severity, hospitalization, and mortality for community-onset pneumonia at 118 us veterans affairs medical centers.," *Journal of general internal medicine*, vol. 37, pp. 3839–3847, 3 2022.
- [7] A. Kline, H. Wang, Y. Li, S. Dennis, M. Hutch, Z. Xu, F. Wang, F. Cheng, and Y. Luo, "Multimodal machine learning in precision health: A scoping review.," *NPJ digital medicine*, vol. 5, pp. 171–, 11 2022.
- [8] B. Shiner, L. W. D'Avolio, T. M. Nguyen, M. H. Zayed, B. V. Watts, and L. D. Fiore, "Automated classification of psychotherapy note text: implications for quality assessment in ptsd care," *Journal of evaluation in clinical practice*, vol. 18, pp. 698–701, 2 2011.
- [9] J. Xu, B. S. Glicksberg, C. Su, P. B. Walker, J.-G. Bian, and F. Wang, "Federated learning for healthcare informatics," *Journal of healthcare informatics research*, vol. 5, pp. 1–19, 11 2020.
- [10] 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.
- [11] Q. Yin, L. Zhong, Y. Song, L. Bai, Z. Wang, C. Li, Y. Xu, and X. Yang, "A decision support system in precision medicine: contrastive multimodal learning for patient stratification," *Annals of Operations Research*, 8 2023.
- [12] B. D. Wissel, H. M. Greiner, T. A. Glauser, F. T. Mangano, D. Santel, J. Pestian, R. D. Szczesniak, and J. W. Dexheimer, "Investigation of bias in an epilepsy machine learning algorithm trained on physician notes.," *Epilepsia*, vol. 60, pp. e93–e98, 8 2019.
- [13] 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.

- [14] B. Ozaydin, F. D. Zengul, N. Oner, and D. Delen, "Text-mining analysis of mhealth research.," *mHealth*, vol. 3, pp. 53–53, 12 2017.
- [15] J. A. Macdonald, G. S. Roberts, P. A. Corrado, A. G. Beshish, K. Haraldsdottir, G. P. Barton, K. N. Goss, M. W. Eldridge, C. J. François, and O. Wieben, "Exercise-induced irregular right heart flow dynamics in adolescents and young adults born preterm," *Journal of cardiovascular magnetic resonance : official journal of the Society for Cardiovascular Magnetic Resonance*, vol. 23, pp. 116–, 10 2021.
- [16] S. J. Shah, R. Cogswell, J. J. Ryan, and K. Sharma, "How to develop and implement a specialized heart failure with preserved ejection fraction clinical program," *Current cardiology reports*, vol. 18, pp. 122–, 10 2016.
- [17] D. Cameron, B. Shiner, A. O'Neill, and M. O'Neil, "Factors associated with engaging in evidence-based psychotherapy during the first year of posttraumatic stress disorder treatment between 2017 and 2019.," *Administration and policy in mental health*, vol. 50, pp. 813–823, 6 2023.
- [18] J. Liu, J. Huang, Y. Zhou, X. Li, S. Ji, H. Xiong, and D. Dou, "From distributed machine learning to federated learning: a survey," *Knowledge and Information Systems*, vol. 64, pp. 885–917, 3 2022.
- [19] R. Avula, "Healthcare data pipeline architectures for ehr integration, clinical trials management, and real-time patient monitoring," *Quarterly Journal of Emerging Technologies and Innovations*, vol. 8, no. 3, pp. 119–131, 2023.
- [20] A. Raghuram, A. Ssentongo, L. M. Pereira, Y. Kuang, J. Uyei, and P. Park, "2378. use of natural language processing to extract published real world data on a covid vaccine and antiviral treatment," *Open Forum Infectious Diseases*, vol. 10, 11 2023.
- [21] S. Yu, K. P. Liao, S. Y. Shaw, V. S. Gainer, S. Churchill, P. Szolovits, S. N. Murphy, I. S. Kohane, and T. Cai, "Toward high-throughput phenotyping: unbiased automated feature extraction and selection from knowledge sources.," *Journal of the American Medical Informatics Association : JAMIA*, vol. 22, pp. 993–1000, 4 2015.
- [22] M. A. Wilcox, D. Coppola, N. Bailey, A. Wilson, A. W. C. Kamaau, P. R. Alba, O. V. Patterson, B. Viernes, D. W. Denhalter, I. D. Solomon, and S. L. DuVall, "Risperdal ® consta ® needle detachment. incidence rates before and after kit redesign: A retrospective study using electronic health records and natural language processing in the department of veterans affairs," *Neurology and therapy*, vol. 8, pp. 95–108, 3 2019.
- [23] 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.
- [24] M.-H. Huang and R. T. Rust, "A strategic framework for artificial intelligence in marketing," *Journal of the Academy of Marketing Science*, vol. 49, pp. 30–50, 11 2020.
- [25] S. Dublin, E. Baldwin, R. L. Walker, L. M. Christensen, P. J. Haug, M. L. Jackson, J. C. Nelson, J. P. Ferraro, D. Carrell, and W. W. Chapman, "Natural language processing to identify pneumonia from radiology reports.," *Pharmacoepidemiology and drug safety*, vol. 22, pp. 834–841, 4 2013.
- [26] J. M. Banda, A. Sarraju, F. Abbasi, J. Parizo, M. Pariani, H. Ison, E. Briskin, H. Wand, S. Dubois, K. Jung, S. A. Myers, D. J. Rader, J. B. Leader, M. F. Murray, K. D. Myers, K. Wilemon, N. H. Shah, and J. W. Knowles, "Finding missed cases of familial hypercholesterolemia in health systems using machine learning.," *NPJ digital medicine*, vol. 2, pp. 1–8, 4 2019.
- [27] D. Myers, R. Mohawesh, V. I. Chellaboina, A. L. Sathvik, P. Venkatesh, Y.-H. Ho, H. Henshaw, M. Alhawawreh, D. Berdik, and Y. Jararweh, "Foundation and large language models: fundamentals, challenges, opportunities, and social impacts," *Cluster Computing*, vol. 27, pp. 1–26, 11 2023.
- [28] V. D. Badal, C. Nebeker, K. Shinkawa, Y. Yamada, K. E. Rentscher, H.-C. Kim, and E. E. Lee, "Do words matter? detecting social isolation and loneliness in older adults using natural language processing," *Frontiers in psychiatry*, vol. 12, pp. 728732–, 11 2021.
- [29] 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.
- [30] A. Belouali, S. Gupta, V. Sourirajan, J. Yu, N. Allen, A. Alaoui, M. A. Dutton, and M. J. Reinhard, "Acoustic and language analysis of speech for suicidal ideation among us veterans.," *BioData mining*, vol. 14, pp. 11–11, 2 2021.

- [31] K. Syed, W. C. Sleeman, K. Ivey, M. P. Hagan, J. Palta, R. Kapoor, and P. Ghosh, "Integrated natural language processing and machine learning models for standardizing radiotherapy structure names.," *Healthcare (Basel, Switzerland)*, vol. 8, pp. 120–, 4 2020.
- [32] A. J. Gawron, Y. Yao, S. Gupta, G. Cole, M. A. Whooley, J. A. Dominitz, and T. Kaltenbach, "Simplifying measurement of adenoma detection rates for colonoscopy," *Digestive diseases and sciences*, vol. 66, pp. 3149–3155, 10 2020.
- [33] D. M. Hilty, C. M. Armstrong, A. Edwards-Stewart, M. T. Gentry, D. D. Luxton, and E. A. Krupinski, "Sensor, wearable, and remote patient monitoring competencies for clinical care and training: Scoping review," *Journal of technology in behavioral science*, vol. 6, pp. 1–26, 1 2021.
- [34] J. R. Rogers, T. J. Callahan, T. Kang, A. Bauck, R. Khare, J. S. Brown, M. G. Kahn, and C. Weng, "A data element-function conceptual model for data quality checks," *EGEMS (Washington, DC)*, vol. 7, pp. 17–17, 4 2019.
- [35] P. R. Alba, J. Lynch, A. Gao, K. M. Lee, T. Anglin-Foote, B. Robison, J. B. Shelton, O. Efimova, O. V. Patterson, and S. L. DuVall, "Using natural language processing (nlp) tools to identify veterans with metastatic prostate cancer (mpca).," *Journal of Clinical Oncology*, vol. 38, pp. 60–60, 2 2020.
- [36] A. Sharma and K. D. Forbus, "Modeling the evolution of knowledge and reasoning in learning systems," in *2010 AAAI Fall Symposium Series*, 2010.
- [37] R. Avula *et al.*, "Data-driven decision-making in healthcare through advanced data mining techniques: A survey on applications and limitations," *International Journal of Applied Machine Learning and Computational Intelligence*, vol. 12, no. 4, pp. 64–85, 2022.
- [38] B. Sharma, D. Dligach, K. Swope, E. Salisbury-Afshar, N. S. Karnik, C. Joyce, and M. Afshar, "Publicly available machine learning models for identifying opioid misuse from the clinical notes of hospitalized patients," *BMC medical informatics and decision making*, vol. 20, pp. 79–, 4 2020.
- [39] B. Falissard, E. L. Simpson, E. Guttman-Yassky, K. A. Papp, S. Barbarot, A. Gadkari, G. Saba, L. Gautier, A. Abbe, and L. Eckert, "Qualitative assessment of adult patients' perception of atopic dermatitis using natural language processing analysis in a cross-sectional study," *Dermatology and Therapy*, vol. 10, pp. 297–305, 1 2020.
- [40] J. Chen, X. Li, B. J. Aguilar, E. Shishova, P. J. Morin, D. Berlowitz, D. R. Miller, M. K. O'Connor, A. H. Nguyen, R. Zhang, A. A. T. Monfared, Q. Zhang, and W. Xia, "Development and validation of a natural language processing system that extracts cognitive test results from clinical notes," *Alzheimer's & Dementia*, vol. 19, 12 2023.
- [41] A. Swaminathan, I. López, R. A. G. Mar, T. Heist, T. McClintock, K. Caoili, M. Grace, M. Rubashkin, M. N. Boggs, J. H. Chen, O. Gevaert, D. Mou, and M. K. Nock, "Natural language processing system for rapid detection and intervention of mental health crisis chat messages.," *NPJ digital medicine*, vol. 6, pp. 213–, 11 2023.
- [42] J. H. Garvin, Y. Kim, G. T. Gobbel, M. E. Matheny, A. Redd, B. E. Bray, P. A. Heidenreich, D. Bolton, J. Heavirland, N. Kelly, R. M. Reeves, M. Kalsy, M. K. Goldstein, and S. M. Meystre, "Automating quality measures for heart failure using natural language processing: A descriptive study in the department of veterans affairs," *JMIR medical informatics*, vol. 6, pp. e5–, 1 2018.
- [43] L. G. Liu, R. H. Grossman, E. G. Mitchell, C. Weng, K. Natarajan, G. Hripcsak, and D. K. Vawdrey, "A deep database of medical abbreviations and acronyms for natural language processing.," *Scientific data*, vol. 8, pp. 149–149, 6 2021.
- [44] E. L. Palmer, S. Hassanpour, J. H. Higgins, J. A. Doherty, and T. Onega, "Building a tobacco user registry by extracting multiple smoking behaviors from clinical notes.," *BMC medical informatics and decision making*, vol. 19, pp. 1–10, 7 2019.
- [45] H. Jain, N. Raj, and S. Mishra, "A sui generis qa approach using roberta for adverse drug event identification.," *BMC bioinformatics*, vol. 22, pp. 330–, 10 2021.
- [46] C. Xiang and M. Abouelyazid, "The impact of generational cohorts and visit environment on telemedicine satisfaction: A novel investigation," 2020.
- [47] V. Castro, Y. Shen, S.-C. Yu, S. Finan, C. T. Pau, V. S. Gainer, C. C. Keefe, G. Savova, S. N. Murphy, T. Cai, and C. K. Welt, "Identification of subjects with polycystic ovary syndrome using electronic health records," *Reproductive biology and endocrinology : RB&E*, vol. 13, pp. 116–116, 10 2015.
- [48] J. Valdez, M. Rueschman, M. Kim, S. Redline, and S. S. Sahoo, "Otm conferences - an ontology-enabled natural language processing pipeline for provenance metadata extraction from biomedical text (short paper).," *On the move to meaningful Internet systems ... : CoopIS, DOA, and ODBASE : Confederated International Conferences, CoopIS, DOA, and ODBASE ... proceedings. OTM Confederated International Conferences*, vol. 10033, pp. 699–708, 10 2016.

- [49] A. Fong, "Realizing the power of text mining and natural language processing for analyzing patient safety event narratives: The challenges and path forward.," *Journal of patient safety*, vol. 17, pp. e834–e836, 12 2021.
- [50] A. J. Afzal, A. Srour, A. Goil, S. Vasudaven, T. Liu, R. Samudrala, N. Dogra, P. Kohli, A. Malakar, and D. A. Lightfoot, "Homo-dimerization and ligand binding by the leucine-rich repeat domain at rhg1/rfs2 underlying resistance to two soybean pathogens," *BMC plant biology*, vol. 13, pp. 43–43, 3 2013.
- [51] J. Sun and P. A. Gloor, "“towards re-inventing psychohistory”: Predicting the popularity of tomorrow’s news from yesterday’s twitter and news feeds," *Journal of systems science and systems engineering*, vol. 30, pp. 1–20, 11 2020.
- [52] 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.
- [53] S. Wang, L. Liu, J. Liu, L. Miao, Q. Zhuang, N. Guo, J. Zhao, Q. Li, and G. Ren, "Characteristics of prescriptions and costs for acute upper respiratory tract infections in chinese outpatient pediatric patients: a nationwide cross-sectional study.," *BMC complementary medicine and therapies*, vol. 20, pp. 1–10, 11 2020.
- [54] R. J. Desai, M. E. Matheny, K. Johnson, K. Marsolo, L. H. Curtis, J. C. Nelson, P. J. Heagerty, J. Maro, J. Brown, S. Toh, M. Nguyen, R. Ball, G. D. Pan, S. V. Wang, J. J. Gagne, and S. Schneeweiss, "Broadening the reach of the fda sentinel system: A roadmap for integrating electronic health record data in a causal analysis framework.," *NPJ digital medicine*, vol. 4, pp. 170–, 12 2021.
- [55] B. Shiner, M. Levis, V. M. Dufort, O. V. Patterson, B. V. Watts, S. L. DuVall, C. J. Russ, and S. Maguen, "Improvements to ptsd quality metrics with natural language processing.," *Journal of evaluation in clinical practice*, vol. 28, pp. 520–530, 5 2021.
- [56] D. Su, Q. Li, T. Zhang, P. Veliz, Y. Chen, K. He, P. Mahajan, and X. Zhang, "Prediction of acute appendicitis among patients with undifferentiated abdominal pain at emergency department.," *BMC medical research methodology*, vol. 22, pp. 18–, 1 2022.
- [57] A. Sharma and K. Forbus, "Modeling the evolution of knowledge in learning systems," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 26, pp. 669–675, 2012.
- [58] 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.
- [59] W. A. Zaghloul and S. Trimi, "Developing an innovative entity extraction method for unstructured data," *International Journal of Quality Innovation*, vol. 3, pp. 1–10, 5 2017.
- [60] 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.
- [61] C. Ramprasad, E. N. Uche-Anyia, and T. M. Berzin, "Artificial intelligence in colorectal cancer screening," *Current Treatment Options in Gastroenterology*, vol. 21, pp. 272–282, 6 2023.
- [62] J. S. Aberdeen, S. Bayer, C. Clark, M. Keybl, and D. Tresner-Kirsch, "An annotation and modeling schema for prescription regimens," *Journal of biomedical semantics*, vol. 10, pp. 1–11, 5 2019.
- [63] F. Liu, A. Jagannatha, and H. Yu, "Towards drug safety surveillance and pharmacovigilance: Current progress in detecting medication and adverse drug events from electronic health records.," *Drug safety*, vol. 42, pp. 95–97, 1 2019.
- [64] M. Song, J. A. O'Donnell, T. Bekhuis, and H. Spallek, "Are dentists interested in the oral-systemic disease connection? a qualitative study of an online community of 450 practitioners," *BMC oral health*, vol. 13, pp. 65–65, 11 2013.
- [65] B. V. Udelsman, E. J. Lilley, M. Qadan, D. C. Chang, K. D. Lillemoe, C. Lindvall, and Z. Cooper, "Deficits in the palliative care process measures in patients with advanced pancreatic cancer undergoing operative and invasive nonoperative palliative procedures," *Annals of surgical oncology*, vol. 26, pp. 4204–4212, 8 2019.
- [66] R. J. Jackson, I. E. Kartoglu, C. Stringer, G. Gorrell, A. Roberts, X. Song, H. Wu, A. Agrawal, K. Lui, T. Groza, D. Lewsley, D. Northwood, A. Folarin, R. Stewart, and R. Dobson, "Cogstack - experiences of deploying integrated information retrieval and extraction services in a large national health service foundation trust hospital," *BMC medical informatics and decision making*, vol. 18, pp. 47–47, 6 2018.
- [67] Y. Shao, Y. Cheng, R. U. Shah, C. R. Weir, B. E. Bray, and Q. Zeng-Treitler, "Shedding light on the black box: Explaining deep neural network prediction of clinical outcomes.," *Journal of medical systems*, vol. 45, pp. 5–, 1 2021.

- [68] X. Zhao and A. Rios, "A marker-based neural network system for extracting social determinants of health.," *Journal of the American Medical Informatics Association : JAMIA*, vol. 30, pp. 1398–1407, 4 2023.
- [69] M. Singh, "Public sentiment and opinion regarding the cares act.," *Business economics (Cleveland, Ohio)*, vol. 58, pp. 24–33, 12 2022.
- [70] A. S. Eldaly, F. R. Avila, R. A. Torres-Guzman, K. Maita, J. P. Garcia, L. P. Serrano, and A. J. Forte, "Simulation and artificial intelligence in rhinoplasty: A systematic review.," *Aesthetic plastic surgery*, vol. 46, pp. 2368–2377, 4 2022.
- [71] D. Goodman-Meza, A. Tang, B. Aryanfar, S. Vazquez, A. J. Gordon, M. Goto, M. B. Goetz, S. Shoptaw, and A. A. T. Bui, "Natural language processing and machine learning to identify people who inject drugs in electronic health records.," *Open forum infectious diseases*, vol. 9, pp. ofac471–, 9 2022.
- [72] D. R. Newman-Griffis, M. B. Hurwitz, G. P. McKernan, A. J. Houtrow, and B. E. Dicianno, "A roadmap to reduce information inequities in disability with digital health and natural language processing.," *PLOS digital health*, vol. 1, pp. e0000135–e0000135, 11 2022.
- [73] P. Mathur, R. Awasthi, K. Maheshwari, J. Cywinski, V. Sabharwal, F. Papay, A. Khanna, and S. Mishra, "302: Quantitative and qualitative evaluation of artificial intelligence-based critical care publications," *Critical Care Medicine*, vol. 51, pp. 137–137, 12 2022.
- [74] P. R. Dexter, J. He, J. Baker, G. J. Eckert, A. Church, and N. J. Zhang, "Extraction of breast cancer biomarker status using natural language processing," *International Journal of Computational Medicine and Healthcare*, vol. 1, no. 1, pp. 112–, 2019.
- [75] Z. Chen, H. Zhang, Y. Guo, T. J. George, M. Prosperi, W. R. Hogan, Z. He, E. Shenkman, F. Wang, and J.-G. Bian, "Exploring the feasibility of using real-world data from a large clinical data research network to simulate clinical trials of alzheimer's disease.," *NPI digital medicine*, vol. 4, pp. 84–84, 5 2021.
- [76] O. F. Hunter, F. Perry, M. Salehi, H. Bandurski, A. Hubbard, C. G. Ball, and S. M. Hameed, "Science fiction or clinical reality: a review of the applications of artificial intelligence along the continuum of trauma care.," *World journal of emergency surgery : WJES*, vol. 18, pp. 16–, 3 2023.
- [77] K. P. Liao, T. Cai, V. S. Gainer, S. Goryachev, Q. Zeng-Treitler, S. Raychaudhuri, P. Szolovits, S. Churchill, S. N. Murphy, I. S. Kohane, E. W. Karlson, and R. M. Plenge, "Electronic medical records for discovery research in rheumatoid arthritis," *Arthritis care & research*, vol. 62, pp. 1120–1127, 8 2010.
- [78] V. Gamerman, T. Cai, and A. Elsässer, "Pragmatic randomized clinical trials: best practices and statistical guidance," *Health Services and Outcomes Research Methodology*, vol. 19, pp. 23–35, 12 2018.
- [79] Y. Li, P. B. Ryan, Y. Wei, and C. Friedman, "A method to combine signals from spontaneous reporting systems and observational healthcare data to detect adverse drug reactions," *Drug safety*, vol. 38, pp. 895–908, 7 2015.
- [80] L. A. Knake, M. Ahuja, E. L. McDonald, K. K. Ryckman, N. Weathers, T. Burstain, J. M. Dagle, J. C. Murray, and P. Nadkarni, "Quality of ehr data extractions for studies of preterm birth in a tertiary care center: guidelines for obtaining reliable data," *BMC pediatrics*, vol. 16, pp. 59–59, 4 2016.
- [81] W. Yu, C. Zheng, F. Xie, W. Chen, C. Mercado, L. S. Sy, L. Qian, S. C. Glenn, H. F. Tseng, G. S. Lee, J. Duffy, M. M. McNeil, M. F. Daley, B. Crane, H. Q. McLean, L. A. Jackson, and S. J. Jacobsen, "The use of natural language processing to identify vaccine-related anaphylaxis at five health care systems in the vaccine safety datalink," *Pharmacoepidemiology and drug safety*, vol. 29, pp. 182–188, 12 2019.
- [82] H. M. Thompson, J. A. Coleman, and P. Kent, "Lgbt medical education: First-year medical students' self-assessed knowledge and comfort with transgender and lgb populations," *Medical Science Educator*, vol. 28, pp. 693–697, 8 2018.
- [83] G. Schmajuk and J. Yazdany, "Leveraging the electronic health record to improve quality and safety in rheumatology.," *Rheumatology international*, vol. 37, pp. 1603–1610, 8 2017.
- [84] T. V. Vleck, L. Chan, S. G. Coca, C. K. Craven, R. Do, S. B. Ellis, J. L. Kannry, R. J. F. Loos, P. A. Bonis, J. H. Cho, and G. N. Nadkarni, "Augmented intelligence with natural language processing applied to electronic health records for identifying patients with non-alcoholic fatty liver disease at risk for disease progression.," *International journal of medical informatics*, vol. 129, pp. 334–341, 7 2019.
- [85] 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.

- [86] A. A. Thomas, C. Zheng, H. Jung, A. Chang, B. Kim, J. Gelfond, J. Slezak, K. Porter, S. J. Jacobsen, and G. W. Chien, "Extracting data from electronic medical records: validation of a natural language processing program to assess prostate biopsy results," *World journal of urology*, vol. 32, pp. 99–103, 2 2013.
- [87] J. Byamugisha, W. Saib, T. Gaelejew, A. Jeewa, and M. Molapo, "Abstract pr-12: Towards verifying results from biomedical deep learning models using the umls: Cases of primary tumor site classification and cancer named entity recognition," *Clinical Cancer Research*, vol. 27, pp. PR–12–PR–12, 3 2021.
- [88] B. Hazlehurst, J. M. Lawrence, W. T. Donahoo, N. E. Sherwood, S. E. Kurtz, S. Xu, and J. F. Steiner, "Automating assessment of lifestyle counseling in electronic health records," *American journal of preventive medicine*, vol. 46, no. 5, pp. 457–464, 2014.
- [89] C. A. Bejan, M. Ripperger, D. Wilimitis, R. Ahmed, J. Kang, K. Robinson, T. J. Morley, D. M. Ruderfer, and C. G. Walsh, "Improving ascertainment of suicidal ideation and suicide attempt with natural language processing.," *Scientific reports*, vol. 12, pp. 15146–, 9 2022.
- [90] R. Avula, "Architectural frameworks for big data analytics in patient-centric healthcare systems: Opportunities, challenges, and limitations," *Emerging Trends in Machine Intelligence and Big Data*, vol. 10, no. 3, pp. 13–27, 2018.
- [91] G. W. Cannon, S. Mehrotra, and B. C. Sauer, "Ab1089 use of natural language processing to enhance retrieval of rheumatoid arthritis disease activity outcomes measures in us veterans," *Annals of the Rheumatic Diseases*, vol. 76, pp. 1435–1436, 2017.
- [92] S. M. Meystre and P. J. Haug, "Automation of a problem list using natural language processing.," *BMC medical informatics and decision making*, vol. 5, pp. 30–30, 8 2005.
- [93] L. Han, S. L. Luther, D. K. Finch, S. K. Dobscha, M. Skanderson, H. Bathulapalli, S. J. Fodeh, B. Hahm, L. Bouayad, A. Lee, J. L. Goulet, C. A. Brandt, and R. D. Kerns, "Complementary and integrative health approaches and pain care quality in the veterans health administration primary care setting: A quasi-experimental analysis.," *Journal of integrative and complementary medicine*, vol. 29, pp. 420–429, 3 2023.
- [94] A. Bompelli, Y. Wang, R. Wan, E. Singh, Y. Zhou, L. Xu, D. Oniani, B. S. A. Kshatriya, J. E. Balls-Berry, and R. Zhang, "Social and behavioral determinants of health in the era of artificial intelligence with electronic health records: A scoping review," *Health data science*, vol. 2021, pp. 1–19, 8 2021.
- [95] E. Hossain, R. Rana, N. Higgins, J. Soar, P. D. Barua, A. R. Pisani, and K. Turner, "Natural language processing in electronic health records in relation to healthcare decision-making: A systematic review.," *Computers in biology and medicine*, vol. 155, pp. 106649–106649, 2 2023.
- [96] V. D. Badal, S. A. Graham, C. A. Depp, K. Shinkawa, Y. Yamada, L. A. Palinkas, H.-C. Kim, D. V. Jeste, and E. E. Lee, "Prediction of loneliness in older adults using natural language processing: Exploring sex differences in speech.," *The American journal of geriatric psychiatry : official journal of the American Association for Geriatric Psychiatry*, vol. 29, pp. 853–866, 9 2020.
- [97] J. A. McCart, D. Finch, J. Jarman, E. J. Hickling, J. D. Lind, M. Richardson, D. J. Berndt, and S. L. Luther, "Using ensemble models to classify the sentiment expressed in suicide notes," *Biomedical informatics insights*, vol. 5, pp. 77–85, 1 2012.
- [98] P. J. Freda, H. R. Kranzler, and J. H. Moore, "Novel digital approaches to the assessment of problematic opioid use.," *BioData mining*, vol. 15, pp. 14–, 7 2022.
- [99] J. Zhang, A. Can, P. M. R. Lai, S. Mukundan, V. M. Castro, D. Dligach, S. Finan, V. S. Gainer, N. A. Shadick, G. Savova, S. N. Murphy, T. Cai, S. T. Weiss, and R. Du, "Morphological variables associated with ruptured basilar tip aneurysms," *Scientific reports*, vol. 11, pp. 2526–2526, 1 2021.
- [100] M. Chernyshev, S. Zeadally, and Z. A. Baig, "Healthcare data breaches: Implications for digital forensic readiness," *Journal of medical systems*, vol. 43, pp. 7–, 11 2018.
- [101] A. Sharma, *Structural and network-based methods for knowledge-based systems*. PhD thesis, Northwestern University, 2011.
- [102] E. A. Wang, J. B. Long, K. A. McGinnis, K. H. Wang, C. Wildeman, C. H. Kim, K. B. Bucklen, D. A. Fiellin, J. Bates, C. Brandt, and A. C. Justice, "Measuring exposure to incarceration using the electronic health record.," *Medical care*, vol. 57, no. Suppl 2, pp. S157–S163, 2019.

- [103] M. Muniswamaiah, T. Agerwala, and C. Tappert, "Data virtualization for analytics and business intelligence in big data," in *CS & IT Conference Proceedings*, vol. 9, CS & IT Conference Proceedings, 2019.
- [104] Z. Nabulsi, A. Selligren, S. Jamshy, C. Lau, E. Santos, A. P. Kiraly, W. Ye, J. Yang, R. Pilgrim, S. Kazemzadeh, J. Yu, S. R. Kalidindi, M. Etemadi, F. Garcia-Vicente, D. S. Melnick, G. S. Corrado, L. Peng, K. Eswaran, D. Tse, N. Beladia, Y. Liu, P.-H. C. Chen, and S. Shetty, "Deep learning for distinguishing normal versus abnormal chest radiographs and generalization to two unseen diseases tuberculosis and covid-19.," *Scientific reports*, vol. 11, pp. 15523–15523, 9 2021.
- [105] K. Chakrapani, M. Kempanna, M. I. Safa, T. Kavitha, M. Ramachandran, V. Bhaskar, and A. Kumar, "An enhanced exploration of sentimental analysis in health care," *Wireless Personal Communications*, vol. 128, pp. 901–922, 10 2022.
- [106] C. Normann, S. Vestring, K. Domschke, J. Bischofberger, and T. Serchov, "Acnp 61st annual meeting: Poster abstracts p271-p540.," *Neuropsychopharmacology : official publication of the American College of Neuropsychopharmacology*, vol. 47, pp. 220–370, 12 2022.